Confidence Reports

Fabrizio Cariani, Paolo Santorio, Alexis Wellwood

Word count: 8,443 (excluding footnotes)

Abstract

We develop a states-based semantics for nominal and adjectival confidence reports like *Ann is confident/has confidence that it's raining*, and their comparative forms. Our account leverages a neodavidsonian analysis of adjectival comparatives in which adjectives denote properties of states and measure functions are introduced compositionally. We hereby provide the first systematic semantics for confidence reports, in addition to providing a needed modal extension to the states-based semantics of comparatives. As we show, the flexibility accorded by the neodavidsonian implementation supports analysis of grammatical constructions with *confident/confidence* that might otherwise be puzzling, and it lends itself to certain natural ideas about the semantics of cross-categorial probabilistic language using e.g. *likely* and *probability*. In the end, we sketch some immediate connections 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

Our goal in this paper is to provide an analysis of confidence reports, i.e. sentences like (1a). A reasonable starting assumption is that these express a confidence relation between an attitude holder and a proposition. Importantly, these attitudes are gradable, as is evidenced by (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 on propositional attitudes, the semantics of confidence reports is under-theorized. This might be because semanticists and philosophers expect that we can simply apply an off-the shelf scalar semantics for gradable adjectives to *confident*. Indeed, in recent years a number of theorists have worked on extending standard degree semantics frameworks to modal adjectives like *likely*.[1] Offhand, we might try to apply an analysis of this sort to attitude ad-

[1] See e.g. Yalcin 2010, Lassiter 2011, 2016.

jectives like *confident*, with perhaps minor adjustments.

A first worry about such a strategy is that it would be insufficiently general. Standard degree-based frameworks cannot be easily generalized to comparative constructions targeting nominals. This means that an off-the-shelf degree-theoretic account of (1) will fail to extend neatly to sentences like (2a)-(2b) involving the nominal *confidence*.

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

Fortunately, the literature includes a proposal to support a unified treatment of adjectival and nominal confidence reports, namely Wellwood's (2014; 2015) treatment of gradability and comparison. She interprets gradable adjectives as properties of states in a broadly Davidsonian setting. This proposal supports the unification of nominal and adjectival comparison, but modal gradable adjectives have not been analyzed in such a framework. Their treatment requires accounting for distinctive complexities, some familiar from the extensive literature on attitudes and modal semantics, and others particular to the Davidsonian interpretation.

A second worry is that the simple, off-the-shelf strategy requires the well-known stipulation of a convert morpheme POS (or something like it) wherever a gradable adjective occurs without overt degree morphology. This is compositionally necessary on the standard degree semantics, but there is little independent evidence to support its existence (see Grano and Davis 2018 for recent discussion and references). Moreover, considering the semantic relationships between gradable adjectives and their nominal counterparts as in (1)-(2), implementing the simple strategy would mean extending this stipulation to the nominal domain as well. We propose instead to integrate the analysis of confidence reports within the POS-free framework independently developed in Cariani et al. (forthcoming).

So we are in the following predicament. On the one hand, there are developed analyses of modal gradable adjectives in degree semantics, but these analyses don't generalize easily to the nominal case and they rely on stipulating POS. On the other hand, there is a unified analysis of adjectival and nominal comparison, but as yet no rigorous attempt has been made to bring modal expressions within its purview. Our goal in this paper is precisely to address this gap. We develop a unified analysis of adjectival and nominal comparison that applies seamlessly to the modal case. Confidence reports are an excellent case study for this integration, given that their occurrence in nominal form is exceedingly commonplace.

A third motivating consideration fueling our proposal is internal to event-based projects in semantics. While there are contributions that aim to integrate neodavidsonian event semantics with standard possible world semantics, there is comparatively little of this work with respect to graded modality and graded propositional attitudes. By comparison, graded modalities and, to a lesser extent, gradable attitudes are well studied within the classical degree framework (e.g. Lassiter 2011, 2016, Pasternak 2019). Given the necessity of an account along these lines, our efforts here profile as the first foray into this territory within the neodavidsonian paradigm.

Here is a brief summary of our proposal. We develop a semantics that interprets *confident* as a property of states, measures of which are introduced compositionally

in constructions like (1b) (Wellwood 2014, 2015). Using states supports a degree of semantic flexibility that we exploit in our analysis of both positive and comparative occurrences of *confident* and *confidence*. According to our proposal, the positive forms of *confident* and *confidence* involve reference to an initial ordering on states introduced by the lexical items (i.e., the "background ordering" of Cariani et al. forthcoming), while the comparative form involves reference to degrees introduced by *more* (i.e., a more familiar "degree ordering").

The resulting semantics directly accomplishes the primary goal for which it was designed: it provides a unified semantics for *confident* and *confidence* across their positive and comparative occurrences. And it does so without recourse to anything like a covert POS morpheme, all while affording a flexibility of composition that supports analysis of other varieties of confidence reports (e.g. *Ann is confident in Bill*, *Ann is a confident person*) that may elude an off-the-shelf scalar semantics. It additionally allows for modeling of confidence ascriptions to agents who fall short of perfect rationality; this is unlikely to be straightforward if we were to simply embed the scalar analysis within recent probabilistic frameworks.

As we flag in the final pages, our states-based analysis can in principle be extended to modal adjectives like *likely* and *probable*, e.g. (3). Differently from existing accounts, this analysis would also generalize seamlessly to nominal probability claims, such as (4).

(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.

