THE UNIVERSITY OF CHICAGO

ON TONGUES: THE GRAMMAR OF EXPERIENTIAL EVALUATION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE HUMANITIES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF LINGUISTICS

BY
PATRICK JOSEPH MUÑOZ

CHICAGO, ILLINOIS
JUNE 2019

To my parents. Of all the things you've given me, I love my mind the most.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

There are too many people to thank for this, and other things – so I'll mention just those that have had the clearest impact on the dissertation itself. First, the committee: to anyone familiar with their work, their influence should be obvious.

Thanks to Chris Kennedy, for being a patient and brilliant advisor. The germ of the idea for this dissertation is in Kennedy (2013), and some of it reflects a broadly 'Kennedian' worldview. In addition to the help with all the topics addressed here – too many to list – thanks for helping me understand that the technical differences between positions are often not as interesting as (or more interesting than) we might think, and that a decompositional analysis is to be preferred wherever possible.

Thanks to Malte Willer, to whom I directly owe a large portion of my technical competence, and whose class in game-theoretic pragmatics in fall 2016 eventually led me to think of linguistic subjectivity in terms of metasemantics, as in Chapter 5. This dissertation was begun in the milieu of Chris and Malte's Subjectivity in Language and Thought grant from the Neubauer Collegium, which made an ideal environment for studying these topics.

Thanks Itamar Francez, for all the teaching and support. Itamar's work on qualities, especially Francez & Koontz-Garboden (2017), gave me a way to think of the basic model-theoretic machinery underlying the treatment of experiences, which is vital to Chapter 2, and on which the rest of the dissertation indirectly depends.

Thanks to Anastasia Giannakidou, for support, for the many classes and discussions, and for the conviction that semantics is cross-linguistic semantics. Anastasia's work on speaker commitment, individual anchors, and the connection between knowledge and epistemic modality, sparked much of the early development of Chapter 4.

Several other grad students working on similar topics were of great help with this material. Tamara Vardomskaya dissertated in the department before me – I owe to her my introduction to the evidentiality literature, which led to my seeing a connection between direct evidentiality and experiential semantics, and large parts of Chapter 4 would not have

been possible otherwise. Carlos Cisneros made me aware of the existence of overt direct evidential markers in Tibetan, which also play a role in Chapter 4. Rachel Rudolph visited UChicago from Berkeley in winter and spring of 2017 – her thoughts on perceptual verbs informed Chapter 3, and her notion of experiential predicates e.g. in Rudolph (forthcoming) is close to mine. Sophia Sklaviadis, while she was in the philosophy department at UChicago, introduced me to the experimental literature on subjective predicates.

Thanks to Ivano Caponigro, for being my first teacher in semantics, and also for teaching a seminar at UC San Diego in spring 2012, during which I was first exposed to the work of Peter Lasersohn, and became puzzled by many of the things that I'd eventually work on.

Thanks to Karma Ngodup, for acting as a Standard Tibetan informant, and helping with some of the examples in Chapter 4.

Thanks to the audiences at the Subjectivity in Language and Thought 2017 Workshop at UChicago, at the Evaluative Meanings conference (DGfS 2018) at Universität Stuttgart, and at the UChicago Linguistics and Philosophy Workshop, to whom earlier portions of this material were presented, for their comments.

Thanks to my parents, David and Victoria, to my brothers, Matthew and Alejandro, and to Sammy, Marty, and Cameron.

And finally, since most of our fortunes come from unknown sources: thanks to my *daimon*, and my lucky star.

ABSTRACT

This dissertation offers a treatment of experiential predicates – those predicates true of an individual just in case it is disposed to stimulate an experience of a certain sort – and demonstrates how their experiential semantics, in concert with other semantic and pragmatic mechanisms, give rise to a certain kind of linguistic subjectivity. In particular, speakers tend to evaluate these predicates autocentrically, or with respect to their own experiential dispositions, despite their apparently not encoding reference to an experiencer in their semantic content.

I characterize experiential predicates as a cross-categorial semantic natural class, which arises compositionally out of (i) a basic lexical semantics encoding a type of experience, combined with (ii) the generic quantification that accompanies individual-level predicates in imperfective aspectual environments. I then demonstrate how this experiential semantics systematically gives rise to a lack of coordination in speaker behavior across the population: (i) the felicity conditions on generic quantification, which specify which situations 'count' for the satisfaction of a generic predicate, are for experiential predicates (ii) constrained by independently-attested norms on direct evidence, according to which speakers are obligated to take themselves to be accurate perceivers of a certain sort. The result is that speakers typically take the conditions under which an individual is generically disposed to stimulate a certain experience *simpliciter* to be the same as the conditions under which it is generically disposed to stimulate that experience with they themselves as experiencer.

The result is a certain kind of 'strong subjectivity,' by which not only do the linguistic conventions not fix how a predicate is to be applied in virtue of 'the way the world is' across the speaker population, but they further 'force' speakers to diverge in their truth-conditional behavior by convention, according as their experiential dispositions differ. This allows the language additional expressive power, by which speakers can use truth-conditional material to express opinions. This is taken to be a model of linguistic subjectivity generally, as a metasemantic phenomenon to which non-hyperintensional grammars are blind.

PREFACE

In the last fifteen years, there's been a resurgence of interest in truth-apt constructions that, even relative to a context of utterance, don't impose unique truth conditions in an obvious way. Such constructions are taken to be 'subjective' in that facts about the state of the world, combined with facts about the communally stable meaning of expressions in a context of utterance, are insufficient to deliver a truth value, at least without commitment to some opinion on a non-factual basis. Just about everything under the sun contributes to subjectivity by these standards, and has been construed that way at some point or another: both epistemic and root modality, vagueness, multidimensionality in gradable predicates, future contingents, epithets, degree morphology of excess and sufficiency, conditionals, evidentials, and all the rest, including of course the topic of this dissertation, experiential predicates.

The idea that's been pursued so far is that since these subjective constructions don't determine truth conditions *simpliciter*, yet nonetheless speakers competently take the contents they express to be true or false, they must determine truth conditions relative to something else in addition to the 'way the world is.' The compositional semantics then encodes supplementary values specifying this 'something else' implicitly – we just have to figure out where and how. Presumably, a similar mechanism will be applicable to the dizzying array of cases.

But this picture has turned out to be illusory. There is no semantic phenomenon of linguistic subjectivity, in the way that there is a semantic phenomenon of tense or number. What we have instead are a variety of unrelated semantic sources of a certain ubiquitous higher-order effect, namely that truth conditions are not uniquely determined by conventions of language use. Linguistic subjectivity is therefore a metasemantic phenomenon, which has to do with these conventions being underdetermined. But this underdetermination depends in every case on the specific semantic constructions that, each for their own reasons, license competent variation in speaker behavior, and the idea that there is a single semantic treatment of the phenomenon to be given – maybe relying on a series of parameter values or lexically encoded arguments that serve to 'fill in' what is missing from the truth conditions,

by anchoring semantic content to a perspective of some sort – needs to be given up. Instead of subjectivity being imposed top-down in the compositional semantics, it needs to be built bottom-up, so that it results as a kind of epiphenomenon, traceable to the fact that the conventions encoded by the compositional semantics are insufficient to force speaker behavior to conform perfectly across a linguistic population.

If we're interested in linguistic subjectivity, then we have to examine the semantic phenomena that give rise to it each on its own terms, and show how it can be derived in each case from independent semantic features of the language, which aren't subjective in of themselves. This dissertation sets out to do that with one particularly prominent source of linguistic subjectivity, namely the semantics of *experiential predicates*, which I call those predicates true of individuals just in case they are disposed to stimulate a certain sort of experience, e.g. *tasty, frightening, look tired, sound funny.* The phenomenon is particularly glaring here, because experiential predicates not only vary across the population in their conventionally correct conditions of application, but this indeterminacy is further enforced by the differing sorts of experiences that speakers are liable to have. And so it seems to the researcher that there is something about experiential predicates which requires them to make implicit reference to an experiencer. I'll argue that isn't the right way to think about things, but even so, it's clear that there is a connection between the experiences of speakers and the way they're pragmatically compelled to apply these predicates.

How experiential predicates are supposed to apply, in virtue of what experiences speakers are disposed to have, is a question of how to resolve the truth conditions of predicates along the dimension of what Kennedy (2013) calls 'qualitative assessment.' That experiential semantics specifically is crucial for this assessment is adumbrated in McNally & Stojanovic (2017) and Bylinina (2017), both of which pay attention to the role of experiencer as a key source of subjectivity, and this has been more explicitly called attention to in Anthony (2016) and Rudolph (forthcoming). But the role of experiencer has been recognized, however implicitly, by just about everyone, since the kinds of examples that have been marshaled in

discussion of linguistic subjectivity have often involved experiential predicates, including the favorite exemplar *tasty*, which will be used as a prominent example in this work as well.

This dissertation takes experiential semantics seriously as a starting point in its own right, and shows how the semantic encoding of experience systematically gives rise to a peculiar flavor of subjectivity, by which the application of predicates tracks speakers' experiences. In doing so, it contributes to the treatment of experiential semantics as a topic of intrinsic interest, as well as giving a proof of concept that we can treat linguistic subjectivity from the bottom up in this way. Hence the two components of the subtitle: 'experiential,' because the source of this subjectivity is experiential semantics, and 'evaluation,' because this subjectivity ultimately results in predicates that are evaluative, in the sense that using them requires speakers to make decisions not made by the semantics itself about how to construe their truth conditions.

The chapters of this dissertation aren't meant to be read separately – each takes the previous ones for granted, and there is a lot of cross-referencing, both forward and back. The chapters are arranged from least to most interesting. The contents, briefly:

Chapter 1 is a ground-clearing exercise that also functions as a literature review. It argues that experiential predicates like *tasty* ought not to be relativized in their extensions to experiencers, at any level of semantic description. This paves the way for the rest of the work, to give a semantics for experiential expressions unhindered by this presupposition, and then to show how the higher-order pragmatic and metasemantic effects that we're interested in derive from this.

The remaining chapters are then split among the main two tasks of the dissertation: Chapters 2 and 3 are about experiential semantics, and Chapters 4 and 5 are about the higher-order effects derived from the semantics in concert with independently active norms on linguistic usage. In particular, Chapter 2 sets the groundwork for the semantics of experiential predicates generally, by exploring deverbal psych predicates in English, formed from object-experiencer psych verbs using the suffix *-ing*, e.g. *shocking, frightening, amusing.*

This chapter has a lot of decompositional semantics, and goes in depth as to how we ought to construe experiences model-theoretically, and how to build up experiential semantics piece by piece, in terms of abstract nominals denoting experiential kinds, psych verbalizers, *-ing*, degree morphology, and generic operators. Chapter 3 then turns to a second type of experiential predicate, derived by composing sensory verbs like *look* with adjectives, e.g. *look tired* and *sound funny*. It's shown how these predicates deal with evidence based on experience and sensory modality, and the chapter is crowned with an in-depth derivation of that most famous of all experiential predicates, the star of the show, *tasty*.

Chapter 4 deals with how speakers commit themselves to having had experiences of a certain sort in virtue of using experiential predicates, and demonstrates how experiential predicates are constitutively related to the evidential category of direct evidence. This in turn has consequences for how speakers evaluate the application conditions of experiential predicates, based on the direct evidence – that is, the experience – that they actually have, what it means to construe speakers as accurate or inaccurate perceivers of experiential properties, and finally the normative consequences of this. Chapter 5 then brings things together, and addresses the questions with which we began: how all of this results in systematic speaker variation in the construal of truth conditions of experiential predicates, why speakers evaluate the truth conditions of these predicates with respect to their own experiential dispositions, how this gives rise to the notion that such predicates reflect 'matters of opinion,' and so what is the origin of the experiential flavor of linguistic subjectivity.

CHAPTER 1

A BARE TREATMENT OF EXPERIENTIAL PREDICATES

This chapter offers a preliminary account of experiential predicates, which serves as the basis for the more detailed treatment of experiential semantics generally in Chapters 2 and 3. It uses the English adjective *tasty* as its sole exemplar, citing other predicates only where necessary for expository purposes. The diet of experiential expressions is expanded in later chapters, and the reader can confirm that the arguments in this chapter pertain equally to these other experiential predicates.

Previous research has nearly unanimously claimed that experiential predicates must make reference to an experiencer at some level of semantic description; disagreement has largely concerned at what level. This chapter examines these claims and rejects them: §1.1 and §1.2 discuss extant views on the role of the experiencer in the semantics of experiential predicates, and §1.3 provides a template for a distinct positive proposal.

§1.1 examines contextualism about experiential predicates, the position that such predicates take an experiencer argument, such that (on composing with this argument) their content at a context of utterance always makes reference to some standard of experience or other. This section argues at length that this position is false.

§1.2 examines relativism about experiential predicates, according to which such predicates, while not composing with an experiencer argument, nonetheless make reference to an experiencer in determining their extension, which is supplied by the value of an intensional parameter. This section demonstrates that removing the experiencer argument resolves the contextualist's problems.

§1.3 offers the basics of a bare semantics for experiential predicates, according to which they denote properties true or false of individuals at worlds of evaluation *simpliciter*, with no reference to an experiencer. This section shows that such an account can recapitulate the advantages of a relativist treatment in a simpler way, so that the experiencer parameter, while empirically harmless as far as the concerns of this chapter go, is superfluous. The bare

semantics is then taken for granted in the rest of the work.

1.1 Contextualism: against the experiencer argument

The contextualist holds that experiential predicates are extensionally dyadic: at a world of evaluation, they denote relations between individuals, as in the following standard denotation for *tasty*.[1]

(1) $[\![\text{tasty}]\!]^w = \lambda x_e.\lambda y_e.tasty''(w)(x)(y)$

x is an experiencer, or individual that sets the relevant experiential standard, and y is a stimulus, or individual that produces experiences. '$tasty''$' is a metalanguage predicate, to be read as follows: $tasty''(w)(x)(y)$ iff at w, y is tasty by the standard of experience set by x (the two apostrophes signify that the predicate is dyadic relative to worlds).[2]

For *tasty*, this talk of 'the experiential standard set by x at w' must track pleasant gustatory experiences in the right way, minimally by determining which individuals are disposed to produce such experiences in x at w (2a). A stronger reading of the metalanguage predicate would instead track not mere dispositions to produce experience, but production of experience *simpliciter* (2b).

(2) a. $tasty''(w)(x)(y)$ iff at w, y is disposed to produce gustatory pleasure in x.

b. $tasty''(w)(x)(y)$ iff at w, y produces gustatory pleasure in x.

1. 'Individual' is used here to cover any object traditionally of type e, whether, atomic, plural, or kind. This allows the contextualist flexibility in characterizing experiential standards.

2. Standards of experience can be modeled in many ways. Taking ς to be a function from world-individual pairs $\langle w, x\rangle$ to the standard of experience set by x at w, one might say that a standard of experience simply is an individual, such that trivially $\varsigma(w, x) = x$, or that it is an abstract object that models the kinds of experiences that x is disposed to have at w: e.g., $\varsigma(w, x)$ might be a function from individuals y to sets of experiences, such that $\varsigma(w, x)(y) = \{z : z \text{ is an experience that } y \text{ is disposed to produce in } x \text{ at } w\}$. The interpretation of the metalanguage predicate can be rephrased to coincide with any of these options, with the same results: e.g., on the latter way of construing experiential standards, (2a) could instead read, '$tasty''(w)(x)(y)$ iff the experience of gustatory pleasure is a member of $\varsigma(w, x)(y)$.' The metalanguage predicate could also be recast as type $\langle s, \langle \sigma, \langle e, t\rangle\rangle\rangle$, where σ is the type of experiential standards, or *tasty* itself could be made to be of extensional type $\langle \sigma, \langle e, t\rangle\rangle$, and saturate its first argument with an experiential standard. This won't make a difference in what follows: the reader can translate where interested.

Experiential predicates alone thus don't denote properties of individuals at any level of semantic description, and at worlds, there are no experiential properties *simpliciter*, but rather only stimuli's production of, or disposition to produce, experiences in experiencers.[3]

The contextualist requires that the experiencer argument to *tasty* be somehow saturated by an individual. The predominant method of doing this is by taking *tasty* to select for an e-type internal argument, whose overt expression is a (usually) *to*-headed experiencer PP, as in (3) (cf. Glanzberg 2007: 11-12, fn. 9; Stephenson 2007a: 519-520; Stojanovic 2007: 701, fn. 15; Sæbø 2009: 337, 339; Schaffer 2011; Pearson 2013; Snyder 2013: 286-289).

(3) Licorice is tasty to Alfonse.

The easiest way to do this (cf. Stephenson 2007a: 500) is to treat the preposition as vacuous (4a), and so allow its DP complement to saturate the first individual argument of *tasty*. (4b) then denotes a property true of those things that produce, or are disposed to produce, gustatory pleasure in Alfonse (at w).

(4) a. $[\![\text{to}]\!]^w = \lambda x_e.x$

b. $[\![\text{tasty [to Alfonse]}]\!]^w = \lambda y_e.tasty''(w)(a)(y)$

Tasty to Alfonse imposes a direct experience constraint on the individual it composes with: (3) implies that Alfonse has tasted licorice. This implication is presuppositional: it can't be canceled (5), can be challenged by discourse moves that target presuppositions (6), and can project out of classical holes (7), where (7a) involves non-external and non-metalinguistic negation. (7b)'s prejacent and (7d)'s antecedent can also allow for local accommodation of the presupposition, and in (7c), even if the presupposition doesn't project, part of what's asked must be whether Alfonse has tasted licorice.

3. §1.1.2 and §1.1.5 entertain an alternative view, where the contextualist saturates the experiencer argument to '*tasty*$''$' in the predicate itself, so that *tasty* denotes a property, but one intrinsically relativized to a standard of experience. A contextualist also might conceivably admit the existence of experiential properties at worlds *simpliciter*, while denying that experiential predicates denote them; for convenience this possibility is ignored.

(5) #Licorice is tasty to Alfonse, but he hasn't tasted it.

(6) A: Licorice is tasty to Alfonse.

B: Wait a minute – Alfonse hasn't tasted licorice.

(7) a. Licorice isn't tasty to Alfonse.

b. Licorice must be tasty to Alfonse.

c. Is licorice tasty to Alfonse?

d. If licorice is tasty to Alfonse, we should get him some.

$\hookrightarrow$ Alfonse has tasted licorice.

One way to account for this is by committing to a modification of the strong interpretation of '*tasty*$''$' in (2b), so that it both enforces that the stimulus has actually produced gustatory pleasure in the experiencer, and allows for presupposition failure when this condition isn't met.

(8) $tasty''(w)(x)(y)$ iff:

at w, the gustatory experience that y produces in x is pleasant.