Relatedly, our analysis can capture inferential relations between confidence reports and reports of probabilistic belief. Here, the target of explanation is the intuition that a confidence report like (1b) sounds roughly equivalent to *Ann thinks that it's more likely to rain than to snow*.

We proceed as follows. We motivate the demand for a unified semantics for *confident* and *confidence* as against an off-the-shelf scalar strategy in §2. We lay out our compositional states-based analysis in §3, and explain the semantics of our target cases and other varieties of confidence reports in §4. Finally, §5 briefly discusses extensions of our approach beyond the domain of confidence reports, including conditionals, attitudinal scale-mates (e.g. *sure*, *doubtful*) and probability operators (e.g. *likely*, *chance*).

2 Confidence reports in a scalar setting

We consider reports with adjectival *confident* and nominal *confidence*, and their interpretation in the positive and comparative forms. A central assumption of our project is that it is desirable to pursue a unified semantics for these reports. This cannot be accommodated on the standard scalar semantics, since it assigns different compositional roles to adjectives and nouns in degree constructions: on classic accounts like those exemplified by Heim 1985, 2000 and Kennedy 1999, gradable

adjectives lexically introduce/express measure functions, but that semantic role is played by a functional item like *much* or *many* with nouns.[2]

By itself, the fact that we have an adjective and a noun that are obviously morphologically related doesn't settle the question of whether a unified semantics is appropriate. Nevertheless, we think that there are strong compositional and logical reasons to pursue unification.

One compositional reason is that the very same comparative morphology applies to both the nominal and adjectival forms. As illustrated by (5), repeated from above, the same form, *more*, is used in comparative confidence reports regardless of the lexical category targeted.

(5) a. Ann is more confident that it will rain than that it will snow.
b. Ann has more confidence that it will rain than that it will snow.

The standard scalar analysis will assign *confident* and *confidence* different semantic types, and so it must postulate systematic ambiguity in the meaning of *more*. A unified account opens up the prospect of a single meaning for *more*.

And there are logical reasons. The inferential connections between nominal and adjectival confidence reports are visible already in the positive form: plausibly, the adjectival *Ann is confident that it's raining* entails the nominal *Ann has confidence that it's raining*. (Whether the entailment runs in the opposite direction is, as we will highlight, a more complex matter.) But the connection is even clearer in the comparative form: the pair in (6a) and (6b) are straightforwardly equivalent.

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

Such reasons strongly suggest to us the desirability of a unified account. Pursuing that goal will, however, steer us away from the standard framework commonly used to model gradability in the modal domain, and towards a framework based on states. It will be important for that move, then, to understand how the standard framework would apply in this case. A typical move is to begin with the adjectival case, and to treat the nominal case separately, if at all. Let us briefly tell a plausible version of such a story.

Applying the classic degree-theoretic account to an adjective first involves drawing attention to the available distributional evidence for gradability. This is, of course, found in abundance for adjectival *confident*: in addition to its comfortable occurrence in the comparative, (1b), it sensibly combines with the full panoply of comparative forms (e.g., *as confident*, *too confident*, *confident enough*, etc.), and with modifiers like *very* and *100%*. Such evidence is typically thought sufficient to posit that the relevant adjective directly lexicalizes a degree semantics.

Assume for now the version of the standard framework which assigns gradable adjectives like *tall* a measure function type (e.g. Kennedy 1999), i.e., the type that maps individuals to degrees.[3] In this framework, *confident* would be assigned the

[2]This item may play its role covertly, given a surface form like *more*, or it may do so overtly, as in surface forms like *as much/many*, etc.

[3]The prominent alternative in the style of Heim 1985, 2000 analyzes them as having a higher quantificational type which embeds a measure function. Everything we say in the text with the simpler types will apply equally well to this alternative, *mutatis mutandis*.

interpretation in (7), where **confident** maps a proposition p to x's degree of confidence in the truth of that proposition, and g is a variable assignment.[4] 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.

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

For occurrences of the adjective without overt degree morphology, this framework typically assumes the presence of a covert morpheme, usually indicated by 'POS'.[5] While the details vary, the general function of this morpheme is to introduce comparison to a contextually-salient "standard" degree. Given these assumptions, the LF of (1a) has the additional structure indicated in (8a), and truth conditions like those in (8b). This says that Ann's confidence in the truth of the proposition expressed by *it's raining* is greater than the standard for confidence in the context.[6]

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

By design, this framework makes simple work of the interpretation of a comparative like (1b). Continuing in the style of Kennedy (1999), adjectival *more* would be interpreted as in (9), the core contribution of which is a strict greater-than relation between degrees. Here, d indicates the semantic value of the *than*-clause, d, and μ indicates a measure function (i.e., minimally, a map from the subject to their degree of μ-ness).

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

Given these assumptions, (1b), repeated in (10a),[7] will have the truth conditions in (10b), so long as *confident* saturates its propositional argument prior to combination with *more*. This says that Ann's degree of confidence in the truth of proposition expressed by *it's raining* strictly exceeds her degree of confidence in the proposition expressed by *it's snowing*.

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

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

[6] The interpretation of POS on the Kennedy-style account we outline looks as in (i) below, assuming that *confident* combines with its propositional complement before combining with POS.

i. $[\![\text{POS}]\!]^g = \lambda \mu_{ed}.\lambda x_e.\ \mu(x) \geq \mathbf{standard}_C(\mu)$

[7] The LF must have structure like that depicted in the tree below, ignoring the copular verb and the internal structure of the *than*-clause. This makes the standard assumption that *more* is discontinuous with the *than*-clause in the string due to obligatory extraposition (Bresnan 1973; cf. Bhatt and Pancheva 2004, Alrenga et al. 2012).