It follows from (8), reading the definite description presuppositionally, that if there is no gustatory experience that y has produced in Alfonse at w, then the value of '$tasty''(w)(a)(y)$' is undefined. Where 'production of gustatory experience' is read as production of experience of a certain phenomenological character by contact with the taste receptors, then this requires tasting, and the presupposition is secured.

Another option is to keep the merely dispositional reading of '*tasty*$''$' in (2a), and posit the direct experience requirement as a domain restriction on the stimulus argument of the predicate itself (cf. Pearson 2013: 122, ex. 45).

(9) $[\![\text{tasty}]\!]^w = \lambda x_e.\lambda y_e : y$ has tasted x at $w.tasty''(w)(x)(y)$

From (9) it follows that *tasty to Alfonse* denotes a partial function, defined only on individuals that Alfonse has tasted. The domain condition must be lexically stipulated, and

vary from predicate to predicate, since different experiential predicates enforce different sorts of direct experience presuppositions: *funny to Alfonse* has no tasting requirement, but does enforce a direct experience presupposition of some sort.

Either approach gets the right result, but commits the contextualist to a substantive claim: all relations denoted by experiential predicates enforce direct experience presuppositions. If this is denied, problems arise in capturing the presuppositions; if it's affirmed, problems arise in contexts in which there are no such presuppositions (cf. §1.1.2).

The question is then how the experiencer argument is saturated when there is no overt experiencer PP, as in (10).

(10) Licorice is tasty.

The contextualist's eponymous answer is that its value is supplied by the context: an occurrence of *tasty* as in (10) has as its literal content at a context *tasty to* x, the value of 'x' somehow contextually determined. Absent an overt argument, experiencer predicates take definite null complements, in Fillmore's (1986: 96) terminology, or anaphoric/context-dependent implicit arguments, in Condoravdi & Gawron's (1996: 2-5); cf. Snyder (2013: 286-289).

This argument is 'context-sensitive' in the wide sense of Partee (2004 [1989]), and can depend for its value on the context of utterance alone, the preceding discourse, or the local sentence context, yielding indexical readings, anaphoric readings (Snyder 2013: 288-289), and bound-variable readings controlled by a local quantifier (Lasersohn 2005: 681; 2008: 323-326; 2017: 116-118; Cappelen & Hawthorne 2009: 132-134; Schaffer 2011: 192-195; Snyder 2013: 285).[4]

4. A fourth sort of reading relies, definite description-like, on accommodation based on prior discourse, without a linguistic antecedent: *Whenever there's a trip to the candy store, the licorice is tasty* can mean that the licorice is tasty to the ones who go to the store, despite there being no preceding expression referring to these people (as expected, this reading is unavailable for explicit pronominal experiencers: *# Whenever there's a trip to the candy store, the licorice is tasty to them*; cf. Partee 2004 [1989]: 270). Implicit donkey experiencers are also possible, as in *Everyone who has a child knows which foods are tasty*, which on the relevant reading has the experiencer covarying with children of parents.

Thus, (11a) can be read with the implicit experiencer indexically valued as Alfonse, if uttered while watching him eat the licorice, delighted, while (11b) has a reading on which the implicit experiencer of *tasty* is anaphoric to *Alfonse*; readings like these that anchor the predicate to a contextually specified experiencer are collectively called *exocentric*. (11c) has a covarying reading on which for every person x such that x went to the candy store, there is a snack y such that x got y and y is tasty to x.

(11) a. The licorice is tasty.

b. Alfonse went to a candy store. The licorice was tasty.

c. Everyone who went to the candy store got a tasty treat.

These interpretations can be captured if the experiential predicate can also select for an empty DP x_ϵ, which acts as a variable valued for an experiencer individual.[5] A semantics for implicit experiencers can then be provided by evaluating denotations relative to a variable assignment g, which maps indices, including the distinguished experiencer index ϵ, to values of the relevant type (here, type e).[6]

(12) $[\![x_\epsilon]\!]^{w,g} = g(\epsilon)$

5. If the contextualist doesn't take this route, it's a non-trivial question how the implicit experiencers are going to work. As Schaffer (2011: 181-191) shows, there is a range of possibilities in logical space, but not all of these are guaranteed to be linguistically plausible. According to the above exposition, the contextualist is committed to experiential adjectives being transitive: giving this up requires another story about how they compose with overt experiencer PPs (cf. §1.1.2). But if the predicates are transitive with overt experiencers, then to deny the same of them with covert experiencers would be to posit a lexical ambiguity, or to require some mechanism for suppressing their selectional behavior. Of course, the syntactic status of implicit arguments generally is a matter of debate: see Bhatt & Pancheva (2006) for an overview (and §5 for a note on 'evaluative predicates'), and Landau (2010b) for the position that implicit arguments are genuinely projected, despite their impoverished syntax. What follows assumes for convenience that implicit experiencers are phonologically null DPs selected for in the ordinary way: something similar is taken for granted in early generative treatments, as in Lakoff (1970: 127), which assumes that implicit experiencers are the result of a transformational deletion. Epstein (1984) argues for the syntactic reality at LF of these experiencer arguments based on the control of PRO by clause-taking experiential predicates like *fun*, and Schaffer (2011: §2.3) and Snyder (2013: 287-288) have revived the claim in the present debate (cf. Lasersohn 2017: 118-119 for a skeptical response). Collins (2013) has argued that experiencer phrases are syntactically adjunctive.

6. This notion is adapted from Stojanovic's (2007: 700) idea of a 'distinguished judge variable' x_T. An alternative in Stephenson (2007a: 503-505) instead posits a range of silent experiencer proforms acting as individual constants; this won't work, since a variable is required to capture quantified readings.

The discourse or utterance context then provides a value for this index. A semantic-pragmatic bridge principle can determine the content of assignment-dependent expressions in a context of utterance c by quantifying over assignments compatible with that context (13b), as restricted by (13a), which ensures that the value of ϵ is ϵ_c, the experiencer relevant in c. κ is a variable over semantic contents, viz. objects of type $\langle s, \tau \rangle$, where τ is an extensional type (so that the semantic contents of sentences relative to contexts are traditional propositions, of type $\langle s, t \rangle$).

(13) a. g is compatible with c only if $g(\epsilon) = \epsilon_c$.

b. $(\!|\alpha|\!)^c = \iota\kappa[\forall g : g \text{ is compatible with } c[\kappa = \lambda w_s.[\![\alpha]\!]^{w,g}]]$

This ensures that unbound occurrences of the implicit experiencer with no discourse antecedent behave indexically. (11a) on the relevant interpretation is interpreted as in (14a); its content in a context c is then as in (14b). Where ϵ_c is Alfonse, the proposition expressed is identical to that expressed by *the licorice is tasty to Alfonse* in c, yielding the right result.

(14) a. $[\![[\text{the licorice}]\ [\text{is}\ [\text{tasty}\ x_\epsilon]]]\!]^{w,g} = tasty''(w)(g(\epsilon))(\iota x[licorice'(w)(x)])$

b. $(\!|\text{the licorice is tasty } x_\epsilon|\!)^c = \lambda w_s.tasty''(w)(\epsilon_c)(\iota x[licorice'(w)(x)])$

There is then nothing problematic in principle about discourse-anaphoric interpretations, like (11b). These can be handled so long as (i) the use of a proper name like *Alfonse* introduces a discourse referent, and (ii) contextual parameters can be equated with discourse referents according to pragmatic principles that become preferred in the right discourse structure, such that e.g. ϵ_c = Alfonse following a use of *Alfonse* (see Condoravdi & Gawron 1996: 14, ex. 18c-d for an example of how to implement such principles in a dynamic framework, using locative anaphoric implicit arguments).

Bound readings as in (11c) are then derived so long as the ϵ index can be abstracted over, or quantificational expressions can quantify into open sentences containing it and bind it off. Assuming a quantifying-in procedure in the spirit of Montague (1973), interpreting traces in the standard way (15a), and treating quantificational DPs as ordinary indexed universal and

existential quantifiers as in (15b)-(15c) (where a is the type of an assignment function, i.e. a function from typed indices to objects in the domain, and intensions are type $\langle a, \langle s, \tau \rangle \rangle$), (16a) can be assigned a structure as in (16b), which is interpreted as in (17), yielding the desired reading.

(15) a. $[\![e_i]\!]^{w,g} = g(i)$

b. $[\![\text{everyone}_i]\!]^{w,g} = \lambda\phi_{a,st}.\forall y[person'(w)(y) \rightarrow \phi(g[i \rightarrow y])(w)]$

c. $[\![\text{a}_i]\!]^{w,g} = \lambda P_{et}.\lambda\phi_{a,st}.\exists x[P(x) \wedge \phi(g[i \rightarrow x])(w)]$

(16) a. Everyone got a tasty treat.

b.

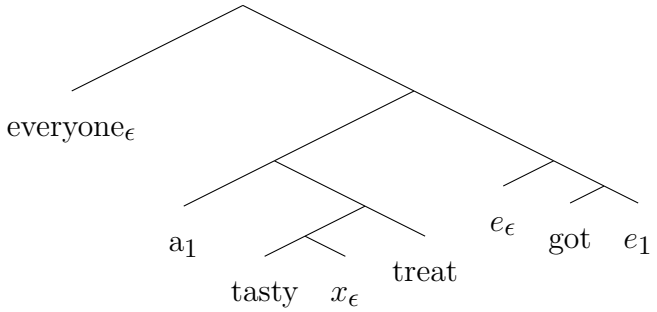

(17) a. $[\![e_\epsilon\ [\text{got}\ e_1]]\!]^{w,g} = get''(w)(g(1))(g(\epsilon))$

b. $[\![[\text{tasty}\ x_\epsilon]\ \text{treat}]\!]^{w,g} = \lambda x_e.tasty''(w)(g(\epsilon))(x) \wedge treat'(w)(x)$

c. $[\![\text{a}_1\ [[\text{tasty}\ x_\epsilon]\ \text{treat}]]\!]^{w,g} =$
$\lambda\phi_{a,st}.\exists x[tasty''(w)(g(\epsilon))(x) \wedge treat'(w)(x) \wedge \phi(g[1 \rightarrow x])(w)]$

d. $[\![[\text{a}_1\ [[\text{tasty}\ x_\epsilon]\ \text{treat}]]\ [e_\epsilon\ [\text{got}\ e_1]]]\!]^{w,g} =$
$\exists x[tasty''(w)(g(\epsilon))(x) \wedge treat'(w)(x) \wedge get''(w)(x)(g(\epsilon))]$

e. $[\![\text{everyone}_\epsilon\ [[\text{a}_1\ [[\text{tasty}\ x_\epsilon]\ \text{treat}]]\ [e_\epsilon\ [\text{got}\ e_1]]]]\!]^{w,g} =$
$\forall y[person'(w)(y) \rightarrow \exists x[tasty''(w)(y)(x) \wedge treat'(w)(x) \wedge get''(w)(x)(y)]]$

Finally, some bound readings relate the bound individuals not to their own experiential standard, but to a standard determined by some individual to which they bear some contextually salient relation. Thus in a context where parents are buying candy for their kids, (18) has a reading on which the parents buy something tasty not to themselves, but to their children.

(18) Every parent got a tasty treat.

These can be captured by allowing x_ϵ to denote not the value that the assignment gives to ϵ, but the result of applying some function to this. Let the assignment also determine such a function, so that $g(f_\epsilon)$ is a function of type $\langle s, \langle e, e\rangle\rangle$, which maps worlds w to individuals x to individuals y such that y bears a relation governing experiencers to x at w. The context can then determine the relevant relation as in (19a), where $f_{c,\epsilon}$ is the contextually relevant function in c governing experiencers, and the denotation for x_ϵ can be updated as in (19c), which is an abbreviation per the definition in (19b).

(19) a. g is compatible with c only if $g(f_\epsilon) = f_{c,\epsilon}$.

b. $g_{f,w}(\epsilon) := g(f_\epsilon)(w)(g(\epsilon))$

c. $[\![x_\epsilon]\!]^{w,g} = g_{f,w}(\epsilon)$

In the relevant context c, (18) is derived because $f_{c,\epsilon}$ is that function which maps worlds w to individuals x to the (relevant) child(ren) of x at w. Composing analogously to (17) then gets the right result relative to this context, and the content of this expression relative to c is as in (20), where '$child_c(w)(y)$' denotes the contextually relevant child(ren) of y at w in c.

(20) $(\!|\text{every}_\epsilon \text{ parent a}_1 \text{ tasty } x_\epsilon \text{ treat } e_\epsilon \text{ got } e_1|\!)^c =$
$\lambda w_s.\forall y[parent'(w)(y) \rightarrow \exists x[tasty''(w)(child_c(w)(y))(x)$
$\wedge treat'(w)(x) \wedge get''(w)(x)(y)]]$

The original bound reading, on which people get treats tasty to themselves, then arises when $f_{c,\epsilon}$ is trivial, i.e. denotes the function from worlds to the identity function on individuals. This is the normal case, and in the following, the simpler denotation for x_ϵ in (12) is assumed where no special non-identity function is being invoked.

Experiential predicates are unlike some lexical items that take anaphoric implicit arguments, in that they can appear freely with no overt PP even in contexts that provide no

antecedent of any kind to which the argument can be anaphoric: (10), for instance, is felicitous out of the blue, with no experiencer made previously salient, no discourse antecedent, no possibility of accommodation, and no local controlling quantifier. In such cases, uses of experiential predicates tend to have to do with the speaker's own experiential standard. In this way they pattern like *local*, whose location argument defaults to the location of the speaker absent any antecedent.

This is intelligible if, as Condoravdi & Gawron (1996: 20) claim, each context determines automatically values for a small range of parameters inherently tied to the concrete circumstances surrounding speech acts: where values of implicit arguments can in principle adopt these values, they will do so absent an antecedent (and will be infelicitous if they cannot). Thus the location of the speech act provides a suitable sort of argument for *local*, and since sentient individuals are a suitable sort of argument for experiential predicates, it can follow that the value of ϵ_c defaults to the speaker of the context.

The contextualist can thus round out the semantic-pragmatic picture with a constraint like the following, where s_c is the speaker in c (or that group, trivially including the speaker, whose experiences are similar to the speaker's in the contextually relevant way).

(21) If there is no salient, antecedent, or accommodable experiencer in c, then $\epsilon_c = s_c$.

§1.1.5 shows that this principle is problematic; but with it the contextualist has a coherent picture of a range of uses of experiential predicates.

The following sections argue that contextualism is fundamentally flawed, and ought to be abandoned. §§1.1.1-1.1.4 show that experiential predicates don't behave the way that the contextualist predicts in a huge array of embedded contexts: §1.1.1 deals with subjective attitude verbs; §1.1.2 with a wide range of intensional contexts, and the presuppositions of direct experience they do or don't impose; §1.1.3 with factive verbs, recapitulating and expanding on observations by Lasersohn (2009); and §1.1.4, with reports of agreement and disagreement, expanding on observations by Cappelen & Hawthorne (2009). §1.1.5, following Ninan (2014), addresses the pragmatics of using experiential predicates to assert, and argues

that a contextualist semantics fails to account for the range of commitments these predicates involve with regard to direct experience: in particular, it doesn't distinguish between direct experience presuppositions and so-called acquaintance inferences.

The data point to a distinction between 'bare' uses of experiential predicates, which make no reference in their content to a specific experiencer, and exocentric uses that do involve such reference. The contextualist can account only for the latter, and so the view is too semantically impoverished to cover the range of uses that these predicates have.

1.1.1 Argument I: subjective attitudes

A contextualist semantics for experiential predicates encounters problems with treating the interpretation of subjective attitude verbs. This section demonstrates this with two English examples, *find* and *consider*.

Find is a verb that embeds (preferably adjectival) small clauses, and places certain restrictions on what can appear in the predicative position of these clauses. The nature of these restrictions is a matter of debate (cf. Stephenson 2007b: 59-62; Sæbø 2009; Bouchard 2012: ch. 3; Kennedy 2013; Hirvonen 2014: ch. 4; Kennedy & Willer 2016; Vardomskaya 2018: ch. 4), but for present purposes only the gross difference in (22) is relevant: *find* felicitously embeds experiential predicates (22a), but not predicates that are non-experiential and in no obvious way 'evaluative' (22b).

(22) a. Alfonse finds licorice tasty.

b. #Alfonse finds licorice vegetarian.

What is unexpected on a contextualist semantics is that experiential predicates with overt experiencer PPs pattern like *vegetarian*, and unlike experiential predicates with no overt PP.

(23) a. #Alfonse finds licorice tasty to Bethany.

b. #Alfonse finds licorice tasty to himself.

Since the contextualist takes occurrences of experiential predicates with overt experiencer PPs to be identical in content at a context to occurrences of those same predicates without these overt PPs (where the implicit experiencer's value is the same as that of the preposition's complement), it's not clear how this difference arises: where *find* enforces semantic restrictions on predicates that it embeds, it treats experiential predicates as semantically distinct depending on whether or not an overt experiencer PP is present, against the contextualist's prediction.[7]

The problem can't be that *find* semantically requires a (*de re*) self-directed attitude, or that the subject of the verb somehow controls the experiencer argument to the predicate: (23b) remains unacceptable, even though the subject and experiencer co-refer.[8] The contextualist thus looks to lack an explanation for *find*'s embedding behavior.

Sæbø (2009: 336-337) offers a treatment of *find* that attempts to explain these restrictions in a way compatible with a contextualist semantics, according to which the verb is semantically vacuous, denoting the identity function on properties of experiencers, as in (24).[9] The subject of the attitude then saturates the experiencer argument of the experiential predicate, so that (25a) is given a structure as in (25b), and interpreted as in (26). This ensures that *find* only selects for clauses denoting properties of experiencers, which can only occur with

7. Might there instead be a syntactic, rather than semantic, reason for *find*'s embedding restrictions? It's hard to see how: since in (22a) and (23) the predicate itself is the same, and heads the small clause, the verb's selectional restrictions are unlikely to be sensitive to this distinction. What would be required is that (i) overt experiencer PPs are syntactically distinct from their covert counterparts; and (ii) the syntactic behavior of the verb is sensitive to this distinction regarding what the predicate that it embeds selects for (perhaps preventing embedding of clauses whose predicative component has a complement). Even if such a syntactic restriction can be formulated, the restrictive behavior of *find* appears not to be syntactically driven in that it cross-cuts the class of predicates without (overt) complements anyway, as in the distinction between *tasty* and *vegetarian*: plausibly, this behavior is driven by the lexical semantics of each adjective.

8. The possibility that the verb requires a *de se* self-directed attitude, possibly contributed by a null proform that differs in semantic function from overt pronouns that enforce *de re* readings, is tempting here. But if such *de se* readings are in general allowed with experiential predicates, contextualism effectively collapses into relativism: cf. fn. 35, and the comments on *de se* attitudes in §1.2.1. See also Stephenson (2007a) for a relativist account that makes a similar move, using the null proform 'PRO_J.'

9. The denotation in (24) may not be sufficient to capture *find*'s restrictions on the embedding of predicates, since the type $\langle e, t\rangle$ isn't fine-grained enough to distinguish properties of experiencers specifically from properties of individuals generally.

experiential predicates.[10]

(24) $[\![\text{find}]\!]^{w,g} = \lambda P_{\langle e,t\rangle}.P$

(25) a. Alfonse finds licorice tasty.

b. 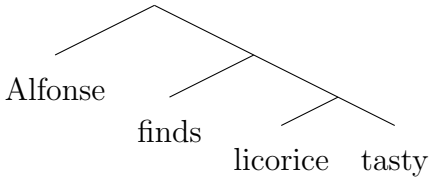

(26) a. $[\![\text{tasty}]\!]^{w,g} = \lambda x_e.\lambda y_e.tasty''(w)(y)(x)$

b. $[\![\text{licorice tasty}]\!]^{w,g} = \lambda y_e.tasty''(w)(y)(l)$

c. $[\![\text{finds [licorice tasty]}]\!]^{w,g} = \lambda y_e.tasty''(w)(y)(l)$

d. $[\![\text{Alfonse [finds [licorice tasty]]}]\!]^{w,g} = tasty''(w)(a)(l)$

This approach entails that the denotation of *tasty* as it occurs in (25b) isn't the standard contextualist one in (1): it composes with the stimulus prior to the experiencer, as in (26a) (where '$tasty''$' is to be read as before).

It follows that this occurrence of *tasty* can't compose with an internal experiencer argument (whatever this argument's syntactic status), for two reasons. First, the composition of *find*-constructions would fail due to type-mismatch, as in (27). Second, this would force *tasty*'s denotation to compose with the experiencer first, getting the wrong interpretation in unembedded uses, as in (28a), which is true just in case Alfonse is tasty to licorice (28b).

10. (25a) gives rise to a direct experience presupposition, much in the way that *Licorice is tasty to Alfonse* does. On this account, this is taken care of automatically, so long as one of the two moves securing these presuppositions for the contextualist mentioned in §1.1 is adopted, since *Alfonse* performs the exact same function as the subject of *find* that it performs as part of the internal argument to *tasty*. Sæbø's own account is broader than its exposition here recognizes: he wants an explanation for the embedding of 'subjective' predicates generally, which may be a broader class than that of experiential predicates.

(27)

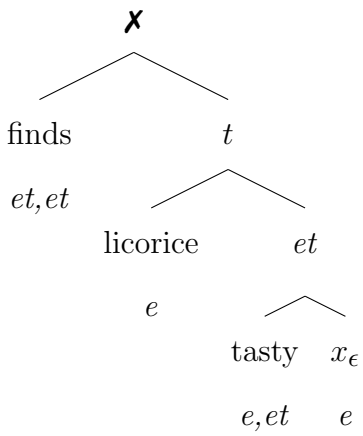

(28) a. Licorice is tasty to Alfonse.

b. $[\![\text{licorice [is [tasty [to Alfonse]]]}]\!]^{w,g} = tasty''(w)(l)(a)$

Sæbø's account thus faces a dilemma in trying to treat both (25a) and (28a): either *tasty* is semantically uniform, and always composes with a stimulus prior to an experiencer, or it is ambiguous, having one denotation (for *find*-reports) with no internal experiencer argument, and another denotation (for elsewhere) that does. The former option leads to a constituent structure for (28a) as in (29).

(29)

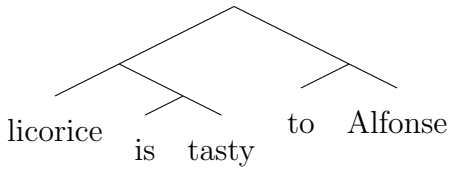

This makes two unwelcome predictions: that *tasty to Alfonse* is not a constituent, and has no denotation assigned to it, and that the predicative copula doesn't compose with a constituent of extensional type $\langle e, t\rangle$.[11]

The latter option effectively amounts to positing a special lexical item for each experiential predicate whose sole purpose is to compose with *find*. Even if two separate lexical entries e.g. for *tasty* are posited, one that composes with an experiencer first, and one that composes with a stimulus first, some restriction has to be enforced for the stimulus-first entry

11. Taking *tasty to Alfonse* not to be a constituent will undoubtedly cause numerous problems, but to name just one, it makes it mysterious how its use as a restrictive nominal modifier, as in *the licorice tasty to Alfonse*, is to be composed.

only to appear beneath *find* as desired. This move effectively abandons a locally compositional account of experiential predicates in these attitude constructions, and commits to the semantics of the predicate differing depending on the environment in which it's embedded.[12]

Perhaps a proposal in Sæbø's spirit can circumvent these problems with minor modifications. (25a) can be given a structure as in (30), consistent with the semantics given in §1.1, such that *tasty* takes a covert experiencer argument whose value is determined by an assignment as usual. The denotation of *find* can then be as in (31a): it serves as an intensional operator that shifts the assignment relative to which the clause it embeds is evaluated, mapping the distinguished experiencer index to the attitude holder. The resulting interpretation is as in (31b): this is the same as the value in (26d).

(30)

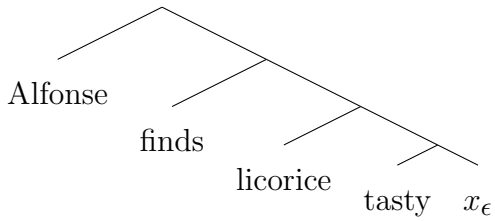

(31) a. $[\![\text{find}]\!]^{w,g} = \lambda\phi_{a,st}.\lambda x_e.\phi(g[\epsilon \to x])(w)$

b. $[\![\text{Alfonse [finds [licorice [tasty } x_\epsilon\text{]]]]}]\!]^{w,g} = tasty''(w)(a)(l)$

The infelicity of *find* with non-experiential, non-evaluative predicates, and with predicates occurring with overt experiencer PPs, would then result from the fact that neither of these contain x_ϵ, making the attitude vacuous, there being no value of ϵ to non-trivially shift (cf. Katzir & Singh 2013, esp. ex. 27 regarding bans on semantically vacuous operators).

But this approach ultimately comes at a high cost: its explanation for the felicity of *find* relies on the presence of a covert experiencer, tying the occurrence of *find* solely to experiential predicates, while the verb has a much wider application than this, and occurs embedding any predicate for whose conditions of application the subject can have a distinguished sort of direct evidence (cf. §5.3.1 & Ch. 5, fn. 46), enforcing a direct experience

12. Sæbø (2009: 339, exx. 25a-b) effectively commits to the ambiguity strategy, though for slightly different reasons, positing one entry for each experiential predicate that composes with a stimulus and then an experiencer, and another that has the experiencer argument saturated '*a priori*' in the predicate itself, and composes only with a stimulus.

requirement as it does so (cf. fn. 10). The contextualist therefore must posit a vast array of such implicit experiential arguments, accompanying everything from *easy* to *helpful* to *decent*, and additionally must take them to be projected by any construction that systematically allows *find*-embedding, as with the degree morpheme *too* (as in, *Alfonse finds the museum too modern*).[13]

Consider embeds a small clause, and enforces a certain special doxastic relation between an agent and the proposition denoted by this clause: very roughly, it requires that the agent have evaluative, as opposed to merely descriptive, beliefs that verify the proposition (cf. Lasersohn 2009: 365-367; Kennedy 2013: 265-266; Kennedy & Willer 2016).

For predicates that are intuitively evaluative to begin with, like *tasty*, the semantic effect of *consider* is often difficult to distinguish from that of ordinary doxastic verbs like *think* (32). Where a predicate is intuitively non-evaluative, embedding beneath *consider* is felicitous, but enforces a distinct sort of belief from *think*: (33b) may report a factual belief that Alfonse has about the material makeup of the table, while (33a) must report some sort of evaluative belief regarding that makeup (e.g. that Alfonse takes its material to count as wood in a case where this is disputable). In cases where it's difficult to determine what sort evaluative belief could possibly be held toward a predicate's application, embedding this predicate beneath *consider* is odd (34) (cf. Lasersohn 2009: 366, ex. 9b).[14]

(32) a. Alfonse considers licorice tasty.

b. Alfonse thinks that licorice is tasty.

(33) a. Alfonse considers the table wooden.

b. Alfonse thinks that the table is wooden.

13. If the contextualist balks at this proliferation of experiential arguments, and instead incorporates the variable-sensitivity that *find* shifts in the lexical semantics of the predicate itself, then the problems to be addressed in §1.1.2 will arise for *find* just as they do with *to*-headed experiencer PPs.

14. Some interpretation could be coerced to repair (34): say, that Bethany is just a quarter-inch shy of six feet tall, and Alfonse is willing to take her to be six feet tall for some intents and purposes. Thus embedding the predicate beneath *consider* is felicitous precisely to the extent that such an evaluative construal of the belief can be accommodated (cf. Vardomskaya 2018: §4.7 for complications with *consider*).

(34) ?Alfonse considers Bethany six feet tall.

What is unexpected on a contextualist semantics is that embedding small clauses beneath *consider* with experiential predicates in predicative position results in very different interpretations depending on whether an overt experiencer PP is present.

(35) a. Alfonse considers licorice tasty.

b. Alfonse considers licorice tasty to Bethany.

c. Alfonse considers licorice tasty to himself.

The most salient reading of (35a) reports Alfonse's opinion regarding the taste of licorice. (35b) and (35c) can't be interpreted in this way: to the extent that they have interpretations, they must mean that Alfonse has an evaluative belief regarding whether, given Bethany's or his own experiential reactions to licorice, licorice counts as tasty to the relevant experiencer by those lights (e.g. if Alfonse's opinion regards how tasty licorice must be to the experiencer to count as tasty to them, or what experiential reactions should count as relevant for *tasty*; cf. *ibid.*, fn. 4).

In other words, there is a reading of (35a) that is not equivalent to any reading of a sentence of the form *Alfonse considers licorice tasty to* x, for any x (not even, as (35c) shows, where x is Alfonse himself). But the contextualist wrongly predicts that this latter interpretation is how (35a) ought to be read, for some value x of the implicit experiencer. Insofar as *consider* differs from *think* depending on the 'evaluativity' of the predicate it embeds, it treats experiential predicates without overt experiencer PPs as evaluative, but not those with overt PPs (which therefore require a special evaluative reading of the belief that *think* doesn't track, as with *vegetarian*). *Consider*-reports thus do not mean what a contextualist semantics predicts that they do.

1.1.2 Argument II: direct experience presuppositions

§1.1 showed that the contextualist can capture direct experience presuppositions like that in (36) by interpreting the metalanguage predicate '*tasty*$''$' so as to make reference to the gustatory experience that the stimulus produces in the experiencer (37a) (repeated from (8)), or by imposing a domain restriction on the predicate itself (37b) (repeated from (9)).

(36) Licorice is tasty to Alfonse.

$\hookrightarrow$ Alfonse has tasted licorice.

(37) a. $tasty''(w)(x)(y)$ iff:

the gustatory experience that y produces in x at w is pleasant.

b. $[\![\text{tasty}]\!]^w = \lambda x_e.\lambda y_e : y$ has tasted x at $w.tasty''(w)(x)(y)$

Recall that the contextualist machinery as it stands commits to the claim that all uses of an experiential predicate to relate experiencers and stimuli enforce such presuppositions. This is a problem, since in many intensional contexts these presuppositions don't arise where expected if this were the case.

In the prejacents of epistemic necessity modals (38a), in the antecedents of conditionals (38b), and in questions, both direct (38c) and indirect (38d), the direct experience presupposition survives when the experiential predicate is accompanied by an overt experiencer PP.

(38) a. The licorice must be tasty to Alfonse.

$\leftrightarrow_{acc}$ It must be that: Alfonse has tasted the licorice, and likes its taste.

b. If the licorice is tasty to Alfonse, he should buy some.

$\leftrightarrow_{acc}$ If Alfonse has tasted the licorice and likes its taste, he should buy some.

c. Is the licorice tasty to Alfonse?

$\leftrightarrow_{acc}$ Is it the case that: Alfonse has tasted the licorice and likes its taste?

d. I wonder whether the licorice is tasty to Alfonse.

$\leftrightarrow_{acc}$ The speaker wonders whether: Alfonse has tasted the licorice, and likes its

taste.

$\hookrightarrow$ Alfonse has tasted the licorice.

The hooked arrow indicates the presupposition, identical to that in (36), that projects in each example when it isn't locally accommodated in the relevant intensional context. For each of (38a)-(38d), there is further a paraphrase, indicated by '$\leftrightarrow_{acc}$,' that arises when the presupposition is locally accommodated. In these cases, the presupposition must be satisfied relative to the shifted point of evaluation: for (38a), in those worlds compatible with some base of information, etc.

The same effect occurs beneath predictive operators (39a) and in belief reports (39b), where local accommodation of the presupposition is obligatory.

(39) a. The licorice will be tasty to Alfonse.

$\leftrightarrow_{acc}$ It will be that: Alfonse has tasted the licorice, and likes its taste.

b. Alfonse believes that the licorice is tasty to Bethany.

$\leftrightarrow_{acc}$ Alfonse believes that: Bethany has tasted the licorice, and likes its taste.

But when experiential predicates lack overt experiencer PPs, the same isn't true, and there are readings for counterparts to each sentence in (38)-(39) with no direct experience presupposition, locally accommodated or otherwise.

(40) a. The licorice must be tasty.

b. If the licorice is tasty, he should buy some.

c. Is the licorice tasty?

d. I wonder whether the licorice is tasty.

(41) a. The licorice will be tasty.

b. Alfonse believes that the licorice is tasty.

All of (40a)-(40d) can be said felicitously with respect to a newly made batch (or kind) of candy that it's mutually known no one has ever tasted, and where what must be, what is

questioned, etc. is not that or whether anyone has ever tasted it.[15] If (41a) is uttered while the licorice is being made, it can be felicitously followed with *It's too bad no one will ever taste it* (cf. Klecha 2014: 451): compare #*The licorice will be tasty to Alfonse; it's too bad no one will ever taste it.* Alfonse's belief as reported in (41b) can be formed with respect to that same batch, even where Alfonse knows it to be untasted.

But the contextualist semantics requires in each of (40a)-(41b) that *tasty* have its experiencer argument saturated implicitly; since the semantic process by which this happens is identical to the one that saturates the argument overtly, the same direct experience presuppositions are predicted to arise with respect to the contextually relevant experiencers. Thus no implicit value for x_ϵ yields the right results, and the contextualist semantics makes the wrong predictions for all of these constructions.

The contextualist thus needs a way to distinguish between those cases in which direct experience presuppositions do and don't arise. One possibility is to enforce the presuppositions via the overt experiencer PPs, and not via implicitly supplied experiencers: this may account for the distinction between (38)-(39) and (40)-(41). Suppose that '*tasty*$''$' retains a merely dispositional interpretation (42a) (repeated from (2a)), and the preposition enforces a domain restriction requiring direct experience (42b).

(42) a. $tasty''(w)(x)(y)$ iff at w, y is disposed to produce gustatory pleasure in x

b. $[\![\text{to}]\!]^{w,g} = \lambda x_e : x$ has direct experience of the relevant kind at $w.x$

This won't work: even if a notion of 'direct experience of the relevant kind' can be articulated (perhaps relative to a context), due to the locality of composition, the experiencer PP is unable to enforce (i) what the experiencer must have direct experience of, since the PP argument has no access to the stimulus, and (ii) what sort of direct experience the

15. A direct question like (40c) sometimes presupposes that the addressee has tasted the licorice: cf. Lasersohn (2005: 673-674). This is because when 'asking is an invitation to assert' (*ibid.*), assertions of expected answers to the question trigger acquaintance inferences (see §1.1.5), and the onus on the addressee to have tasted the licorice arises. This parallels 'interrogative flip' in languages with grammaticized direct evidentiality: cf. §4.2.1.

experiencer must have, since it has no access to the predicate.

The licorice is tasty to Alfonse enforces grammatically, and not merely contextually, that Alfonse has experience specifically of the licorice (due to *the licorice*), and that this experience is specifically gustatory (due to *tasty* – again compare *funny*, which has no tasting requirement). An entry like (42b) is thus unable to provide a coherent method of composing direct experience presuppositions.

This might be remedied if the contextualist takes overt experiencer PPs to be modifiers of, rather than arguments to, the predicate. This requires that the experiencer argument of '*tasty″*' be saturated in the lexical semantics of the predicate itself (43a), and that the experiencer PP be an intensional shifter of this standard (43b), where $g[\epsilon \to x]$ is that assignment just like g except that $g[\epsilon \to x](\epsilon) = x$.[16] *Tasty to Alfonse* is then as in (43c), supposing that 'direct experience of the kind relevant' to the intension of *tasty* is gustatory experience.

(43) a. $[\![\text{tasty}]\!]^{w,g} = \lambda x_e.tasty''(w)(g(\epsilon))(x)$

b. $[\![\text{to}]\!]^{w,g} = \lambda x_e.\lambda P_{a,\langle s,et\rangle}.$
$\lambda y_e : x$ has direct experience of y of the kind relevant to P at w.
$P(g[\epsilon \to x])(w)(y)$

c. $[\![\text{tasty [to Alfonse]}]\!]^{w,g} = \lambda y_e : a$ has tasted y at $w.tasty''(w)(a)(y)$

This analysis allows the direct experience presupposition to compose, because a modifier to the predicate can 'see' both the stimulus (y) and the predicate (P) to enforce the appropriate conditions on them. Retaining the dispositional interpretation of '*tasty″*' from (42a), it follows that where there is no overt experiencer PP, there is no direct experience presupposition, only an asserted dispositional requirement, as desired.[17]

16. This echoes Kölbel's (2009: 394) proposal that the contextualist might treat such phrases as monstrous operators. Here there are no context shifts, only assignment shifts, as with traditional predicate abstraction or quantifiers, sidestepping objections to monsters.

17. It seems no contextualist has explicitly advocated this modifier view. Glanzberg (2007: 13, ex. 15a) has a denotation similar to (43a), and Sæbø (2009: 331) takes him at his word, but Glanzberg (2007: 11-12,

This analysis doesn't get the right results, since it predicts that the distinction between experiential predicates enforcing and not enforcing direct experience presuppositions patterns exactly with whether they occur with overt experiencer PPs or not. But this isn't so: in exocentric uses of experiential predicates (cf. §1.1), the presupposition persists in the absence of any overt PP. Thus if (44) (repeated from (11a)) is uttered when picking out a gift to get for Alfonse from the candy shop, it has a reading that presupposes that Alfonse has tasted the licorice.[18]

(44) The licorice is tasty.

↪ Alfonse has tasted the licorice.