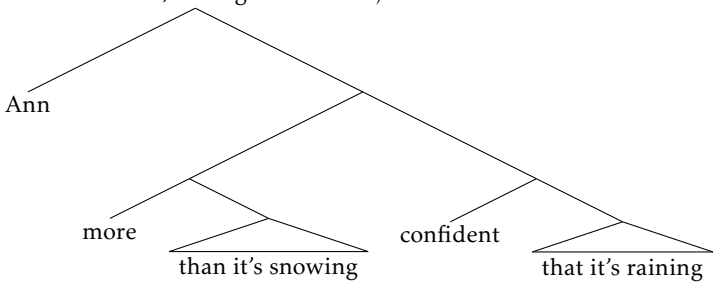

(10) a. Ann is more confident that it will rain than that it will snow.
b. $[\![(1b)]\!]^g$ = true iff $\mathbf{confident}_a(\mathbf{rain}) > \mathbf{confident}_a(\mathbf{snow})$

This proposal is simple and elegant. As it stands, however, it is not obvious how nominal confidence reports should be analyzed within the same framework. While we would not say that extending this analysis of adjectival reports to their nominal counterparts is impossible, no fully developed extension of this sort yet exists.[8] Meanwhile, *Ann has confidence that it will rain* just seems like it will require different compositional resources than what the standard account provides for the adjectival case. And while the account of nominal comparatives with *confidence* will overlap in important respects with better-considered nominal cases like *Ann has coffee* (though cf. Francez and Koontz-Garboden 2017), its abstract and modal character are different and will require care.

In the next section, we show how the picture looks if we start with an account of the nominal reports. In this direction of assimilation, we are able to leverage an existing alternative proposal towards a unified semantics—one that starts with gradability in the nominal domain and extends it to the adjectival case. As it stands, we know of no attempt to assimilate in the other direction, and the challenges facing such an attempt would be substantial.

3 The states-based analysis

Unlike the off-the-shelf degree-theoretic analysis just discussed, our state-based analysis takes as basic the nominal form, exemplified by (11a) and (11b).

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

For some initial indications of its semantic role, observe that *confidence* is a mass noun.[9] On a first pass, this is 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).[10]

(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.

[8] This does not seem to be an accident, either. Some of the compositional worries that flow from the degree-theoretic treatment of gradable adjectives—e.g., the necessity of a covert POS or POS-like element whenever the GA occurs in the positive form—would be multiplied for any noun that occurs in the comparative form (e.g., *confidence*, *coffee*, *toys*) and any verb (e.g., *run*, *jump*, *like*).

[9] Plausibly, gradable adjectives like *confident* nominalize as mass because they bear the 'subinterval property', e.g. P has the subinterval property just in case whenever P is true over an interval i, P is also true over non-arbitrary sub-intervals of i. The closely-related notion of 'divisiveness' generally holds of concrete mass nouns like *water*.

[10] (12a) is negative because such occurrences of *much* have an NPI-like distribution; see Solt 2015.

In this section, we do three things. First, we introduce a standard neodavidsonian analysis of abstract mass nouns. Second, we introduce Wellwood's (2014, 2015) extension of the neodavidsonian framework to gradable adjectives. Third, we modify and update that analysis by showing how it can produce plausible truth conditions for the positive form. As we observe, the updated Wellwood framework we land on is silent on the treatment of modals and propositional attitudes. Addressing that task will be the focus of the rest of the paper.

The analysis of abstract mass nouns

Mass nouns like *height*, *swagger*, *confidence*, *pleasure*, and *ambition* denote properties of states, the 'mass' subtype of the type of eventualities, v.[11] For example, (13).

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

Bare ascriptions of the form *S has F* where *F* is an abstract noun are interpreted along standard neodavidsonian lines.[12] Specifically, *Ann has swagger* is interpreted as an existential claim about states involving Ann as the 'holder' or 'bearer' of the state; see (14).

(14) $\exists s_v[\mathbf{holder}(s, a) \wedge \mathbf{swagger}(s)]$

The comparative targeting an abstract mass noun, then, looks just like a comparative targeting any other mass noun, modulo a distinction between entities and eventualities carried by the basic 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 the context-sensitivity of measures[13]), 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 (16) is interpreted as in (17), where δ abbreviates the *than*-clause degree. This may be read, 'there is a state of swagger s that Ann is in, the measure of which is greater than δ.' Given standard assumptions about the derivation of the meaning of the *than*-clause, the value of δ in (17) is equivalent to $max(\lambda d.\exists s_v[\mathbf{holder}(s, k) \wedge \mathbf{swagger}(s) \wedge g(\mu)(s) \geq d])$.

[11] We choose the formulation 'states' for simplicity. Related possibilities for valuing the *s*s are, for example, 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.

[12] 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 our target case—a noun like *confidence*.

[13] Wellwood 2014, 2019 discusses empirical and theoretical issues surrounding the handling of this component of *more*'s indeterminacy.

(16) Ann has more swagger than Kim.

(17) $\exists s_v[\textbf{holder}(s,a) \wedge \textbf{swagger}(s) \wedge g(\mu)(s) > \delta]$

The selection of values for $g(\mu)$ must meet one further condition which guarantees that the (strict) ordering relations on the measured domain are preserved in the corresponding degree ordering. This condition may be stated as in (18) (cf. Schwarzschild 2006, Nakanishi 2007, and Wellwood 2014, 2015), which says that if any two states are strictly ordered in a certain way, then their μ-measures are ordered in the same way in the degree ordering.