The reader can confirm that this presupposition projects or is locally accommodated in all the constructions described above, if the relevant context is held fixed.

Perhaps the contextualist can appeal to another silent shifter of the experiencer, that behaves like an experiencer PP but is covert, and is inserted into the structure in exocentric contexts to guarantee the direct experience presupposition. This can be done by reimagining x_ϵ as a modifier, as follows.

(45) a. $[\![x_\epsilon]\!]^{w,g} = \lambda P_{a,\langle s,et\rangle}.$

$\lambda y_e : g(\epsilon)$ has direct experience of y of the kind relevant to P at $w.P(g)(w)(y)$

b. $[\![\text{tasty } x_\epsilon]\!]^{w,g} = \lambda y_e : g(\epsilon)$ has tasted y at $w.tasty''(w)(g(\epsilon))(y)$

But this also fails: it predicts that the silent modifier can appear or not, and so that regardless of the value of $g(\epsilon)$, the direct experience presupposition ought to be able to be

fn. 9) makes clear this is not his preferred position. Sæbø (2009: 337, 339) posits that the experiencer must sometimes be saturated '*a priori*' and have no syntactic reality, suggesting a similar analysis, but he does this only to patch a technical problem with his analysis of the verb *find*, which was rejected in §1.1.1.

18. (44) can be uttered, and sincerely agreed to by all the interlocutors that are gift-shopping, even if none of them like the licorice themselves, but are only assessing based on their knowledge of Alfonse's tastes. Thus it's to be read as anchored to Alfonse specifically, and not as a general recommendation to get the licorice because, since it's tasty *simpliciter* (or to the interlocutors), Alfonse is likely to enjoy it too. If getting this reading is difficult, suppose that before (44) is uttered, another of the interlocutors asks, *What does Alfonse like?*, with (44) being a direct response to the question. These exocentric readings are brought out more clearly in 'cat food' examples, where the interlocutors are aren't likely candidates for the relevant experiencer, but some distinct individual or group is: cf. Stephenson (2007a: 499, 504).

enforced or not. This isn't how it works: where the contextually anchored experiencer is e.g. Alfonse, the direct experience presupposition is obligatory.[19]

The problem that a contextualist semantics faces is that the data show two kinds of occurrences of experiential predicates: first, there are predicates in exocentric contexts or occurring alongside an overt experiencer PP, and which simultaneously introduce an experiencer into their content and enforce a direct experience presupposition; second, there are 'bare' occurrences, which do neither. The contextualist, requiring that the experiencer always be present in the content, is unable to capture this pattern.

1.1.3 Argument III: factives

Lasersohn (2009: 369-372; 2017: 158-161) presents a criticism of contextualist semantics on the basis of the behavior of factive attitude reports; this section briefly expands on it.[20]

Factive verbs like *realize* unproblematically embed clauses that look to predicate experiential properties of individuals, as in (46a), and to presuppose the truth of their complements in an ordinary way: compare (46a) to (46b).

(46) a. Alfonse realizes that licorice is tasty.

↪ Licorice is tasty.

b. Alfonse realizes that licorice is vegetarian.

↪ Licorice is vegetarian.

There are three commitments to keep track of in considering an utterance of (46a): those on the part of the speaker, the addressee, and the subject of the attitude.

19. There are two moves a contextualist might try to make here. The first is to recast (45a), so that it does not merely enforce the direct experience presupposition, but also shifts the value of the variable to that of yet another index (say ϵ'): this would effectively require two contextually relevant experiencers – one intrinsic to the predicate, and one intrinsic to the exocentric modifier – and perhaps only the latter might be able to be valued for individuals non-identical the speaker. This would require tying ϵ_c, the 'non-exocentric' experiencer, unacceptably closely to the speaker, causing problems with capturing conversational disagreement with experiential predicates (cf. §1.1.5). The second is to plead that x_ϵ must occur in certain contexts – say, when ϵ_c is non-identical to the speaker, which is to demand the insertion of a lexical item depending on contextual parameters.

20. Pearson (2013: 113-114) makes some related comments.

First, in uttering (46a), the speaker commits to the truth of the presupposed proposition, and ergo for the assertion to be sincere and felicitous, must take this proposition to be true for conversational purposes. Second, the addressee becomes committed to the presupposed proposition if there is uptake on the illocutionary force of the assertion. Finally, the lexical entry for *realize* enforces commitment on the part of Alfonse, the attitude's subject: Alfonse can't realize that licorice is tasty without believing this.

These three commitments can be enforced by a denotation like (47b), where belief is characterized by doxastic alternatives defined as usual (47a).[21]

(47) a. $Dox_{x,w} := \{w' : w' \text{ is compatible with } x\text{'s beliefs at } w\}$

b. $[\![\text{realize}]\!]^{w,g} = \lambda\phi_{a,\langle s,t\rangle} : \phi(g)(w).\lambda x_e.\forall w' \in Dox_{x,w}[\phi(g)(w')]$

(47b) composes with a proposition and an individual (agent), and commits the former to belief in the latter. The domain condition on the first argument enforces the presupposition that the proposition denoted by its complement is true (at the assignment and world of evaluation).

Suppose that (46a) is uttered in a situation in which both the speaker and Alfonse like the taste of licorice, and both the speaker and addressee know this. Further, Alfonse mistakenly believes that the speaker is repulsed by the taste of licorice. In such a situation, (46a) is felicitous, and if the addressee also likes the taste of licorice, the speaker and addressee can competently agree to it. Further, on its being uttered and accepted into the common ground, the speaker, the addressee, and Alfonse, in virtue of being committed to the embedded proposition, typically become committed to liking (or being disposed to like) the taste of licorice.

The question is then: which proposition is embedded beneath *realize*, and so presupposed

21. This entry for *realize* is incomplete, since it only requires belief, presupposed to be true, of the attitude's subject. Plausibly *realize* also has an epistemic component, like *know*, which requires that the subject be acquainted with the truth of what's believed in the right way (cf. MacFarlane 2014: 159-160). The domain restriction also only forces adoption of the presupposition within the local intensional context, which is enough for present purposes: it doesn't predict the possibility of full projection in e.g. *Alfonse might realize that licorice is tasty*, but only the local accommodation reading.

by (46a), in this context of utterance? The obvious answer, viz. the proposition *that licorice is tasty* (*simpliciter*), is unavailable to the contextualist, since there is no such proposition expressible by the complement clause. Since there is no overt experiencer PP, there must be a covert experiencer to saturate *tasty*'s first argument: so the proposition is *that licorice is tasty to* ϵ_c, for some value of ϵ_c. But apparently no candidate value can predict the commitments of both the speaker and Alfonse. There are four options:

1. ϵ_c is some group containing Alfonse but not the speaker
   (possibly just Alfonse himself).

2. ϵ_c is some group containing the speaker, but not Alfonse
   (possibly just the speaker herself).

3. ϵ_c is some group containing both Alfonse and the speaker
   (possibly just Alfonse and the speaker themselves).

4. ϵ_c is some group containing neither Alfonse nor the speaker.

None of these get the right result: [1] and [4] fail to commit the speaker to liking the taste of licorice, [2] and [4] fail to commit Alfonse to liking the taste of licorice, and [3] commits Alfonse to thinking that the speaker likes the taste of licorice, contrary to the example.

MacFarlane (2014: 159) suggests that option [3] can be made to work, if the commitment that the speaker and attitude holder share is to the proposition that licorice is tasty to those with tastes relevantly similar to Alfonse: Alfonse can believe this *de dicto*, and trivially will so long as he believes *de re* of himself that he likes licorice; and the speaker, in recognizing that both she and Alfonse like licorice, is committed to this belief as well.

To implement this, x_ϵ (the covert complement to *tasty*) must not world-invariantly denote an individual. This would result in a *de re* reading to the effect that Alfonse believes of the group that shares his tastes that licorice is tasty by its standards: but he has no such belief, since this group includes the speaker, and per the example he believes *de re* of the speaker

(incorrectly) that she doesn't like the taste of licorice. x_ϵ must therefore have a world-relative denotation capable of scoping beneath *realize*, so that Alfonse can have a merely *de dicto* belief about the group that shares his tastes.

This can be accomplished via the machinery introduced in §1.1, which effectively maps the value of ϵ_c in c to some individual standing in a relevant relation to it, as in (48a) (repeated from (19c)): *tasty* in (46a) can then take x_ϵ as its complement, and at a context in which ϵ_c is Alfonse, and $f_{c,\epsilon}$ is that function that maps worlds w to individuals x to the largest group that shares x's tastes at w in the contextually relevant way, the right result is achieved (where MAX$x[\phi]$ is the maximal, possibly plural or kind, individual such that ϕ is true on an assignment of that individual to x).[22]

(48) a. $[\![x_\epsilon]\!]^{w,g} = g_{f,w}(\epsilon)$

b. $[\![$Alfonse [realizes [that [licorice [is [tasty x_ϵ]]]]]$]\!]^{w,g} =$
$\forall w' \in Dox_{a,w}[tasty''(w')(g_{f,w}(\epsilon))(l)]$,
if $tasty''(w')(g_{f,w}(\epsilon))(l)$;
else undefined

c. Where $\epsilon_c =$ Alfonse and
$f_{c,\epsilon} = \lambda w_s.\lambda x_e.MAXy[y$ has tastes similar to x at w in the way relevant for $c]$,
$(\!|$Alfonse realizes that licorice is tasty $x_\epsilon|\!)^c =$
$\forall w' \in Dox_{a,w}[tasty''(w')$
(MAX$y[y$ has tastes similar to a at w in the way relevant for $c])(l)]$,
if $tasty''(w')$(MAX$y[y$ has tastes similar to a at w in the way relevant for $c])(l)$;
else undefined

But this won't work: suppose the example varies slightly so that the addressee does *not*

22. What is 'the contextually relevant way' that the tastes must be shared? It seems in this case it has to be narrow enough that it only tracks tastes with respect to licorice: for (46a) is felicitous even if the speaker, addressee, and Alfonse all realize that, in liking licorice, Alfonse differs from any relevant group one might pick out that *otherwise* shares his tastes. This apparently commits the contextualist to rapid shifts in these fine-grained similarity relations, such that e.g. *Alfonse realizes that licorice is tasty and that marzipan isn't tasty* must track licorice-similarity in the first embedded conjunct and marzipan-similarity in the second.

like the taste of licorice, even though the speaker does. In accepting the assertion in this context, the addressee is still typically committed to liking the taste of licorice as well. But this commitment doesn't follow if the content of (46a) in such a context is that Alfonse realizes that licorice is tasty *to the group with tastes relevantly similar to Alfonse's*: then the addressee in accepting what *this* proposition presupposes need make no commitment regarding her own tastes, since she can felicitously and without pretense commit to licorice being tasty to *that* group, even while not committing to it being tasty to herself.[23]

Thus the situation in which all three characters share a commitment in virtue of the assertion and acceptance of the factive report is inexplicable, and the contextualist has no plausible way of explaining the behavior of factivity with respect to experiential predicates in attitude reports.

This sort of difficulty isn't exclusive to factive verbs, but rather is symptomatic of a general problem with contextualist semantics, involving the tracking of mutual commitment to experiential dispositions in virtue of mutual commitment to the truth of a single content. Factive verbs happen to track such mutual commitment via a lexical item, but periphrastic constructions can be used to argue for the same points. A sentence like *Alfonse thinks licorice is tasty, and it is*, where the attitude report and commitment on the speaker's part are tied to distinct conjuncts, can be used to run the same arguments (cf. Lasersohn 2009: 372 for a similar point about truth-sensitive adverbs like *correctly*).

1.1.4 Argument IV: agreement and disagreement

The verbs *agree* and *disagree* characterize doxastic attitudes that distinct agents hold towards a single content (cf. Cappelen & Hawthorne 2009: 54-67). *Agree*- and *disagree*-reports

23. It does no good to say that the addressee becomes committed to liking the taste of licorice when committing to being in the licorice-liking group, and not when not committing to being in that group: normally the addressee becomes committed to liking the taste of licorice in virtue of accepting the speaker's assertion, *regardless* of whether the addressee actually likes licorice: where the addressee is independently committed to disliking licorice, accepting the speaker's assertion in such a context is anomalous or insincere, and the present account can make no sense of this.

that embed experiential predications are often felicitous and judged true based only on the experiential reactions (and resulting opinions) of the agents in question. In many cases, there is no plausible single proposition for these agents to agree on or disagree over that the embedded clause can denote on a contextualist semantics, and so the correctness of these reports is mysterious.

Doxastic verbs like *think* allow distributive readings with complex subjects that relate distinct agents to distinct propositional contents. Suppose Alfonse and Bethany are testing the properties of licorice by feeding it to animals: Alfonse feeds it to cats, while Bethany feeds it to dogs. They both conclude that licorice is tasty to the sort of animal they've been feeding it to, and so (49a) and (49b) are true.

(49) a. ✓Alfonse thinks that licorice is tasty to cats.

b. ✓Bethany thinks that licorice is tasty to dogs.

Suppose further that both Alfonse and Bethany are under the impression that licorice is *not* tasty to the sort of animal they didn't feed: thus Alfonse thinks that licorice isn't tasty to dogs, and Bethany thinks it isn't tasty to cats. Their beliefs are summarized in (50).

(50) **Who is licorice tasty to?**

	cats	dogs
Alfonse	✓	✗
Bethany	✗	✓

If a curious interlocutor asks how the investigation went, (51a) can be read as a true report of Alfonse's and Bethany's conclusions – (51b) can follow this up to clarify the intended target of belief for each member of the complex subject.

(51) a. ✓Alfonse and Bethany (both) think that licorice is tasty.

b. (Alfonse, to cats, and Bethany, to dogs).

Where true, (51a) means that both Alfonse and Bethany think *that licorice is tasty to the sort of animal they were investigating.* The report doesn't attribute a common belief to

Alfonse and Bethany, appropriately, since in this situation they have no relevant belief in common.[24]

This sort of reading is captured when the complex subject quantifies into an open sentence and binds off the experiencer index ϵ as follows, where *Alfonse* and *Bethany* are Montague-lifted to denote quantifiers (54a)-(54b), *and* denotes standard generalized conjunction (54c), as in Partee & Rooth (1983), *think* has an ordinary denotation relating agents to embedded propositions in terms of doxastic alternatives (53a), and x_ϵ makes use of the function index $f_{c,\epsilon}$ introduced in §1.1.

(52)

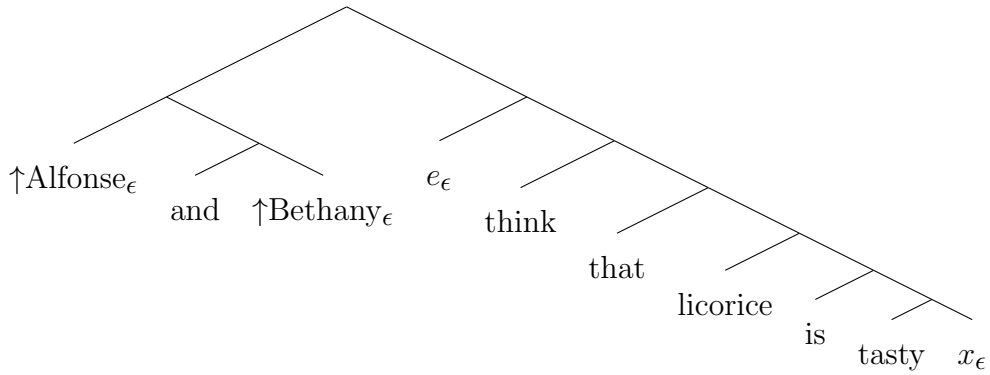

(53) a. $[\![\text{think}]\!]^{w,g} = \lambda\phi_{a,st}.\lambda x_e.\forall w' \in Dox_{x,w}[\phi(g)(w')]$

b. $[\![\text{licorice [is [tasty } x_\epsilon]]]\!]^{w,g} = tasty''(w)(g_{f,w}(\epsilon))(l)$

c. $[\![e_\epsilon \text{ [think [that [licorice [is [tasty } x_\epsilon]]]]]]\!]^{w,g} =$
$\forall w' \in Dox_{g(\epsilon),w}[tasty''(w')(g_{f,w}(\epsilon))(l)]$

(54) a. $[\![\uparrow\text{Alfonse}_\epsilon]\!]^{w,g} = \lambda\phi_{a,st}.\phi(g[\epsilon \to a])(w)$

b. $[\![\uparrow\text{Bethany}_\epsilon]\!]^{w,g} = \lambda\phi_{a,st}.\phi(g[\epsilon \to b])(w)$

c. $[\![\text{and}]\!]^{w,g} = \lambda\alpha.\lambda\beta.\alpha \sqcap \beta$

d. $[\![\uparrow\text{Alfonse}_\epsilon \text{ [and } \uparrow\text{Bethany}_\epsilon]]\!] = \lambda\phi_{a,st}.\phi(g[\epsilon \to a])(w) \wedge \phi(g[\epsilon \to b])(w)$

24. Assume that Alfonse and Bethany have no opinions either on the taste of licorice themselves, or on the reaction to licorice of any other third party. The attempts in e.g. Moltmann (2010: 190-191) and Lasersohn (2017: 151) to use distributive reports of this sort as a diagnostic for shared propositional content of attitudes fail for this reason. Similar observations hold for anaphoric implicit arguments generally, e.g. in *Alfonse and Bethany (both) think that a local bar serves sours.*

(55) $[\![[\uparrow\text{Alfonse [and } \uparrow\text{Bethany}]]_\epsilon\ [e_\epsilon\ [\text{think [that [licorice [is [tasty } x_\epsilon]]]]]]]\!]^{w,g} =$
$\forall w' \in Dox_{a,w}[tasty''(w')(g_{f,w}(a))(l)] \wedge \forall w' \in Dox_{b,w}[tasty''(w')(g_{f,w}(b))(l)]$

In a context c such that $f_{c,\epsilon}$ is the function that maps worlds w to individuals x to the individual (kind) that x was investigating the tastes of at w, the desired result is achieved, where small caps denote the relevant kind.[25]

(56) $(\!|\uparrow\text{Alfonse}_\epsilon \text{ and } \uparrow\text{Bethany}_\epsilon\ e_\epsilon \text{ think that licorice is tasty } x_\epsilon|\!)^c =$
$\lambda w_s.\forall w' \in Dox_{a,w}[tasty''(w')(\textsc{cat})(l)] \wedge \forall w' \in Dox_{b,w}[tasty''(w')(\textsc{dog})(l)]$

But *agree* differs from *think* in disallowing such distributive readings (cf. Cappelen & Hawthorne *ibid.*: 56-57). There's no reading of (57) that can be appropriately judged true in relation to the situation described by (50). Alfonse and Bethany don't share any relevant beliefs regarding the tastiness of licorice, and so there is no candidate propositional content denoted by the embedded clause that both believe.[26]

(57) ✗Alfonse agrees with Bethany that licorice is tasty.