(18) $\forall s, s' \in Dom(\langle D, \gtrsim \rangle)$, if $s > s'$, then $g(\mu)(s) > g(\mu)(s')$.

The general idea is that the function expressed by a (measurable) abstract mass noun has a structured domain (of states), and the structure of such domains 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 other cases: the analysis of *more coffee*, for instance, assumes that *coffee* expresses a function whose domain is ordered by a part-of relation on portions of coffee, as maintained by much prior literature, and that permissible values of μ in the comparative (e.g., a volume or weight measure) respect strict part-whole relations on those portions.

The upshot, then, is that ascriptions of abstract properties are states-based. Comparative confidence reports in the nominal domain will 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.

Gradable adjectives in a states-based framework

As anticipated, we mostly build on the neodavidsonian semantics for gradable adjectives in Wellwood (2014, 2015), on which they express properties of ordered states. We first illustrate these details using *tall*.

Following Wellwood, computing the interpretation of a simple sentence like (19) involves interpreting *tall* as a property of states, as in (20a), with the resulting logical form for (19) in (20b), according to the usual neodavidsonian assumptions.

(19) Mary is tall.

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

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

The advantage of this states-based analysis is that it makes quick work of the comparative case.[14] The denotation for *more* provided in (15) composes just as

[14] Wellwood (2014, 2015) cites general evidence for states in adjectival denotations, in line with which it is easy enough to observe that sentences 'about' John's confidence, whether reported with adjective or noun, introduce something that can be the antecedent for anaphors like *that* and *it*. Going beyond her diagnostics, we note that: (i) those things must be of the sort that can figure into explicitly causal language, contributing either cause (*John's confidence/being confident that it would rain made it snow*) or effect (*The cloud's appearance made John confident/gave John confidence that it would rain*); and (ii) their presence can straightforwardly help explain scopal interactions between confidence reports and *because*-clauses (i.e., the evident ambiguity of sentences like *Ann is confident/has confidence that Mary is in Paris because Gary is in Paris*).

easily with the denotations that emerge from this account of gradable adjectives. (In fact, according to this account these denotations have the exact same type.)

(21) Mary is taller (=more tall) than Beth.

(22) $\exists s_v[\textbf{holder}(s,m) \wedge \textbf{tallness}(s) \wedge g(\mu)(s) > \delta]$
$\delta = max(\lambda d.\exists s_v[\textbf{holder}(s,b) \wedge \textbf{tallness}(s) \wedge g(\mu)(s) > \delta])$

An initial concern about this framework is that it leaves open exactly what counts as a state of having tallness. If having any degree of height is enough to count as having tallness, the meaning for the positive form will be too weak. (19) would end up saying that Mary has a state of height; since even the shortest person in the world has some degree of height, we should expect that (19) is always true so long as Mary exists. Alternatively, **tallness** could be strengthened so as to apply to states of height that exceed a certain threshold—say six feet. The problem then is that this would generate unwanted entailments for the comparative form: a comparative of the form *a is F-er than b* would be predicted to entail that at least *a* is *F*.

One solution to this problem would be to import a move from the scalar playbook and introduce a special POS morpheme, which produces the right meaning for the positive form, but which works with the types presumed in the nominal setting. Instead, we explore an expansion of the framework that solves this problem without POS.[15]

Positive form without positive morphemes

The standard approach to capturing the meaning of adjectives in the positive form involves introducing a null morpheme, POS or ABS, to operate on the degree value or argument of a lexically-introduced measure function. On a Kennedy-style semantics, for example, the sentence in (19) would be interpreted as in (23), with $\textbf{standard}_C(\mu_{\textbf{height}})$ a contextually-provided threshold on a scale representing height.

(23) $\mu_{\textbf{height}}(m) > \textbf{standard}_C(\mu_{\textbf{height}})$

So far, the semantics in (20b) says only that Mary instantiates some state of tallness, without specifying quite what that means. It is straightforward work, however, to postulate an appropriately-outfitted POS that revises (20b); see (24).[16]

(24) $\exists s_v[\textbf{holder}(s,m) \wedge \textbf{tallness}(s) \wedge g(\mu)(s) \geq \textbf{standard}_C(g(\mu))]$

Such a POS operator would, compositionally, turn a predicate of states (e.g. *tall*) into a different predicate of states, one which holds when a state has the relevant property to a sufficient degree in context *C*.

Granting that such a move is available, and that the worst consequence of adopting it would be that the POS assumption spreads from the adjectival to the

[15] In this, we follow Cariani et al. (forthcoming). As shown there, this approach to eliminating POS receives a natural formulation on the states-based approach but is not exclusively available to it; a variant can be incorporated easily enough into the standard picture.

[16] Take *more* in (15), replacing > with ≥, and filling in the appropriate assumptions about the right-hand side of the relation instead of saturation by a *than*-clause denotation, etc.

nominal case, we pursue a novel characterization of the positive form as advanced in Cariani et al. (forthcoming). This approach capitalizes on the order-theoretic properties of gradable stative predicates. To foreshadow the general approach, we draw a basic distinction between states of height and states of having tallness. In particular, every state of having tallness is a state of height but not vice versa. States of height are ordered with respect to one another. Positive occurrences of *tall* apply to height states in the 'positive region' of this ordered set, while forms like *taller* apply to any state of height.

In this 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, (25). (The subscript '**height**' flags that the relevant states are states of height—more on this momentarily.)

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

Crucially, we make a distinction between 'height states' and 'having tallness states'. So for example, the tallest individuals in a context will instantiate states of height and states of having tallness, while the shortest individuals will instantiate states of height but not states of having tallness. Thus, the meaning of *tall* isolates which of the height states count as having 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.[17] 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 states of having tallness.

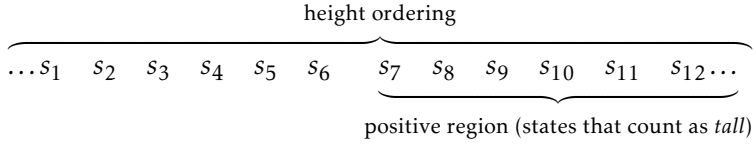

Under this approach, the well-known context-dependence of gradable adjectives is reflected in how the positive region of the background ordering is determined. As a reminder, the central datum here is that a sentence like *Ann is tall* is true just in case Ann exceeds some standard for having tallness in the context. We implement this context-sensitivity in the states-based framework via 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**.

Concretely, we assume the lexical entry for *tall* in (26a), with **tallness** indexed by C. In general, the truth-conditional effect of this is given in (26b), for any gradable property, **g-ness**, associated background ordering $\succsim_{\textbf{g-ness}}$, and state s. The function **contrast** maps the target state s to a salient *contrast state* s'. In the case of

[17] 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. See also Francez and Koontz-Garboden 2017, ch.3 for a related proposal with a different implementation.

(19), we now have the logical form in (27), which says that (19) 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.[18] (For most ordinary gradable adjectives, the expressive capacity provided by **contrast** is overkill, but it will be crucial to our model of gradable attitudes and for the further extensions to gradable modal properties.)

(26) 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 $s \gtrsim_{\textbf{g-ness}} \textbf{contrast}_C(\textbf{g-ness})(s)$

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

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

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 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.

Revisiting the comparative form

Our analysis avoids POS by building extra information into the relevant lexical item: their entries distinguish a background ordering and a 'positive region' within that ordering. Comparative morphology uses part of this information, and discards the rest. In particular, comparative morphology discards the information encoded in the **contrast** function used to delimit the positive region. This blocks any inference to the positive form, and allows 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 having tallness to a background ordering of states of height.

For ease of exposition, it's convenient to 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 (28), **background** is invoked in a presupposition on the state argument: the relevant state is presupposed to be part of the background ordering.

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

For illustration, the compositional process underlying (29) is sketched in (30). Abbreviating the contribution of the *than*-clause using δ, the property expressed

[18] 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.

by the degree phrase, *taller than Sue*, is a property of states in the background ordering of $\textbf{tallness}_C$ whose $g(\mu)$-measure is greater than δ, (30c). Combining the rest, (29) is interpreted as in (31), which says that Mary is in a height state (i.e., a state in the domain of the background ordering for $\textbf{tallness}_C$), the $g(\mu)$ measure of which is greater than that of a corresponding state of Sue.[19]

(29) Mary is taller than Sue is.

(30) a. $[\![\text{than Sue}]\!]^g = \delta$

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

c. $[\![\text{[tall [er than Sue]]}]\!]^g = \lambda s_v :\ s \in \textbf{background}(\textbf{tallness}_C).\ g(\mu)(s) > \delta$

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

States and attitudes

Let us take stock. We initially focused on the case of abstract mass nouns. We took notice of the standard approach in terms of states. Next, we presented Wellwood's proposal to the effect that it is possible to give a unified, states-based treatment of abstract mass nouns and gradable adjectives. Finally, we presented a new account of the relation between the positive and comparative forms within Wellwood's framework.

There is one last bridge left to cross before we can provide an account of confidence reports. A key difference between *swagger* and *confidence* is that states of confidence are propositional attitudes. The last module of our account of confidence reports, then, will aim to capture how the contents of confidence states figure in their truth conditions.

4 The semantics of confidence reports

Comparative confidence reports

Assume that the schematic comparative sentences in (32) have the same interpretation, i.e. (33). 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 (32)). 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.

(32) a. Ann is more confident that p.

b. Ann has more confidence that p.

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

The relevant background ordering ranks states by how confident a holder a is in the truth of the state's theme, θ. Thus, the background ordering includes states of high as well as low confidence. Both *confident* and *confidence* single out

[19] We have indicated the restriction to the background ordering on height states as a restriction on the domain of the existential quantifier in (31).

a positive region of this ordering that includes the upper bound. The ordering is tracked, as before, via presupposition, as shown in (34). Here, the superscript $\mathbf{h}(s)$ abbreviates $\mathbf{holder}(s)$ and the subscript **conf** indicates that the ordering concerns states of confidence.[20]

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

In the end, logical forms like (33) 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 (35) have the same interpretation—i.e., (36). 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.

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

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

The fact that this analysis predicts that (35a) and (35b) 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 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.

[20] 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.

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, (37) doesn't require that the weatherman assign zero confidence to the possibility of rain in order to be judged true.

(37) 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 predicts *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 (35a) and (35b) 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.

Varieties of confidence reports

Semanticists often puzzle over variant uses of *confident*/*confidence* when considering confidence reports that appear, at least on the surface, to require a different treatment. In this section, we briefly show how our analysis would approach this broader distribution. We will not provide anything like a complete account of these varieties, however we think the states-based analysis in particular highlights promising avenues for analysis.