Agree composes with an individual denoted by a *with*-headed PP, a proposition, and a (subject) individual, and requires that both individuals believe the proposition denoted by the embedded clause. It thus allows for two distinct agents to be related to the intension

25. Assume for simplicity that at all worlds, the kind of animal that Alfonse and Bethany investigate remains constant, i.e. cats and dogs, respectively.

26. Instead of sentences like (57), Cappelen & Hawthorne make use of complex-subject *agree*-reports, e.g. *Alfonse and Bethany agree that licorice is tasty.* This obscures the crucial difference in the argument structure of *think* versus *agree* (compare (53a) to (58)). For the latter, neither internal argument to the verb needs to be overt: either or both can be implicit and anaphoric:

(1) a. Alfonse agrees with Bethany.
b. Alfonse agrees that licorice is tasty.
c. Alfonse agrees.

This suggests that the former complex-subject cases have the same argument structure as (57), with an implicit individual internal argument that's preferably reciprocal. That it's not inherently reciprocal is seen from e.g. *Gamela thinks that licorice is tasty. Alfonse and Bethany (both) agree*: this can read that Alfonse and Bethany agree with Gamela, not with each other. Nor is *agree* inherently anti-distributive: e.g. as expected on the present treatment, *Alfonse and Bethany (both) agree with Gamela that licorice is tasty* can be read as true where both Alfonse and Gamela think licorice is tasty to cats, and both Bethany and Gamela think licorice is tasty to dogs.

of the embedded clause without a complex subject, and from this the lack of a distributive reading resulting from binding off the experiencer index follows.

(58) $[\![\text{agree}]\!]^{w,g} = \lambda x_e.\lambda\phi_{a,st}.\lambda y_e.\forall w' \in Dox_{x,w}, Dox_{y,w}[\phi(g)(w')]$

(59)

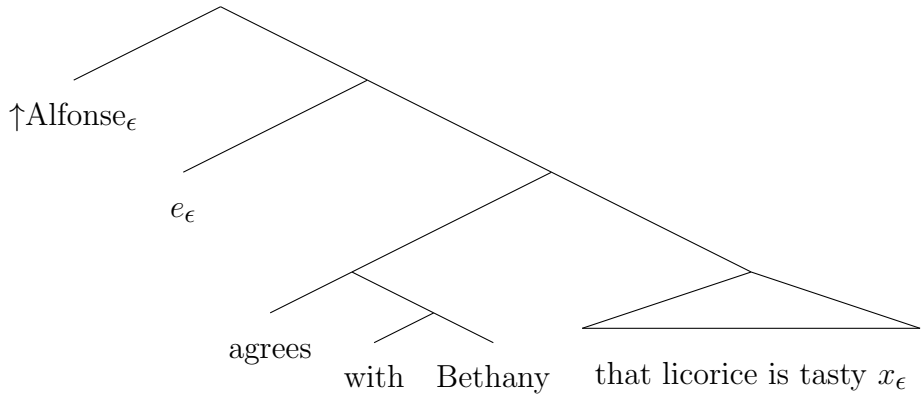

(60) $[\![\uparrow\text{Alfonse}_\epsilon\ [e_\epsilon\ [[\text{agrees [with Bethany]}]\ [\text{that [licorice [is [tasty } x_\epsilon]]]]]]]\!]^{w,g} =$
$\forall w' \in Dox_{a,w}, Dox_{b,w}[tasty''(w')(g_{f,w}(a))(l)]$

Depending on the value of $f_{c,\epsilon}$ in the context of utterance, the relevant proposition asserted to be believed by both Alfonse and Bethany may differ, but relative to any single context the same proposition must be believed by both. And so the falsity of (57) as opposed to (51a) is intelligible.

But *agree*-reports involving experiential predicates can be read as true, even where the contextualist semantics predicts that there is no plausible proposition that two distinct agents believe and that can be denoted by the embedded clause. Suppose that both Alfonse and Bethany like the taste of licorice (and on this basis alone feel justified in asserting *Licorice is tasty*), but that each mistakenly believes that the other is repulsed by it. The situation is analogous to that in (50), and their beliefs are summarized as follows.

(61) **Who is licorice tasty to?**

	Alfonse	Bethany
Alfonse	✓	✗
Bethany	✗	✓

In such a situation, a distributive belief report can be read as true.

(62) ✓Alfonse and Bethany (both) think that licorice is tasty.

This is expected on the contextualist semantics: the structure in (52) and the denotation in (55) work fine, and in a context in which $f_{c,\epsilon}$ is vacuous, denoting the (function from worlds to) the identity function on individuals, the resulting reading is that Alfonse and Bethany both think *that licorice is tasty to themselves*. Thus even though they share no relevant belief, the distributive reading can relate them to distinct contents and be true.

What is unexpected is that, unlike with the situation in (50), the corresponding *agree*-reports can be read as true as well. Despite being misinformed about each other's beliefs, Alfonse and Bethany can truly be reported as agreeing that licorice is tasty (even though they don't know that they do).

(63) ✓Alfonse agrees with Bethany that licorice is tasty.

But what is the proposition they both believe in common? The obvious answer, *that licorice is tasty* (*simpliciter*), isn't available on the contextualist semantics, since the embedded clause can't denote such a proposition: in c, it must be that they both believe that licorice is tasty by the standard of ϵ_c (or some individual functionally related to it). But there is no candidate for such an individual that accurately describes their beliefs. There are four options as to what it could be.

1. ϵ_c is some group containing Alfonse, but not Bethany
   (possibly just Alfonse himself).

2. ϵ_c is some group containing Bethany, but not Alfonse
   (possibly just Bethany herself).

3. ϵ_c is some group containing both Bethany and Alfonse
   (possibly the group consisting of just Bethany and Alfonse themselves).

4. ϵ_c is some group containing neither Alfonse nor Bethany.

[1] and [3] wrongly attribute to Bethany the belief that Alfonse likes licorice; [2] and [3] wrongly attribute to Alfonse the belief that Bethany likes licorice; [4] gives an irrelevant reading false in this scenario (cf. fn. 24). Appealing to $f_{c,\epsilon}$ won't help, since the output of the function will be faced with the same quadrilemma. The contextualist semantics thus fails to predict the behavior of *agree* with respect to experiential predicates.[27]

A parallel problem arises for *disagree*, which relates two agents to a single propositional content, and requires that one believe, and the other disbelieve, that content. Suppose that Alfonse and Bethany are again conducting their licorice investigations: Alfonse comes to believe that licorice is tasty to cats, while Bethany comes to believe that it isn't tasty to dogs. Further, each accepts the others' conclusions, and so they agree as to all the results: their beliefs are then as in (64).

(64) **Who is licorice tasty to?**

	cats	dogs
Alfonse	✓	✗
Bethany	✓	✗

If a curious interlocutor asks as to the results of Alfonse's and Bethany's investigations, a *think*-report like (65) can be judged true, but a *disagree*-report as in (66) can't.

(65) ✓Alfonse thinks that licorice is tasty, but Bethany doesn't.

(66) ✗Alfonse disagrees with Bethany that licorice is tasty.

This is unsurprising, given the above: (65) can be read to mean that Alfonse, but not Bethany, thinks *that licorice is tasty to the animal they were investigating.* (66), by contrast, can't be read as true, because there is no plausible proposition that the embedded clause could denote over which Alfonse and Bethany disagree. A denotation for *disagree* capturing this behavior would be as in (67).

27. A strategy analogous to the one tried in §1.1.3 – relating Alfonse and Bethany both to a *de dicto* belief regarding a group of experiencers relevantly like themselves – won't work either, as the reader can confirm.

(67) $[\![\text{disagree}]\!]^{w,g} = \lambda x_e.\lambda\phi_{a,st}.\lambda y_e.\forall w' \in Dox_{x,w}[\phi(g)(w')] \wedge \forall w' \in Dox_{y,w}[\neg\phi(g)(w')]]$

But in a situation in which Alfonse likes the taste of licorice, and on this basis alone feels justified in uttering *Licorice is tasty*, and Bethany doesn't like the taste of licorice, and on this basis alone feels justified in uttering *Licorice isn't tasty*, then the two can be felicitously and truly reported as disagreeing, and this is so even if they are entirely aware of each other's tastes (and everyone else's), so that they have no disagreements at all as to what is tasty to whom.

(68) **Who is licorice tasty to?**

	Alfonse	Bethany
Alfonse	✓	✗
Bethany	✓	✗

(69) ✓Alfonse disagrees with Bethany that licorice is tasty.

The situation parallels that in (64), and yet the behavior of the *disagree*-report differs. Even if Alfonse's and Bethany's *de re* beliefs regarding what is tasty to whom entirely align, as do their *de dicto* beliefs about what is tasty to those groups relevantly similar in taste to themselves, they can still be reported as disagreeing, and so the contextualist can't account for the behavior of this verb.

A final note on the doxastic nature of *disagree*: Huvenes (2012: 167) makes use of this verb as an explicit marker of disagreement, and on this basis suggests that disagreement need not be explained in terms of non-cotenability of semantic content, as in (70), where explicit disagreement is possible even though the interlocutors' assertions are non-contradictory.

(70) Alfonse: I like this licorice.

Bethany: I disagree – it's too sweet for me.

If *disagree* in (70) is the same verb that has been discussed above, and so has the same argument structure, then it has here two anaphoric implicit arguments, one of which is propositional (cf. fn. 26): to the extent that this exchange is felicitous, it is so because

this argument is valued for some proposition toward which the interlocutors have conflicting doxastic attitudes (if e.g. in liking the licorice, Alfonse commits to thinking *that the licorice is tasty*, and Bethany commits to the opposite in not liking it).

If the propositional argument is overt, then targeting what Alfonse says results in infelicity, precisely because there is no non-cotenability between Bethany's belief state and the proposition asserted: *#I disagree that you like this licorice – it's too sweet for me.* If the overt propositional complement makes explicit the doxastic requirements of *disagree*, this is problematic for contextualists and their fellow-travelers that want to explain disagreement involving experiential predicates generally by moving away from conflict in doxastic attitudes (cf. Huvenes 2014: 151-152; Wyatt 2018: 270-273). Doxastic conflict is precisely what *disagree* is sensitive to, and conflict at some other, non-doxastic level is not enough to make *disagree*-reports felicitous or true.

1.1.5 Argument V: acquaintance inferences

Assertions of experiential predications commit the speaker to having had direct experience of the relevant sort with the individual predicated of (cf. Pearson 2013: 117-118; Klecha 2014: 450-451; MacFarlane 2014: 3, Ninan 2014; Bylinina 2017: 300).

(71) Licorice is tasty.

⇝ The speaker has tasted licorice.

Pearson (2013: 122) takes this inference to be presuppositional: this is unsurprising on a contextualist semantics, which independently requires that experiential predicates enforce direct experience presuppositions with respect to their experiencer argument (cf. §1.1), and can posit that the value of the experiencer is by default anchored to the speaker (cf. (21)).

The proposition implied is also identical to that implied in the presuppositional cases. The inference in (71) thus reduces to the inference in (72), if (71) is uttered when no experiencer is salient, antecedent, or accommodable, so that $s_c = \epsilon_c$.

(72) Licorice is tasty to me.

↪ The speaker has tasted licorice.

(73) a. $(\!|71|\!)^c = (\!|\text{licorice is tasty } x_\epsilon|\!)^c = \lambda w_s.tasty''(w)(\epsilon_c)(l)$

b. $(\!|72|\!)^c = (\!|\text{licorice is tasty to me}|\!)^c = \lambda w_s.tasty''(w)(s_c)(l)$

But Ninan (2014) has demonstrated that this isn't right: (71) involves a distinct *acquaintance inference* whose behavior deviates from direct experience presuppositions in several respects.

First, acquaintance inferences can't be canceled beneath external negation, while presuppositions can be (cf. *ibid.*: 297-298; the italicization of the negation marks focus: this and the justification clause are meant to prime the acceptability of external negation).

(74) a. Licorice *isn't* tasty to me, because I haven't tasted it.

b. #Licorice *isn't* tasty, because I haven't tasted it.

Second, acquaintance inferences can't be accommodated into the local intensional environment, unlike presuppositions (cf. *ibid.*: 299-300). The reader can confirm that this same pattern holds for the range of intensional contexts in (38)-(40).

(75) a. Licorice must be tasty to me.

$\leftrightarrow_{acc}$ It must be that: the speaker has tasted licorice, and likes its taste.

b. Licorice must be tasty.

$\not\leftrightarrow_{acc}$ It must be that: the speaker has tasted licorice, and likes its taste.

Third, acquaintance inferences are canceled beneath predictive operators, rather than being projected into the future, as presuppositions are (cf. Klecha 2014: §1.3).

(76) a. The licorice will be tasty to me.

$\leftrightarrow_{acc}$ It will be that: the speaker tastes the licorice, and likes its taste.

b. The licorice will be tasty.

$\not\leftrightarrow$ It will be that: the speaker tastes the licorice, and likes its taste.

But the contextualist holds that *tasty* in (71) composes with a covert experiencer, in the same way that it would compose with an overt experiencer PP, so that all occurrences of experiential predicates should give rise to direct experience presuppositions in the same way, with respect to the predicate's experiencer argument.

The problem that the contextualist faces with acquaintance inferences is thus twofold. First, assertions like (71), with no salient, antecedent or accommodable experiencer, don't give rise to direct experience presuppositions. Second, these assertions give rise to a different sort of inference for which a contextualist semantics has no account, because it treats all instances of experiential predicates on a par in virtue of composing with an experiencer argument.

The contextualist semantics can accommodate these facts by altering itself to fit the suggestions in §1.1.2, according to which experiencer PPs, and covert experiencers appearing in exocentric contexts, are modifiers of the experiential predicate that shift the assignment function, and introduce direct experience presuppositions ((77a)-(77b); repeated from (43b), (45a)). Experiential predicates themselves then inherently encode reference to an experiencer (77c) (repeated from (43a)).

(77) a. $[\![\text{to}]\!]^{w,g} = \lambda x_e.\lambda P_{a,\langle s,et\rangle}.$
$\lambda y_e : x$ has direct experience of y of the kind relevant to P at w.
$P(g[\epsilon \rightarrow x])(w)(y)$

b. $[\![x_\epsilon]\!]^{w,g} = \lambda P_{a,\langle s,et\rangle}.$
$\lambda y_e : g(\epsilon)$ has direct experience of y of the kind relevant to P at w.
$P(g)(w)(y)$

c. $[\![\text{tasty}]\!]^{w,g} = \lambda x_e.tasty''(w)(g(\epsilon))(x)$

This allows the contextualist a distinction between unmodified and modified uses of experiential predicates, with only modified uses imposing presuppositions, as desired. Bare uses would then merely be constrained by a dispositional reading of the predicate, and these

could conceivably give rise to acquaintance inferences anchored to ϵ_c, e.g. for epistemic reasons as Ninan (2014: 302) suggests.

But acquaintance inferences in assertions are solely speaker-directed.[28] Since virtually all unmodified, non-exocentric experiential predications give rise to acquaintance inferences, this means that for such a move to work, the contextualist is forced to claim that ϵ_c must invariably be the speaker, and that the only way to shift the experiencer is by means of the same modifiers that introduce direct experience presuppositions.

(78) $\epsilon_c = s_c$

But this raises a problem with respect to the contextualist treatment of conversational denial, for which anchoring bare uses of experiential predicates intrinsically to the speaker yields the wrong results.

Echoic denials of experiential predications with no overt experiencer PP (79a) are in general felicitous: these involve a rejective discourse particle like *no* followed by a counter-assertion using the same sentence but with sentential negation (79b). This isn't so for echoic denials of experiential predications with an overt experiencer PP including a speaker-indexical (80); cf. Kölbel (2002: §3.3); Lasersohn (2005: §2).[29]

Let both (79) and (80) take place in a context in which there is no salient, antecedent, or accommodable experiencer, and let both Alfonse and Bethany make their assertions sincerely in each case only on the basis of their own experiential reactions to licorice, so that both their assertions trigger acquaintance inferences. Further, let both Alfonse and Bethany be aware of and take for granted the other's attitude towards licorice.

28. In exocentric contexts, as when uttering *The licorice is tasty* while Alfonse is the salient experiencer (cf. (11a)), there results not an acquaintance inference anchored to Alfonse, but rather a direct experience presupposition anchored to him: the reader can confirm this by running the tests in this section on the relevant case.

29. (80) can be read as felicitous on some intonations in certain contexts. To block this, (79) and (80) are to be read with the intonation characteristic of an echoic denial, which marks the difference in felicity obviously. The same point can also be made from the behavior of elliptical rejection: In response to *Licorice is tasty*, the clearest reading of *No, it's not* that anchors the elided predicate to the property expressed by the first speaker does not deny that licorice is tasty by the speaker's lights, as an inherently speaker-oriented view requires.

(79) a. ALFONSE: Licorice is tasty.

b. BETHANY: No, licorice isn't tasty.

(80) a. ALFONSE: Licorice is tasty to me.

b. BETHANY: #No, licorice isn't tasty to me.

This is a problem for a contextualist semantics that adopts (78): this move reduces (80) to (79), the only predicted difference between the two being (if the above recommended modifications are made) that the former triggers distinct acquaintance inferences on the part of the speakers, while the latter triggers distinct speaker-oriented direct experience presuppositions.

As has been discussed at length by Kölbel and Lasersohn (*ibid.*) and elsewhere, no independently plausible account of the felicity of echoic denials makes sense of the felicity of the exchange in (79) on an inherently speaker-oriented account, since such an account results in (79a) and (79b) not being assertions of conflicting contents. There is no issue with accommodating both that licorice is tasty by Alfonse's standard, and that it is not tasty by Bethany's, into the common ground. (79) is expected to pattern like (80) in this respect, yielding infelicity for denial targeting compatible contents, contrary to fact.

There is an extensive literature disputing whether a contextualist semantics can make sense of an exchange like (79) (cf. the references in Cappelen & Huvenes 2018), but an inherently speaker-oriented semantics for unmodified uses of experiential predicates is unacceptable for this task.[30] An account of acquaintance inferences that requires commitment to such a faulty semantics is equally unacceptable.

§1.2 shows that the contextualist's problems discussed above are dispelled if experiential predicates are not forced to take an experiencer semantic argument. As such, contextualism

30. A method of salvaging an intrinsically speaker-oriented contextualist semantics may be found in a strain of contextualist thought that accounts for aspects of disagreement involving experiential predicates by attributing semantic blindness or incompetence to speakers: cf. Stojanovic (2007: §2), Cappelen & Hawthorne (2009: 118). Blindness theory will be briefly discussed at the start of Chapter 5. López de Sa (2008) and Gutzmann (2016) offer intrinsically speaker-oriented contextualist semantics, but with additional machinery to handle the intelligibility of conversational disagreement.

has little going for it as a semantic hypothesis, and pending some undiscovered good linguistic reason to pick it up again, it ought to be abandoned for the sorts of constructions that the literature has so far discussed.

1.2 Relativism: the experiencer as intensional parameter

The relativist holds that the extension of an expression is evaluated relative to a standard of experience, in addition to classical intensional parameters, like a world of evaluation. As on a contextualist semantics, there are many ways to model a standard of experience (cf. fn. 2). The following adopts the treatment in Lasersohn (2005) and Stephenson (2007a), which take these standards to be type-e experiencers (their 'judges'). The intensional type of an expression α in general is thus $\langle s, \langle e, \tau_\alpha \rangle \rangle$, where τ_α is the extensional type of α, and the intensional type of an expression of extensional type t is $\langle s, \langle e, t \rangle \rangle$, that of a 'relativist-proposition.'[31]

Experiential predicates are 'experiencer-sensitive,' in that their extensions are determined non-trivially relative to the experiencer parameter at some world (81a)-(81b). A standard denotation for the experiencer-sensitive *tasty* is as in (82a). Non-experiential expressions are experiencer-insensitive, such that the experiencer parameter is idle in determining their extensions (82b) (cf. Sæbø 2009: 335, ex. 15).

(81) a. f of type $\langle s, \langle e, \tau \rangle \rangle$, where τ is an extensional type,
is experiencer-sensitive at w iff $\exists x, y[f(w)(x) \neq f(w)(y)]$;
otherwise, f is experiencer-insensitive at w.

31. More recent relativist treatments incorporate world- and time-sensitivity into experiential standards, casting them as world-time-individual triples (Lasersohn 2017: 95) or abstract standards determined relative to such triples (MacFarlane 2014: 151). This fixes experiential standards by worlds and times independent of the world and time of evaluation: the relativist needs this to adequately treat (i) tense in relation to how experiencers' dispositions to experience change over time (cf. MacFarlane 2014: 149, 163-165; Lasersohn 2017: 139-141), and (ii) certain modal evaluations of experiential predicates (cf. the worry in Lasersohn 2005: 663, fn. 13; MacFarlane 2014: 165-166). The following ignores these issues, as well as non-world, non-experiencer intensional parameters where irrelevant.

b. α is experiencer-sensitive iff $\exists w'[\lambda w_s.\lambda x_e.[\![\alpha]\!]^{w,x}$ is experiencer-sensitive at $w']$; otherwise, α is experiencer-insensitive.

(82) a. $[\![\text{tasty}]\!]^{w,x} = \lambda y_e.tasty''(w)(x)(y)$

b. $[\![\text{vegetarian}]\!]^{w,x} = \lambda y_e.vegetarian'(w)(y)$

The metalanguage predicate '*tasty''*' in (82a) is the same one used in the standard contextualist denotation for *tasty* (1), and is to be read in the same way (cf. (2)). Relativism is a dyadic semantics in the sense that it invokes a metalanguage predicate that is dyadic relative to worlds: as on a contextualist account, there are no experiential properties of individuals *simpliciter* at worlds, but rather only stimuli's production of, or disposition to produce, experiences in experiencers.

But the relativist denotation for an experiential predicate is extensionally monadic, not dyadic: *tasty* composes only with a stimulus argument, and not with an experiencer argument. The experiencer argument to the metalanguage predicate (x in (82a)) is saturated not by something with which the object language predicate must compose, but rather with the value of an intensional parameter (though see below: the relativist can also allow composition with experiencers via modifiers).

This has consequences for the semantic values of the expressions involved, and for the semantics-pragmatics interface.[32] For instance, some truth-apt sentences are experiencer-sensitive.

(83) a. Licorice is tasty.

b. $[\![\text{licorice [is tasty]}]\!]^{w,x} = tasty''(w)(x)(l)$

c. $(\!|\text{licorice is tasty}|\!)^c = \lambda w_s.\lambda x_e.tasty''(w)(x)(l)$

Where at w, licorice produces, or is disposed to produce, gustatory pleasure in x, but

32. Stojanovic (2007: §4) demonstrates a weak equivalence between contextualist and relativist semantics in the assignment of truth values relative to parameters, but this equivalence doesn't extend to the assignment of contents to expressions: compare e.g. (83c) to its contextualist near-counterpart (14b) (cf. Lasersohn 2008: 317-318).

not in y, $[\![\text{licorice is tasty}]\!]^c(w)(x) = true$, while $[\![\text{licorice is tasty}]\!]^c(w)(y) = false$.

This raises the question of how the truth of a relativist-proposition is to be assessed at a world at which it is experiencer-sensitive, since it has no truth value *simpliciter* at such a world. To assess a proposition as true is to take it to be true: this primitive notion has pragmatic consequences, as one who assesses a proposition as true will, pending confidence in that assessment, commit to its truth, and may thereby come to believe it, act as though it were true, assert it felicitously under Gricean sincerity conditions, etc. The same holds *mutatis mutandis* for assessing a proposition as false.

Speakers operate under norms governing which values for intensional parameters are to be supplied as the arguments to a proposition in assessing it for truth. This can be explicated using the notion of 'contexts of assessment' (MacFarlane 2007: 26 ff; 2014: 60 ff; Lasersohn 2017: 93-94). Informally, these are concrete situations in which truth assessments are made; formally, they can be treated as tuples of parameter values c' (and so as objects of the same kind as contexts of utterance), including at least $w_{c'}$, the world of assessment, and $a_{c'}$, the assessor (the one who assesses for truth).

Two aims of truth assessment can be distinguished: assessment is 'correct' when the assessor 'gets it right' factually, and it is 'proper' when the assessor behaves in accordance with pragmatic norms, in some suitably restricted sense. Leaving experiential predicates (and other problematic expressions) to the side, it's trivial to state on a non-relativist semantics, where propositions ϕ are of type $\langle s, t\rangle$, how these operate: correct and proper truth assessment both track the truth value of the proposition at the world of assessment.[33]

(84) a. i. ϕ is correctly assessed as true in c' iff $\phi(w_{c'})$
(and is incorrectly assessed as true in c' otherwise).

ii. ϕ is correctly assessed as false in c' iff $\neg\phi(w_{c'})$

33. For any actual act of assessment in c', $w_{c'}$ is the actual world. Formulating the norms with respect to the world of assessment generally takes care of counterfactual correct and proper assessment automatically: if things were different, it would be correct and proper to assess with respect to the way things would be. Lasersohn (2013: §7) comments on the world of evaluation as fixed by the context of assessment.

(and is incorrectly assessed as false otherwise).

b. i. ϕ is properly assessed as true in c' iff $\phi(w_{c'})$
(and is improperly assessed as true in c' otheriwse).

ii. ϕ is properly assessed as false in c' iff $\neg\phi(w_{c'})$
(and is improperly assessed as true in c' otherwise).

Truth assessment is correct just in case it tracks the truth of the proposition at the world of assessment, and pragmatic norms bid speakers to assess relative to the world of assessment, and not relative to counterfactual situations. Thus ϕ is correctly assessed as true in c' iff it is properly assessed as true in c'. Further, since the assessor c'_a plays no role in these definitions, correct and proper truth assessment converge for all assessors at a world.

The question is then how speakers are to choose values for the intensional parameters in assessing the truth of relativist-propositions, and in particular for those that are experiencer-sensitive at the world of assessment. The relativist's answer is as follows, where 'ϕ' is a variable over relativist-propositions of type $\langle s, \langle e, t\rangle\rangle$ (and *mutatis mutandis* for falsity).

(85) a. ϕ is true *simpliciter* at w iff $\forall x[\phi(w)(x)]$.

b. ϕ is false *simpliciter* at w iff $\forall x[\neg\phi(w)(x)]$.

(86) a. i. ϕ is correctly assessed as true in c' iff ϕ is true *simpliciter* at $w_{c'}$.

ii. ϕ is incorrectly assessed as true in c' iff ϕ is false *simpliciter* at $w_{c'}$.

b. ϕ is properly assessed as true in c' iff $\phi(w_{c'})(a_{c'})$
(and is improperly assessed as true otherwise).

(86a) says that correct truth assessment tracks the truth of the relativist-proposition at the world of assessment for all experiencers (cf. Lasersohn's 2017: 102, ex. 141 'absolute' truth predicate). Where ϕ is experiencer-insensitive at w, this condition effectively collapses into its non-relativist counterpart in (84a), since by definition ϕ is then true or false *simpliciter* at w. But on this definition, relativist-propositions that are experiencer-sensitive

at w can neither be correctly nor incorrectly assessed as true at w: there is no 'fact of the matter' at w as to whether they hold or not.[34]

(86b) says that proper truth assessment tracks the truth of the relativist-proposition at the world of assessment relative to the assessor: it is *autocentric*, and so sensitive to the assessor's own (disposition to) experience. Where ϕ is experiencer-insensitive at w, autocentric truth assessment of it is vacuous at w, since then value of the experiencer intensional parameter plays no non-trivial role. But where ϕ is experiencer-sensitive at w, proper truth assessment at w crucially depends on the assessor's (disposition to) experience: thus if at w, licorice produces, or is disposed to produce, gustatory pleasure in x, but not in y, then it is proper for x to assess *that licorice is tasty* (83c) as true at w, and proper for y to assess it as false at w.[35]

For experiencer-insensitive relativist-propositions, the norms in (86) effectively collapse into the non-relativist norms in (84): correct and proper truth assessment converge, and they do so for all assessors at the world in question. But experiencer-sensitive propositions deviate: correct assessment of them is a non-issue, and proper assessment varies from experiencer to experiencer depending on the relevant (dispositions to) experience. Not to assess in accord with one's own experience is a pragmatic error.[36]

---

34. If the standard of experience parameter is simply an experiencer, as (86a) assumes, a relativist-proposition can be true or false *simpliciter* at a world just because all experiencers happen to be the same at that world with respect to the relevant sorts of (dispositions to) experience. If this 'experiential unanimity' effect is an undesirable consequence, the relativist can avoid it by complicating the intensional parameter to be world-sensitive (cf. fn. 31), and make quantification in (85) range over these world-bound experiential standards, so that the experiencer-sensitivity of a relativist-proposition at a world doesn't depend on the states of experiencers at that world.

35. An equivalent way of saying this (cf. Egan 2010: 276-279) is that relativist-propositions are properties, in the vein of Lewis (1979) (note that their type, $\langle s, \langle e, t\rangle\rangle$, is that of the intension of a property of individuals on a non-relativist semantics), and that speakers self-attribute those properties. Thus, *that licorice is tasty* denotes the property of liking the taste of licorice, and to commit to its truth is to take oneself *de se* to be a liker of the taste of licorice. This is the same as casting the content of declarative clauses as sets of centered worlds, in which the experiencer plays the role of center (possibly along with other roles: cf. Stephenson 2010).

36. This error can be grounded in factual errors, where assessors are ignorant about their own experiences: cf. Kölbel (2009: 383, fn. 12); Lasersohn (2017: 93-94). Lasersohn (2005: 670-674; 2009: 364; 2017: §7.2, §7.5) allows for proper truth assessment sometimes to be non-autocentric, such that the proper value for the experiencer parameter in assessing for the truth of ϕ in c' is not $a_{c'}$, but rather some individual contextually

With this machinery for truth assessment in place, speakers can assert, reject, retract, etc. relativist-propositions governed by norms on sincerity and evidence as standard: there's no need for special relativist pragmatic norms governing speech acts, as e.g. Lasersohn (2005: 670), Stephenson (2007a: 509), and MacFarlane (2014: 103, 108, 110) suggest.

Competent speakers are aware that the norm governing proper truth assessment in (86b) holds for speakers generally, and not just for themselves.[37] Whether one speaker takes another to be assessing the truth of a relativist-proposition properly doesn't always track whether the first speaker agrees with the second regarding the truth of that proposition: where x likes the taste of licorice and y doesn't, and both are fully knowledgable about the other's tastes, both x and y will take the other to be assessing properly just in case they take the other to disagree with them regarding the truth of *that licorice is tasty*. In such situations, it's not possible for competent speakers to take no issue whatsoever with one another: either they both assess properly, and so disagree, or they agree, and so at least one assesses improperly.[38]

Finally, the relativist can allow experiential predicates to compose with experiencers, by treating lexical items that introduce experiencers as modifiers that shift the experiencer parameter (cf. Lasersohn 2005: 666, ex. d; 2017: 103-106; MacFarlane 2014: 153-156). This can be done either overtly, to capture experiencer PPs (87a), or covertly, to capture

---

salient in c'. This is meant to treat exocentric uses of experiential predicates. In dealing with exocentricity, §2.3.2 argues that this possibility should be abandoned.

37. There are linguistic reflexes of this fact: it is typically the speaker's experiences that a non-exocentric assertion of *Licorice is tasty* (putting aside complications with inaccurate perceptions – §4.2.1) must track in order to be felicitous, not just according to the speaker, but to any interlocutor: thus, such an utterance by Alfonse implies both that Alfonse has tasted licorice, and that Alfonse likes its taste (as part of the acquaintance inference: cf. §1.1.5, §4.1). Similar effects arise for attitude holders in ordinary doxastic reports, as in *Alfonse thinks that licorice is tasty*, which implies these same things (cf. §4.1.4).

38. Such cases of disagreement are thus 'faultless,' in Kölbel's (2004) terminology, in two senses: (i) neither interlocutor assesses incorrectly for truth (86a); and (ii) both interlocutors assess properly for truth (86b) (and would be committing an impropriety to hold the contrary belief). The disagreement is not faultless in one sense: each interlocutor takes the other to have a false belief, and so each takes the other to be making an 'error of taste' in the sense of Lasersohn (2017: 210, ex. 49). Lasersohn (2009: 364; 2017: §7.6) also allows for interlocutors to take an 'acentric stance,' and refuse to assess experiencer-sensitive contents for truth, which if possible may serve to remove even the intuition of an error of taste.

exocentric readings (87b) (allowing extensions to be determined relative to an assignment g, as in §1.1). Presumably, overt experiencer PPs are infelicitous with non-experiential predicates due to their vacuity (cf. the similar point for contextualist *find* in §1.1.1).

(87) a. $[\![\text{to}]\!]^{w,x} = \lambda y_e.\lambda P_{s,\langle e,et\rangle}.\lambda z_e.P(w)(y)(z)$

b. $[\![x_\epsilon]\!]^{w,x,g} = \lambda P_{a,\langle s,\langle e,et\rangle\rangle}.\lambda y_e.P(g)(w)(g(\epsilon))(y)$

This allows the relativist access to every reading captured by a contextualist semantics: the result of modifying an experiential property in this way is the same as the result of saturating the experiencer argument of the predicate for the contextualist (compare e.g. (88) with (4b)). Note that the modified denotation is experiencer-insensitive.

(88) $[\![\text{tasty [to Alfonse]}]\!]^{w,x} = \lambda z_e.tasty''(w)(a)(z)$

Whatever semantic machinery the contextualist employs to obtain indexical, anaphoric, and bound readings of experiential predicates is therefore equally available to the relativist, and the same holds for the machinery employed to derive direct experience presuppositions. In this sense, a relativist semantics is strictly more expressive than a contextualist semantics: it allows for all the same readings where an experiential predicate composes with an experiencer, as well as additional readings on which it doesn't (and where the value for the experiencer is therefore supplied by the intensional parameter).[39]

What follows demonstrates that this increase in expressive power deflates the problems that afflict a contextualist semantics. §1.2.1 briefly revisits these challenges, and shows how an extensionally monadic treatment of experiential predicates resolves them.

39. And so Kneer et al. (2017) and Schaffer (2011: 192-195) are mistaken to think that 'perspectival plurality' and bound readings of experiential predicates respectively favor a contextualist semantics. Special technical tools to capture relativist bound readings have been invoked by Lasersohn (2008: 323-326) and Zeman (2015), but these are unnecessary, since there is no technical or empirical impediment to binding modifiers using the same tools as in §1.1.

1.2.1 The contextualist's problems redux

A relativist semantics removes the experiencer argument from experiential predicates, and allows for two sorts of uses of those predicates: unmodified uses don't make reference to a specific experiencer in their content (83c), while modified uses (either exocentric or with overt experiencer PPs) do. This dispels the contextualist's problems discussed in §§1.1.1-1.1.5.

First, as to subjective attitude verbs (cf. §1.1.1), a sufficient condition on predicates being embeddable beneath *find* may be that they are experiencer-sensitive (81b), as unmodified experiential predicates are (82a); those occurring with overt experiencer PPs are by contrast experiencer-insensitive, and so fail to meet this condition (88), making intelligible their infelicity beneath *find*. Since experiencer-sensitivity need not be a necessary condition, this further allows *find* to have a wider range of application than just for experiential predicates. It's also expected that *consider*-reports differ in interpretation between unmodified and modified uses of experiential predicates, since the two have different meanings: (35a) reports that Alfonse considers licorice tasty *simpliciter*, and (35b) and (35c) that he considers it tasty to some specific individual.

Second, the distribution of direct experience presuppositions (cf. §1.1.2) is intelligible. Unmodified uses of experiential predicates don't give rise to such presuppositions (cf. (40), (41)), while modified uses do (cf. (38), (39) for uses with overt experiencer PPs, and (44) for an exocentric use). The relativist captures this distinction if, as suggested for the contextualist in §1.1.2, the experiential predicate itself makes use only of the dispositional reading of '*tasty*$''$' (2a), and the direct experience presupposition is encoded only on the experiencer PP or exocentric modifier. The denotations for *to* and x_ϵ (modified from (87)) are then as follows.

(89) a. $[\![\text{to}]\!]^{w,x} = \lambda y_e.\lambda P_{s,\langle e,et\rangle}.$
$\lambda z_e : y$ has direct experience of x of the kind relevant to P at $w.P(w)(y)(x)$

b. $[\![x_\epsilon]\!]^{w,x,g} = \lambda P_{a,\langle s,\langle e,et\rangle\rangle}.$

$\lambda y_e : g(\epsilon)$ has direct experience of x of the kind relevant to P at w.

$P(g)(w)(g(\epsilon))(y)$

Refusing to make this move, and having '*tasty*$''$' encode direct experience presuppositions as in (8), wrongly predicts that such presuppositions always arise according to the assessor with respect to himself when autocentrically assessing the truth of experiencer-sensitive relativist-propositions.[40]

This further allows acquaintance inferences (cf. §1.1.5) in principle to be captured as inferences independent of direct experience presuppositions, since no direct experience presuppositions are predicted for unmodified uses of experiential predicates, and given their distinct semantics, some distinct reason may be found (again, perhaps ala Ninan 2014) for the appearance of the non-presuppositional implication.