First, it is often observed that, in addition to its occurrences with clausal complements, *confident* (like *confidence*) can surface with prepositional complements, (38), and without any apparent complement whatsoever, (39).

(38) Marie is confident in Bill.

(39) Marie is confident.

The *in*-phrase in (38) doesn't obviously stand in for any propositional content. (39), in contrast, can be read just like (38) or as proposition-involving like our target cases, depending on what is salient in the context. For that matter, (39) can also be read as saying that Marie is a confident person, simpliciter, or that she is confident as an athlete, etc.

Our focus here has been on the proposition-involving uses, however we think it is plausible to see such data as reflective of the kind of flexibility afforded by our neodavidsonian analysis. Conceiving of *confident*/*confidence* as introducing a set of states allows us to suppose that such a set may be restricted in various ways; such restrictions are the function of elaboration by thematic predicates like *in Bill*

or modifiers like *as an athlete*. Consider, for an analogy, that *Al ran to the park* is telic and supports different inferences than *Al ran in the park* which is atelic (cf. Rothstein 2016). Elaborating the ways in which the relevant sets can be restricted is a matter for semantic typology; regardless, we might expect that any such possible refinements can be inferred in the absence of any thematic expression, as we are plausibly observing with uses like that in (39).[21]

It is inviting to think about the contrast between *be confident* and *feel confident*, as well, but this will require a semantics for *feel* which is clearly beyond the scope of the paper. Still, we may say as preliminary remarks that it appears that the assertibility conditions for *feel confident* line up with those for *be confident*. It strikes us (and an anonymous reviewer) bad to assert the one but deny the other, (40a)-(40b).

(40) a. ? Marie is confident that it's raining but she doesn't feel confident that it is.

b. ? Marie isn't confident that it's raining but she feels confident that it is.

Yet the reason for the badness of (40) is not the logic of *feel X*, because sentences like those in (41) don't strike us as bad in the same way.

(41) a. Marie is tired but she doesn't feel tired.

b. Marie isn't tired but she feels tired.

Plausibly, people bring to bear some kind of default assumption to the effect that, typically, as something general that we know about the world, experiencing the feeling of confidence comes close to being confident and vice versa; this is not true of what we know about being tired.

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[22] that, under determinate circumstances, subjects appear to routinely violate this constraint. As a result, it appears that, under the right circumstances, the sentences in (42) can be true together.

(42) a. John is not confident that Linda is a bankteller.

b. John is confident that Linda is a feminist bankteller.

Our semantics is 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.

[21] The same likely goes for thinking about Individual-level vs Stage-level attributions of confidence. Compare *She's a confident person* and *She's a confident swimmer*.

[22] 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 (43a) to (43b) is validated by our semantics. Let us introduce σ as a metalinguistic variable ranging over individuals, and continue to use p and q as metalinguistic variables for propositions.

(43) a. σ is confident that p, σ is more confident of q than of p
b. σ is confident that q

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

Transitivity. (44a) and (44b) entail (44c).

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

Antisymmetry. (45a) and (45b) entail (45c).

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

Connectedness. (46) is a logical truth.

(46) Either σ is at least as confident of p as of q, or σ 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 (42) 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 (47) appears to be contradictory.

(47) 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 gradable adjectives and the comparative will vindicate the entailment in (43). Conversely, the pattern in (48) requires more controversial assumptions. Suppose that the contrast function invoked by the positive form of the adjective always maps a proposition to its negation (a corollary of this is that one cannot be both confident that p and confident that $\neg p$). Then we predict the entailment in (48): (48a) requires that confidence-that-p states be ranked higher than confidence-that-not-p states, and *more* will always map higher-ranked states to higher degrees, in line with its monotonicity condition.

(48) a. σ is confident (/has confidence) that p.

b. σ 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 its 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.[23]

Second, there seem to be intuitive cases where an agent is confident in a proposition, even though they are more confident of its negation.[24] 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. (49) seems true in this context, despite the mass of Clara's confidence favoring some team other than the Warriors.

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

5 Beyond *confidence*

Having presented our analysis of confidence reports, we should like to quickly touch on a variety of possible extensions. Some of these extensions are immediate, and can be resolved entirely within the scope of our discussion. Others, as we discuss, will require separate development.

Conditional confidence

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 (50a). Such conditional ascriptions sound roughly equivalent to self-ascriptions of confidence in the conditional, as in (50b).

(50) 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.

[23] 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 (48) 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.

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

This is plausibly an instance of a more general phenomenon: doxastic attitudes, as well as other kinds of attitudes, may be restricted by conditional antecedents (Blumberg and Holguín, 2019; Jerzak, 2019). There are plausible off the shelf modifications one might make to the semantics to capture this sensitivity to conditional antecedents. One option would be to endow confidence reports with a modal base and allow conditional antecedents to restrict it. Another option might be to make the background ordering sensitive to an information state i which can be operated on by conditional antecedents. Specifically, revise (34) to (51)—indexing the ordering $\succsim$ with i. Then, assume a semantics for conditionals (broadly in the style of Kratzer 1991; Yalcin 2007; Kolodny and MacFarlane 2010) in which conditional antecedents restrict i. Since (51) makes the ordering $\succsim$ dependent on i, this predicts that one can assert (50a) even if one isn't unconditionally confident that Lisa is at the lab. Choosing between these options is, of course, beyond the scope of the present investigation.