Third, the fact that the relativist allows for relativist-propositions of the form *that licorice is tasty (simpliciter)* to be denoted by embedded clauses makes factives (cf. §1.1.3) and reports of agreement and disagreement (cf. §1.1.4) easy to treat. The relativist need only say that what Alfonse realizes according to (46a) is that licorice is tasty *simpliciter*, and likewise for what Alfonse and Bethany agree on in (63) and disagree over in (69).

For this to work, the relativist needs a notion of belief in experiencer-sensitive relativist propositions *simpliciter*. This is achievable if, following Stephenson (2007a: 496, ex. 28), the relativist takes belief to be inherently autocentric. Doxastic alternatives can then be recast as world-experiencer pairs (90a), and belief in a relativist-proposition can be taken to mean that all an agent's doxastic alternatives verify it (90b).

40. It further wrongly predicts spurious addressee-directed presuppositions: an unmodified utterance of *Licorice is tasty*, as heard by its addressee, ought to be extensionally equivalent on such an account to *Licorice is tasty to you*. That this is wrong is shown by the fact that the former cannot be heard by the addressee as encoding a presupposition that the addressee has tasted licorice, unlike the latter, which must be so heard.

(1) A: Licorice is tasty to you.
B: Wait a minute – I've never tasted licorice.

(2) A: Licorice is tasty.
B: #Wait a minute – I've never tasted licorice.

(90) a. $Dox_{x,w} := \{\langle w', y\rangle : \text{it is compatible with } x\text{'s beliefs at } w \text{ that } x \text{ is } y \text{ in } w'\}$

b. $[\![\text{think}]\!]^{w,x} = \lambda\phi_{s,et}.\lambda y_e.\forall\langle w', z\rangle \in Dox_{y,w}[\phi(w')(z)]$

To believe a relativist-proposition ϕ is thus just to take it to be properly assessed as true by oneself *de se*, viz. true at the world of thought relative to the thinker.[41] Denotations for *realize*, *agree*, and *disagree* (modified from (47b), (58), and (67)) that get the right results by importing this notion of belief are then unproblematic.

(91) a. $[\![\text{realize}]\!]^{w,x} = \lambda\phi_{s,et} : \phi(w)(x).\lambda y_e.\forall\langle w', z\rangle \in Dox_{y,w}[\phi(w')(z)]$

b. $[\![\text{agree}]\!]^{w,x} = \lambda y_e.\lambda\phi_{s,et}.\lambda z_e.\forall\langle w', v\rangle \in Dox_{y,w}, Dox_{z,w}[\phi(w')(v)]$

c. $[\![\text{disagree}]\!]^{w,x} = \lambda y_e.\lambda\phi_{s,et}.\lambda z_e.$
$\forall\langle w', v\rangle \in Dox_{y,w}[\phi(w')(v)] \wedge \forall\langle w'', u\rangle \in Dox_{z,w}[\neg\phi(w'')(u)]$

In assenting to the relativist-proposition *that licorice is tasty (simpliciter)*, a speaker thus accepts that this relativist-proposition is true as assessed autocentrically, and so must commit to liking, or being disposed to like, the taste of licorice. The reader can confirm that this allows speakers to track mutual belief or lack thereof in the same proposition, in terms of mutual (dispositions to) experience or lack thereof, in a way appropriate for the cases in §1.1.3 and §1.1.4.

1.3 Bare semantics: no experiencer

Granted that the elimination of the experiencer argument is empirically advantageous, the question remains what advantage there is in holding onto the experiencer as an intensional

41. A treatment of belief as in Lasersohn (2005: §6.2), which takes doxastic verbs effectively to have an additional argument slot for an experiencer, recapitulates all of the contextualist's problems with respect to doxastic agreement and disagreement, since it reduces all belief in experiencer-sensitive relativist-propositions to *de re* belief about the experiential reactions of experiencers. To believe that licorice is tasty autocentrically is then just to have a *de re* belief about one's own experiences, and these sorts of beliefs do not track agreement and disagreement based on converging and diverging experiences in the right way, since the single relativist-proposition *that licorice is tasty* can't be believed or disbelieved *simpliciter* (though oddly, it can be asserted *simpliciter*). Lasersohn (2009: 367; 2017: §8.6) comes to recognize the need for the relativist to adopt a Stephensonian *de se* notion of belief *simpliciter*, but never fully abandons this other treatment (Lasersohn 2017: 156, ex. 188).

parameter. The experiencer parameter serves two functions: it allows for experiencer phrases to work as intensional operators in the compositional semantics (87), and it allows speakers to assess the truth of experiencer-sensitive relativist-propositions autocentrically (86b).

But this parameter is necessary for neither of these functions, and so while it is empirically harmless at this level of description, it can be abandoned. §1.3.1 shows that experiential intensional operators can be treated without the parameter, and §1.3.2 shows that it is unneeded for autocentric assessment of traditional propositions. Given this, there is no reason to reference an experiencer in the semantics of experiential predicates at any level of description, and so the rest of this work adopts a 'bare' semantics for these predicates. A preliminary denotation for *tasty* is as follows.[42]

(92) $[\![\text{tasty}]\!]^w = \lambda x_e.tasty'(w)(x)$

'$tasty'$' is distinct from its dyadic counterpart '$tasty''$,' used in §1.1 and §1.2. Chapter 2 elaborates on (92), and provides a richer compositional semantics for experiential predicates. Until then, the entry can be read 'disquotationally:' $tasty'(w)(x)$ iff x is tasty at w.

The result of this section is that the difference between a relativist and bare semantics is trivial on the grounds considered in this chapter.[43] A bare semantics inherits the virtues of a relativist semantics, in virtue of not being committed to an experiencer argument, and in virtue of distinguishing between modified and unmodified uses of experiential predicates (cf. §1.2.1). It thus captures all of the data so far considered, and in a simpler way, since it has no need to complicate the treatment of attitudes by making them inherently *de se* (cf.

42. This position is rare in the literature. Cf. Collins (2013: 51): "The truth or falsity of tokens of the sentences [...] depends upon a relevant experiencer or judge [...] To the best of my knowledge, so much is not denied by anyone." But cf. Schafer (2011), Clapp (2015), Hirvonen (2016), and Wyatt (2018), which take some sort of bare semantics seriously, and possibly Iacona (2008).

43. Cf. Cappelen & Hawthorne (2009: 137): "Suppose we are simple-minded realists about predicates of personal taste and are presented with a supposed card-carrying relativist [...] There is quite a natural translation algorithm available to us [...] it will be natural to interpret the relativist's talk of some proposition being true at a standard of taste index as expressing the claim that the proposition is true by such and such standards [...] Meanwhile, it will be very natural to interpret the relativist's disquotational truth predicates as expressing the very properties that the realist expresses by 'true' and 'false.' According to this proposed translation manual, the so-called relativist and realist do not differ at all!"

(90)).

1.3.1 Experiential intensional operators

Lasersohn (2005: 656, 2017: 34-35) worries that a bare semantics can't compose experiencer PPs with experiential predicates, since without an experiential standard somewhere in the semantics, these PPs can't provide a value for that standard (as on a semantics referencing an experiencer: cf. (4a), (87a)).

But intensional operators can anchor properties to experiencers by shifting the world parameter. A related practice is ubiquitous in the adoption of Hintikka's (1969) approach to propositional attitudes as encoding a set of worlds anchored to an agent, viz. those worlds compatible with the agent's attitudes: e.g., a standard denotation for *think* (93b) relativizes the intension of the verb's complement clause to the agent's set of doxastic alternatives, at which the intension is evaluated for truth.

(93) a. $Dox_{x,w} := \{w' : w' \text{ is compatible with } x\text{'s beliefs at } w\}$

b. $[\![\text{think}]\!]^w = \lambda\phi_{st}.\lambda x_e.\forall w' \in Dox_{x,w}[\phi(w')]$

A notion of *experiential alternatives* can be defined on analogy with (93a).

(94) $Exp_{x,w} := \{w' : w' \text{ is compatible with } x\text{'s experiences at } w\}$

The denotation of *to* is then as in (95a):[44] it composes with an experiencer (individual) to yield a property modifier, which relativizes a property's intension to the experiencer's set of experiential alternatives. An analogous treatment could be given of exocentric modifiers, and direct experience presuppositions could be encoded on the modifier as above for the contextualist and relativist semantics.

The result of relativizing *tasty* to Alfonse is as in (95b): this denotes a property true of individuals that are tasty in all of Alfonse's experiential alternatives, i.e. tasty according to Alfonse's experiences.

44. This answers Lasersohn's (2017: 35, ex. 27) challenge to produce such a denotation.

(95) a. $[\![\text{to}]\!]^w = \lambda x_e.\lambda P_{s,et}.\lambda y_e.\forall w' \in Exp_{x,w}[P(w')(y)]$

b. $[\![\text{tasty [to Alfonse]}]\!]^w = \lambda y_e.\forall w' \in Exp_{a,w}[tasty'(w')(x)]$

An explanation is needed of what an experiential alternative is, along with a demonstration of how the notion can be employed, in combination with a semantics of experiential predicates, to yield the correct distribution of experiencer PPs and the right compositional result: these are provided in §2.3.3, where the denotation of *to* is also refined.

1.3.2 Autocentric assessment

Using experiential alternatives (cf. §1.3.1), a bare approach to experiential predicates can implement autocentric truth assessment, e.g. by holding to the norms in (96). These correspond to their relativist counterparts (cf. (86); again *mutatis mutandis* for falsity).

(96) a. ϕ is correctly assessed as true in c' iff $\phi(w_{c'})$
(and is incorrectly assessed as true in c' otherwise).

b. ϕ is properly assessed as true in c' iff $\forall w' \in Exp_{a_{c'},w_{c'}}[\phi(w')]$
(and is improperly assessed as true in c' otherwise).

(96a) is the non-relativist notion of correct truth assessment (cf. (84)). (96b) says that proper truth assessment tracks truth at the assessor's experiential alternatives at the world of assessment.

Let the constraints in (97)-(98) on experiential alternatives hold, where ϕ' is the 'relativist-proposition counterpart' of the proposition ϕ.[45]

(97) ϕ is experiential at w iff ϕ' is experiencer-sensitive at w (cf. 81a)
(and is non-experiential at w otherwise).

45. ϕ' (an object of type $\langle s, \langle e, t\rangle\rangle$) is the relativist-proposition counterpart of ϕ (an object of type $\langle s, t\rangle$) iff there is an expression $\rho = \ulcorner\lambda w_s.\lambda x_e.\psi\urcorner$ in the metalanguage – ψ an expression in the metalanguage of type t – that denotes ϕ', and if ρ is altered such that all instances of type-$\langle e, t\rangle$ expressions $\ulcorner\alpha''(w)(x)\urcorner$ in ψ ('α' a variable over strings corresponding to English object-language predicates), where 'w' and 'x' are bound by the leftmost occurrences in ρ of 'λw_s' and 'λx_e' respectively, are replaced with $\ulcorner\alpha'(w)\urcorner$, and the leftmost occurrence 'λx_e' in ρ is removed, the result is an expression that denotes ϕ. The reader can confirm that this procedure yields the correlation of the two translations of *that y is tasty* to follow.

(98) a. If ϕ is non-experiential at w,

i. $\phi(w)$ iff $\forall x[\forall w' \in Exp_{x,w}[\phi(w')]]$.

ii. $\neg\phi(w)$ iff $\forall x[\forall w' \in Exp_{x,w}[\neg\phi(w')]]$.

b. If ϕ is experiential at w, $\forall w' \in Exp_{x,w}[\phi(w')]$ iff $\phi'(w)(x)$.

(98a) says that non-experiential propositions are experiential-alternative-insensitive, in the same way that their relativist-proposition counterparts are experiencer-insensitive. Where ϕ is non-experiential at w, principles (96a) and (96b) converge: ϕ is correctly assessed as true at w just in case it's properly assessed as true at w, and correct and proper truth assessment converge for all individuals at w. This mirrors the relativist situation in §1.2.

(98b) says that experiential propositions are experiential-alternative-sensitive in the same way that their relativist-proposition counterparts are experiencer-sensitive. E.g., the relativist-proposition counterpart of (99a) (*that y is tasty*) is (99b).

(99) a. $\lambda w_s.tasty'(w)(y)$

b. $\lambda w_s.\lambda x_e.tasty''(w)(x)(y)$

The former is true at every member of $Exp_{x,w}$ just in case the latter is true at w and x: $\forall w' \in Exp_{x,w}[tasty'(w')(y)]$ iff $tasty''(w)(x)(y)$. Experiential propositions are properly assessed as true just in case they're true according to the assessor's experiences at the world of assessment, where this notion tracks exactly the function of the relativist's experiencer parameter. Where ϕ is experiential at w, proper truth assessment varies between individuals at w as their experiences relevantly diverge, again mirroring the relativist situation in §1.2.

So a bare treatment can distinguish between correct and proper truth assessment, and have them pattern together and apart depending on the proposition, in the same way that a relativist treatment does.[46] The norms in (96) pragmatically compel speakers to assess

46. Cf. Kölbel's (2004: 59-60) comments on mitigated realism, Schafer's (2011: 272-275) appeal to second-order norms on belief formation, and his distinction between being mistaken and being at fault, and Hirvonen's (2016: 60) distinction between an assertion's being true and justified.

propositions as true or false, and so to agree or disagree about them, in the same circumstances as do their relativist counterparts, with respect to the relevant relativist-proposition counterparts.[47]

A difference between the two views is that (96a) entails that where ϕ is experiential at w, it is correctly or incorrectly assessed as true at w, depending on whether $\phi(w)$; a relativist treatment denies the correlate of this (cf. (86) ff.). This might cause worry that a bare view entails that individuals have experiential properties at worlds as a 'matter of fact' in some objectionable sense. By contrast, the relativist treatment allows a formalization of the idea that there is not necessarily such matter of fact with respect to experiential properties (cf. (85)-(86) ff.). This 'lack of a matter of fact' is ultimately a consequence of the metasemantics of experiential predicates, to be addressed in Chapter 5, where this worry will be dispelled.

The relativist's experiencer parameter is a shorthand for dealing with a cluster of semantic-pragmatic phenomena surrounding experiential predicates. The rest of this work explains these phenomena in more fundamental terms, dispensing with the need for this shorthand. The behavior of experiencer phrases is cast in terms of the experiencer role in certain experiential states (in Chapters 2 and 3). Autocentric assessment is seen to be an epiphenomenon of experiential semantics (again as explicated in Chapters 2 and 3) interacting with conventions on direct evidentiality (to be explicated in Chapters 4 and 5). A relativist semantics is thus empirically unobjectionable on most grounds (though see §4.2.2 for an exception), but the move to a bare semantics will allow for the derivation as an epiphenomenon of what the relativist has stipulated.

47. This defuses Lasersohn's (2005: 655) worry that a bare semantics (his 'Option 3c;' cf. Lasersohn 2017: 32-34) can't account for speakers' justification in asserting, without uncertainty and based on their own experiences, that experiential properties hold of individuals: speakers can depend on these experiences for propriety of assessment and hence assertion regardless of the adicity of the metalangauge predicate.

CHAPTER 2

DEVERBAL PSYCH PREDICATES AND EXPERIENTIAL KINDS

Chapter 1 offered a treatment of experiential predicates according to which their extensions are not relativized to an experiencer. Nevertheless, experiential predicates do have denotations that make reference to experiences in a way that distinguishes them from non-experiential predicates. The present chapter, and Chapter 3, provide an account of the semantics of experiential predicates, spell out what exactly makes them experiential, and elucidate what their relationship to experiencers is.

Experiential predicates form a semantic natural kind: roughly, they are those predicates that denote a property true of an individual just in case that individual is disposed to produce experience of the relevant sort – *simpliciter*, and not in any particular experiencer or other. There are a host of lexical and phrasal compositional processes that make it perfectly predictable whether certain predicates are experiential, and in what way.

The subject of this chapter is a large group of adjectives in English, termed here *deverbal psych predicates*, which are formed regularly when the suffix *-ing* composes with object-experiencer psych verbs. Adjectives in this category include e.g. *stunning*, *frightening*, *interesting*, *confusing*, and *horrifying*. Providing a compositional account of these results in a templatic semantics for experiential predicates generally, which can be retrojected onto other apparently experiential predicates that exist simplex in the lexicon.

That deverbal psych predicates ought rightly to be considered experiential is shown by the fact that they have all the crucial features of experiential predicates outlined in Chapter 1, and the interested reader can run them through all the arguments presented there to confirm that they apply equally to them in all cases. To illustrate and recap, using the exemplar *frightening*: deverbal psych predicates occur with *to*-headed PPs acting as overt experiencers (100a), and in predicate position embedded beneath *find* (100b), and in each

case a presupposition of direct experience relevant to the predicate is enforced.

(100) a. The movie is frightening to Alfonse.

b. Alfonse finds the movie frightening.

↪ Alfonse has experienced fear of the movie.

They further exhibit no implications of direct experience when embedded in a wide variety of intensional contexts (cf. §1.1.2; see Ch. 1, fn. 15 on adressee-directed presuppositions involved in direct questions like (101c), and read (101e) as said of a movie still in production).

(101) a. The movie must be frightening.

b. If the movie is frightening, he shouldn't watch it.

c. Is the movie frightening?

d. I wonder whether the movie is frightening.

e. The movie will be frightening.

f. Alfonse believes that the movie is frightening.

They also give rise to exocentric readings, evaluated with respect to a contextually anchored experiencer ((102a)-(102b), reading (102b) as being uttered while the speaker is looking at Alfonse, who is clearly terrified, watching a movie), and these experiencers can be bound by local quantifiers ((102c), whose interpretation is, 'Every person x saw a movie frightening to x').

(102) a. Alfonse saw a movie in the theater. It was frightening.

b. The movie is frightening.

c. Everyone saw a frightening movie.

And they give rise (when read non-exocentrically) to speaker-oriented acquaintance inferences (cf. §1.1.5), which enforce non-presuppositional constraints on what experiences the speaker has had of the subject of predication, which mirror the content of the direct experience presuppositions illustrated above.

(103) The movie is frightening.

⇝ The speaker has experienced fear of the movie.

Finally, speakers tend to evaluate these predicates using their own experiential reactions, or dispositions to react, as outlined in §1.3.2, and so evaluate them by default autocentrically (again, where they are not read exocentrically), and the assessment of the truth of experiential propositions based on one's own experiences is tracked as expected by the behavior of truth-sensitive and belief-sensitive lexical items, like factives and verbs of agreement (cf. §1.1.3 and §1.1.4).

What makes English deverbal psych predicates especially worth studying is that they exist in a rich morphological domain, where they interact with object-experiencer psych verbs (e.g. *frighten*), nouns denoting *experiential kinds* (e.g. *fright*), adjectival passives (e.g. *be frightened*), and various other nominals (e.g. *a fright*). The way these predicates are treated is therefore constrained by the way these related expressions are treated, and the compositional relations that hold between all of them.