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

certain and *certainty*

The kind of analysis for *confident*/*confidence* we proposed can be extended to a variety of other adjectives which report gradable attitudes and have nominal counterparts. This is not surprising. Confidence reports are embedded in a network of inferential and semantic connections with reports of certainty, doubt, and subjective plausibility. For instance, (52a) asymmetrically entails (52b), which in turn appears inconsistent with (52c). Similarly, (53) is evidently defective.

(52) a. Ann is certain that the dress is blue.
b. Ann is confident / has confidence that the dress is blue.
c. Ann doubts that the dress is blue.

(53) ? Ann is certain that the dress is blue, but she isn't confident.

Our approach can be generalized smoothly to these other gradable attitudinal adjectives. To show this, we will focus here on the adjectives *certain* and *sure*, and on the noun *certainty*. We will also offer some pointers to an analysis of attitude verbs such as *doubt*.

Focusing first on *certain*, we can first observe that *certain* leads a double life, only one of which can concern us here. On the one hand, like *confident*, it is a gradable attitudinal adjective, i.e. it is used to describe the mental state of a subject. The mark of this usage of *certain* is the presence of a non-expletive subject, as in (54), which is repeated from above.

(54) Ann is certain that the dress is blue.

On the other hand, *certain* may be used with an expletive subject to produce an 'impersonal' reading that is more akin to an epistemically modalized claim, an example of which is in (5).

(55) It is certain that the dress is blue.

We set this impersonal reading aside and focus on the attitudinal use.

As for *certainty*, there appears to be a full analogy with the attitudinal uses of *certain*. Having certainty that p appears to asymmetrically entail having confidence that p, as is shown by the data in (56). So far as we can tell, 'impersonal' uses of *certainty*, as in (57), are also allowed.

(56) a. Bob has confidence, but not certainty, that the dress is blue.
b. ? Bob has certainty, but not confidence, that the dress is blue.

(57) There is certainty that the dress is blue.

In light of these data, we hypothesize that *certain* and *certainty* exploit the same background ordering of states as *confident* and *confidence*—similarly to e.g. *hot* and *warm*. We keep calling this ordering the **confidence** scale. The difference between the two adjective/noun pairs is that they are assigned different positive regions within that ordering. In particular, *certain* and *certainty* are assigned a narrower positive region than *confident* and *confidence*, as shown by Figure 2.

Figure 2: The ordering for *confident* and *certain*.

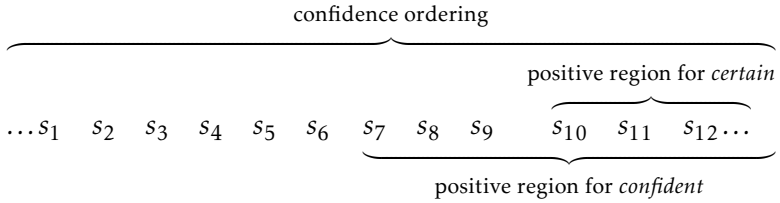

The resulting truth conditions for a bare ascription of certainty are in (58).

(58) a. A is certain that p.
b. A has certainty that p.

(59) $\exists s_v[\mathbf{holder}(s,a) \wedge \mathbf{certainty}_C(s) \wedge \theta(s,p)]$
$\equiv \exists s_v[\mathbf{holder}(s,a) \wedge s \gtrsim^{\mathbf{h}(s)}_{\mathbf{conf}} \mathbf{contrast}_C(\mathbf{certainty})(s) \wedge \theta(s,p)]$

These assumptions immediately predict the pattern of asymmetric entailment that we have observed, i.e. that one can ascribe confidence without certainty, as in (56a), but not certainty without confidence, as in (53) and (56b). Moreover, we have the clear defectiveness of predicating certainty while denying confidence, as in (52a), yet the clear acceptability of predicating confidence while denying certainty, as in *Ann is confident that the dress is blue, but she isn't certain*.

One nice feature of this systematic approach is that it readily explains why the comparative forms of these attitude-reporting adjectives are so close in truth-condition. That is, we can explain why the sentences in (60) sound approximately equivalent to each other. Because the **contrast** function (and thus sensitivity to the positive region) is disabled in the comparative form, these sentences express exactly the same comparison.

(60) a. A is more confident that p than that q.
b. A is more certain that p than that q.

sure

sure appears to be roughly synonymous with *certain*. In particular, it seems that one cannot be certain without being sure, and *vice versa*, (61).

(61) a. ? Ann is certain but not sure that the dress is blue.
b. ? Ann is sure but not certain that the dress is blue.

The only exception to this pattern that we have noticed is that *sure*, but not *certain*, may be used to assent to a claim just made by another speaker.[25] But this difference appears to be idiomatic, and we set it aside. (In particular, notice that both *sure* and *certain* are fine if they are used as part of a full sentence, as in *I'm sure*/*I'm certain*.)

(62) A: The dress is blue.
B: Sure./ ?Certain.

Given this evidence, we suggest that *sure* and *certain* are synonymous, and that our analysis of *certain* also applies to *sure*.

convinced, *doubtful*, and other attitudinal predicates

Similar strategies could be pursued for other kinds of attitudinal adjectives, such as *convinced* or *doubtful*. In particular, there is evidence for the idea that *convinced* and *doutbful* operate on the same background ordering as *confident* and *certain*. For example, the conjunctions in (63a)-(63c) suggest that *convinced* or *doubtful* also enter the patterns of entailments with *confident* and *certain* that we have discussed above. On the other hand, in our view, these similarities do not amount to a knock-down argument, so we remain uncommitted to this extension, pending a more systematic investigation of their entailment relations.

(63) a. Ann is confident but not convinced that the dress is blue.
b. # Ann is convinced but not confident that the dress is blue.
c. # Ann is confident but doubtful that the dress is blue.

A reviewer also indicates that there may be systematic connections between expressions expressions of confidence and other expressions of suspicions and doubt, exploiting the entailments in 64. We defer discussion of such links to future work:

(64) a. # Ann is confident that the dress is blue but she doesn't suspect that it is.
b. # Ann is confident that the dress is blue but she doubts that it is.

From *confident* to *likely*

Confidence and certainty reports can be 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

[25] Thanks to David Barner for this pointer.

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 (65).[26]

(65) 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*. The same contrast shows up with the pairs in (66) and (67) that involve our target expressions.

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

(67) 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.

While these examples show that probability operators are importantly different from attitude verbs, there are also 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 (68) sound roughly equivalent, and similarly for the sentences in (69).[27]

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

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

We emphasize that this bridging observation is merely a 'rough equivalence', since we have encountered a variety of judgments on the matter. While some people judge them to be equivalent, others hear (68a) as somewhat stronger than (68b). This intuition is supported by the observation that (70) can be heard as consistent.

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

[26] 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.

[27] We do not, of course, deny that there is another reading of (68a) on which this equivalence fails. Given this other reading, "I am confident that p" means roughly that I have faith that p will happen.

While we do not deny these complications, we also note that the intuitions in favor of equivalence are much stronger with the comparative sentences in (69). If those intuitions were taken at face value, they would support the need to forge an inferential link between gradable attitude expressions and probability operators.

As a reviewer points out, there are deep analogies between *think/believe likely* on the one hand and *confident* on the other when it comes to their evidentiality requirements. Thus both (71b) and (71c) require that Susanna has not tried the pasta, while (71a) is used—and in fact *preferably* used — in contexts in which she has tried the pasta.

(71) a. Susanna thinks the pasta is delicious.
b. Susanna thinks the pasta is likely to be delicious.
c. Susanna is confident that the pasta is delicious.

There are also more direct considerations for extending the states-based approach to probability operators. The main one is that, as for confidence reports, related probability claims can be expressed using nominal forms, for example (72).

(72) 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.[28]

Moreover, the array of data that suggested the presence of an eventuality argument in confidence reports can be replicated for likelihood claims. For example, these nominal forms can introduce causes, (73), and effects, (74). This suggests that, as for confidence reports, that the sentences in (73) and (74) involve concrete entities linked to likelihood that can enter into causal relations.

(73) 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.
(74) 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.

Threading together the similarities and differences between confidence reports and probability operators with regards to their semantic profiles is a complex project. Our hope here has just been to provide some initial direction and focus for that project.

[28] 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.

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. The central insight of the approach encodes a division of labor with regards to degree semantics. Expressions like *confident* carry along information about the sorts of state structures they are relative to. We exploit these background orderings in our formulation of a (degree-less) positive form, exploiting scales (i.e., ordered sets of degrees) only when dealing with the comparative form.

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 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.

Kyle Blumberg and Ben Holguín. Embedded attitudes. *Journal of Semantics*, 36(3): 377–406, 2019. doi: 10.1093/jos/ffz004.

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

Fabrizio Cariani, Paolo Santorio, and Alexis Wellwood. Positive gradable adjectives ascriptions with states, not degrees. In *Proceedings of Sinn und Bedeutung 2022*, forthcoming.

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.

Itamar Francez and Andrew Koontz-Garboden. Semantic variation and the grammar of property concepts. *Language*, 91(1):533–563, 2015. doi: 10.1353/lan.2015.0047.

Itamar Francez and Andrew Koontz-Garboden. *Semantics and Morphosyntactic Variation: Qualities and the Grammar of Property Concepts*. Oxford Studies in Theoretical Linguistics. Oxford University Press, Oxford, UK, 2017.

Thomas Grano and Stuart Davis. Universal markedness in gradable adjectives revisited: The morpho-semantics of the positive form in Arabic. *Natural Language and Linguistic Theory*, 36:131–147, 2018.

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. Degree operators and scope. In Brendan Jackson and Tanya Matthews, editors, *Proceedings of SALT X*, pages 40–64, Cornell University, Ithaca, NY, 2000. CLC Publications. doi: dx.doi.org/10.3765/salt.v10i0.3102.

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

Ethan Jerzak. Two ways to want? *Journal of Philosophy*, 116(2):65–98, 2019. doi: 10.5840/jphil201911624.

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. doi: 10.1007/BF00351812.

Ewan Klein. The interpretation of adjectival comparatives. *Journal of Linguistics*, 18(1):113–136, 1982. doi: 10.1017/S0022226700007271.

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

Angelika Kratzer. Modality. In Arnim von Stechow and Dieter Wunderlich, editors, *Semantics: An international handbook of contemporary research*, pages 639–650. Mouton de Gruyter, Berlin and New York, 1991.

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.

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.

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.

Robert Pasternak. A lot of hatred and a ton of desire: Intensity in the mereology of mental states. *Linguistics and Philosophy*, 42:267–316, 2019.

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. Oxford University Press, 2016.

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

Susan Rothstein. Aspect. In Maria Aloni and Paul Dekker, editors, *The Cambridge Handbook of Formal Semantics*, chapter 2, pages 342–368. Cambridge University Press, 2016.

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.

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.

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.

Alexis Wellwood. *The meaning of* more. Studies in Semantics and Pragmatics. Oxford University Press, Oxford UK, 2019.

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

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