Most important for present purposes is the relationship between experiential kind terms, object-experiencer psych verbs, and deverbal psych adjectives. The morphological relatedness of these expressions is obvious, and comes in various superficial forms.

(104)

	Kind Term	Psych Verb	Psych Adjective
Type I	∅	*stun*	*stunning*
Type II	*fright*	*frighten*	*frightening*
Type III	*interest*	*interest*	*interesting*
Type IV	*confusion*	*confuse*	*confusing*
Type V	*horror*	*horrify*	*horrifying*

Some deverbal psych predicates have no corresponding experiential kind term (Type I); others are associated with an object-experiencer psych verb that relates morphologically to a corresponding kind term, with either the verb being apparently more complex (Type II), the verb and noun being identical in form (Type III), the noun being apparently more complex

(Type IV), or the verb and the noun apparently sharing a root (Type V). The behavior of *-ing* is highly regular, and virtually every object-experiencer psych verb in English can take it as a suffix, with the same semantic effect.[1] What follows in this chapter examines this regular route to experiential predicates.

§2.1 examines experiential kind terms, and discusses their formal properties, how they encode experiential states, and how these states ultimately relate to actual episodes of experience in experiencers. It is shown that experiences can be modeled as mass individuals of a certain sort, portions of which can be located 'in' experiencers and directed 'at' stimuli, and which can be ranked according to their size. An experience being located in this sense constitutes an experiential state, which itself is an eventuality triggered by an experiential episode, and which in virtue of persisting enforces a disposition to similar experiential episodes.

§2.2 discusses object-experiencer psych verbs, and shows how they can be treated as relating experiencers and stimuli, in order to assert the existence of experiential states as described above. It's shown that there are both stimulus-stative and causal-eventive readings of these verbs, and that it is the former rather than the latter that are relevant to the composition of experiential predicates. A sketch is then provided of a compositional semantics for stimulus-stative psych verbs, which accounts for their precise relation to the experiences that an experiencer has had and is disposed to have, and which accounts for their gradability in terms of the size of portions of experiential kinds they encode.

§2.3 then treats deverbal psych predicates themselves, and so provides a model for experiential predicates generally. The semantics of *-ing* is discussed, and how it desaturates the experiencer of the psych verb. These experiential predicates are then related to generic quantification over experiential states, which are disposed to arise in ideally normal conditions on the production of experience (yet with no intrinsic reference to an experiencer), and

1. Where there are apparent exceptions to this pattern, there is always an adjective of distinct form with a meaning identical to that which *-ing* would have produced: thus, **delighting* is preempted by *delightful*, **scaring* by *scary*, **awing* by *awesome*. It can thus either be maintained that *-ing* is regular in the domain of object-experiencer psych verbs except where lexically blocked by an irregular deverbal form, or that it is indeed perfectly regular, but the suffix has an idiosyncratic morphophonological realization on these verbs.

the relation between exocentric and non-exocentric predicates is elucidated on these terms. Finally, the role of overt experiencer PPs is clarified.

2.1 Experiential kind terms

By 'experiential kind term' is meant a noun that denotes a 'kind' of experience that it is possible for an experiencer to undergo. The class of these expressions is enormous, with little morphological homogeneity. Below a sample of twenty-five of these nouns, to which the present treatment is taken to apply.

(105) **Experiential kind terms**:

amazement, anger, astonishment, bewilderment, cold, concern, confusion, consternation, delight, displeasure, fear, fright, hate, horror, hunger, interest, love, pain, pleasure, puzzlement, resentment, satisfaction, shock, stupor, terror

Some of these are signaled by abstract nominalizers like *-ment* and *-or*, others are diagnosed by their relation to a psych verb, either subject-experiencer (e.g. *hate / hate*) or object-experiencer (e.g. *pleasure / please*). Still others (e.g. *consternation*) participate in no relevant morphological relations to other expressions. For present purposes, the most crucial of these terms are those that have a corresponding object-experiencer psych verb.

§2.1.1 explores the formal properties of experiential kind terms, treating them as abstract mass nouns that denote portions of experience. §2.1.2 then establishes how these relate to experiential states pertaining to a kind, and how these in turn relate to experiential episodes of the phenomenological character associated with that kind.

2.1.1 Formal properties

Experiential kind terms have three core properties that any treatment of them must take account of. First, they can be associated with stimuli, or individuals 'at whom' experience

is directed. Many kind terms are inherently capable of such direction, and encode this by selecting for PPs denoting a stimulus.

(106) a. fear of spiders

b. displeasure with the proceedings

c. interest in mathematics

d. shock at the news

Whether such a PP is available, and which preposition must head it, is apparently lexically idiosyncratic, and preserved in morphologically related constructions. Thus the appearance of the stimulus-phrase is apparently not semantically-governed (compare (107) with (106a)); some kind terms allow multiple prepositional heads (108); and adjectival passives tend to preserve the lexical properties of their associated kind terms (compare (109) with (108)).

(107) *fright of spiders

(108) anger {with / at} Alfonse

(109) angry {with / at} Alfonse

The stimulus relation can be described using the relation $stim''$, which is dyadic relative to worlds: '$stim''(w)(x)(y)$' is to be read, 'x is the stimulus of y at w,' and the term is defined only if y is an experience.

Second, experiential kind terms can be associated with experiencers, or individuals 'in whom' they reside, and who undergo experiences as a result of this. Experiencers are encoded by the genitive relation, e.g. via the Saxon genitive.

(110) a. Alfonse's anger (at Bethany)

b. Bethany's fear (of spiders)

The experiencer relation can be described using exp'', which is also dyadic relative to worlds: '$exp''(w)(x)(y)$' is to be read, 'x is the experiencer of y at w,' where again the term is defined only if y is an experience.

What these two relations amount to will be elucidated in what follows. One crucial feature must hold between them: where an experience has a stimulus, it must also have an experiencer, as it is impossible for an experience to be stimulated but not experienced.[2]

(111) If $stim''(w)(x)(y)$, then there is a z such that $exp''(w)(z)(y)$.

The third crucial feature of experiential kind terms is that they are mass nouns, as diagnosed by the usual tests. Unlike (uncoerced) count nouns, they occur with mass-sensitive determiners (112); with partitives (113); and with quantity questions (114).

(112) a. {little / much / no} {anger / pleasure / shock}

b. ?{little / much / no} {dog / phone / building}

(113) a. a {lot / bit / ton} of {terror / pain / satisfaction}

b. ?a {lot / bit / ton} of {party / dream / book}

(114) a. How much {interest / awe / bewilderment} did Alfonse experience?

b. ?How much {horse / car / cloud} did Alfonse see?

They are also typically strange with atomic-sensitive determiners, like *a(n)* and *every*. In (115), the only felicitous interpretation is that the noun refers to a certain kind of anger.

(115) a. ?angers

b. ?an anger

c. ?every anger

2. Other possible restrictions on these relations are not enforced. For instance, there is no guarantee that an experience is stimulated or experienced at any given world: this is because counterfactual reasoning may pertain to experiences that are not actually being experienced or stimulated. Further, it may be possible that an experience has an experiencer at a world, but no stimulus at that same world: this allows for coherent denotations for one reading of phrases like *anger at nothing*, and leave open the possibility of experiential kinds that are not inherently directable at stimuli.

Where there are apparent exceptions to this generalization, the experiential kind term tends to take on an alternate denotation. Thus the nouns in (116a) refer to individuals that tend to produce the kind horror, not to that kind itself (and so these nouns are experiential predicates in their own right). And those in (116b) (as in, *Alfonse received a shock*) have an eventive interpretation, and refer to countable experiential episodes of shock, again not to the experiential kind (cf. §2.1.2).

(116) a. horrors, a horror, every horror

b. shocks, a shock, every shock

As with other mass nouns, experiential kind terms can thus be taken to denote the property of being a 'portion' of a certain kind: just as *water* denotes the property of being a portion water, *interest* denotes the property of being a portion of interest, and so on.

So let there be a range of experiential kinds k, denoted in the metalanguage by terms in small caps, including FEAR, CONFUSION, INTEREST, SHOCK, etc. The denotations of experiential kind terms are then schematic, as below, where e.g. $fear'(x)$ iff x is a portion of fear, the kind denoted in the metalanguage by 'FEAR.'[3]

(117) a. $[\![\text{fright}]\!]^w = \lambda x_e.fear'(x)$

b. $[\![\text{shock}]\!]^w = \lambda x_e.shock'(x)$

c. $[\![\text{interest}]\!]^w = \lambda x_e.interest'(x)$

For simplicity, these terms are taken to lack internal arguments, of the sort that might be encoded by e.g. *of spiders*: a full treatment ought to incorporate these as well, to account for their lexical idiosyncrasy and to describe how they are preserved in other constructions, e.g. with adjectival passives (*frightened of spiders*).

3. For simplicity, this 'portion-of' property is taken to be world-invariant: thus, a portion of k is necessarily such a portion. As with many mass nouns, experiential kind terms in argument position with a null determiner seem to denote the kind itself, rather than the property of being a portion of this kind: thus, *Fear is a powerful emotion*, where $[\![\varnothing \text{ fear}]\!]^w$ = FEAR, or some such. The text won't address these constructions, but presumably the treatment of mass nouns denoting kinds should generalize here.

Experiences can then be construed as individuals, of type e, that are mereologically structured like the individuals to which the denotations of typical mass nouns apply. Following Link (2002 [1983]: 130), let the domain of individuals $\mathcal{D}_e$ be closed under the binary operation $\oplus$, where '$x \oplus y$' is to be read, 'the individual sum of x and y,' and denotes that individual comprised of x and y, and only these, as mereological parts. $\oplus$ has the usual properties of commutativity, associativity, and idempotence, and forms a complete join-subsemilattice of the set of atomic individuals (those without proper mereological parts in this sense) of $\mathcal{D}_e$.

A binary 'individual part' relation $\sqsubseteq$ can then be defined as follows (cf. *ibid.* 130-131), where '$x \sqsubseteq y$' is to be read, 'x is a mereological part of y.'

(118) $x \sqsubseteq y$ iff $x \oplus y = y$.

The predicates in (117) can then allow for cumulative and divisive reference with respect to experiential kinds: for any two individuals that are a portion of an experiential kind, their mereological sum is a portion of that kind; and for any individual that is a portion of an experiential kind, all of its mereological parts are portions of that kind as well.

(119) Where $P(x)$ iff x is a portion of experiential kind k:
$P(y)$ and $P(z)$ iff $P(y \oplus z)$.

For convenience, it can also be stipulated that experiential kinds are disjoint, such that for any distinct P, P' that denote the property of being portions of distinct experiential kinds k, k', there is no x such that $P(x)$ and $P'(x)$.

Finally, a mereological maximal operator σ can be introduced (cf. *ibid.* 131, ex. 17), that picks out the maximal individual that satisfies some descriptive condition. It is defined as follows, where ϕ is a metavariable over formulae, and '$\phi[y \to x]$' denotes the value of ϕ on an assignment of the individual y to all free occurrences of the individual variable 'x.'

(120) $\sigma x[\phi]$:= that unique individual y such that:
a. $\phi[y \to x]$;

b. For no z such that $y \sqsubset z$: $\phi[z \to x]$.

In other words, $\sigma x[\phi]$ is the unique mereologically maximal individual such that ϕ. With this machinery in place, experiences that are portions of certain kinds, and which have certain stimuli and experiencers, can be denoted using type-e expressions. For example, the composition of *Alfonse's fear of spiders* runs as follows, for simplicity taking *of spiders* to be a modifier, in line with the caveat about internal arguments above, and taking the genitive to act as a maximizing definite descriptor.

(121) a. $[\![\text{of}]\!]^w = \lambda x_e.\lambda P_{et}.\lambda y_e.P(y) \wedge stim''(w)(x)(y)$

b. $[\![\text{of spiders}]\!]^w = [\![\text{of}]\!]^w([\![\text{spiders}]\!]^w)$
$= \lambda P_{et}.\lambda y_e.P(y) \wedge stim''(w)(\textsc{spider})(y)$

c. $[\![\text{fear of spiders}]\!]^w = [\![\text{of spiders}]\!]^w([\![\text{fear}]\!]^w)$
$= \lambda y_e.fear'(y) \wedge stim''(w)(\textsc{spider})(y)$

d. $[\![\text{-'s}]\!]^w = \lambda x_e.\lambda P_{et}.\sigma y[P(y) \wedge exp''(w)(x)(y)]$

e. $[\![\text{Alfonse's fear of spiders}]\!]^w = [\![\text{-'s}]\!]^w([\![\text{Alfonse}]\!]^w)([\![\text{fear of spiders}]\!]^w)$
$= \sigma y[fear'(y) \wedge stim''(w)(\textsc{spider})(y) \wedge exp''(w)(a)(y)]$

Thus, *Alfonse's fear of spiders* denotes the mereologically maximal individual that is a portion of the kind FEAR, which is stimulated by the kind SPIDER, and which is experienced by Alfonse.

Experiential kind terms further denote portions of *qualities*, in the sense of Francez & Koontz-Garboden (2017). This can be seen from the fact that they occur in a range of constructions that other abstract mass nouns denoting qualities do, but that concrete mass nouns do not (cf. *ibid.* §6.2 for these and related diagnostics). They occur with *what*-questions, both exclamatives and questions, that target the size of a portion (122); with modifiers targeting an upper end of a size ordering (123); with modifiers of large size, like *great* (124); and with verbs of changing size, like *grow* (125).

(122) a. What {excitement / amazement} Alfonse felt!

b. I didn't realize what {excitement / amazement} Alfonse felt.

c. ?What {snow / water} Alfonse saw!

d. ?I didn't realize what {snow / water} Alfonse saw.

(123) a. {utter / complete / absolute / outright} {horror / astonishment}

b. ?{utter / complete / absolute / outright} {sand / chalk}

(124) a. great {puzzlement / confusion}

b. ?great {iron / gold}

(125) a. Alfonse's {anger / grief} grew.

b. ?Alfonse's {tea / sauce} grew.

Francez & Koontz-Garboden propose that it is distinctive of nouns denoting abstract qualities, as opposed to concrete mass nouns, that their lexical entries encode an inherent size ordering, which these expressions are capable of targeting. Experiential kind terms apparently fall into this category.

To capture this, a 'size' ranking on qualities relative to a kind k, $\preceq_k$, is introduced, where '$x \preceq_k y$' is to be read, 'x is at least as small as y with respect to k.' This ranking preorders individuals according to how 'big' of an experience of a certain kind they comprise. The lack of antisymmetry on the preordering allows 'ties' among indivdiuals, which are distinct but nevertheless of the exact same size, e.g. when Alfonse's fear is just as much as Bethany's, but nonetheless *Alfonse's fear* and *Bethany's fear* are not coreferential.

(126) $\preceq_k$ is a total preorder on $\mathcal{D}_e$ such that:

a. $\neg\exists y[y \sqsubseteq x \wedge y$ is a portion of $k]$ iff x is a minimal element according to $\preceq_k$;

b. $\preceq_k$ preserves $\sqsubseteq$: if $x \sqsubseteq y$, then $x \preceq_k y$.

The preorder on qualities is constrained in two ways. (126a) says that any individual that is composed of no part that is a portion of kind k (that is, which has no such part as a mereological constituent) is minimally ranked by the preorder. And so an individual that

contains no portion of a kind at all is as small as possible with respect to the 'size' of that kind – it follows trivially that any individual that is not an experience, or has no experience as a mereological part, is so minimally ranked. (126b) says that an individual is always at least as big according to the size preorder as all of its mereological parts (cf. Francez & Koontz-Garboden *ibid.* 39, B).

These two conditions do not determine a unique total preorder on the domain of individuals. This is as it should be, since such a preorder cannot be deduced from the formal properties of experiences alone. Further decisions must be made as to how to 'rank' individuals according to 'how big' a portion of experience is. This might have to do, among other things, with the intensity of experiential episodes produced in an experiencer, their frequency, the effort required in triggering them, the extent to which they approximate 'canonical' experiences of the kind phenomenologically, and so on. These matters are likely not linguistically encoded at all, and different speakers will have different opinions as to how exactly the preorder functions – and in many cases, it may be left indeterminate or even inconsistent according to changing communicative needs. This indeterminacy of the preorder on qualities is thus a source of 'weak subjectivity,' as characterized in §5.3.2.

This formal setup ultimately allows for a way to capture truth conditions comparing sizes of experiences, as with *Alfonse's fear of spiders is greater than Bethany's.* But an incorporation of these kinds of comparative and degree constructions requires a treatment of degrees constructed out of this preorder: cf. §2.2.2.

There are two extant metaphors that receive a lot of use cross-linguistically in encoding the relation between experiencers and experiences, using experiential kind terms. The first is locative: portions of these kinds are said to be located 'in/on/at' individuals, as in Irish (127).[4] The second is possessive: individuals are said to 'have' portions of these kinds, as in

4. For Landau (2010a), these locative constructions are telling, and he takes many experiencers acting as objects of psych verbs also to be syntactically locative: they are oblique, and undergo locative inversion. The invocation of these spatial metaphors in the text shouldn't be taken to endorse these stronger syntactic claims. A second, rarer sort of locative metaphor, discussed e.g. in Arad (1998: 228), has the experiencer located in a portion of the experiential kind, rather than vice-versa: thus, *Alfonse is in (love / shock).* The

the familiar range of limited Romance and Germanic examples, e.g. in Spanish (128).[5]

(127) Tá eagla orm.
is fear on-me
'I am afraid.' [lit. 'Fear is on me.']
[Irish. McCloskey & Sells (1988: 182, ex. 78a)]

(128) Tengo miedo.
have.1SG.PRES *fear*
'I am afraid.' [lit. 'I have fear.']
[Spanish.]

These metaphors can be employed more or less grammatically: where not grammaticized, they often appear in a number of unsystematic periphrastic constructions. The present treatment will assume that both encode the very same relation, and that this is the same relation encoded by exp''. Thus '$exp''(w)(x)(y)$' can also equally be read, 'y is in x at w' or 'x has y at w,' where again its definedness conditions require that y be an experience. The genitive relation therefore defaults to exp'' where experiences are concerned, as in (121d).

As to the relations between stimuli and experiences, these are typically encoded using some directional metaphor (129) or relational marker (130).

(129) Tá eagla roimh Y ar X.
is fear before Y *on* X
'X is afraid of Y.' [lit. 'There is fear before Y on X.']
[Irish. McCloskey & Sells (1988: 181, ex. 77a)]

(130) Tengo miedo de morir.
have.1SG.PRES *fear of die*
'I am afraid of dying.' [lit. 'I have fear of dying.']
[Spanish.]

---

text won't address this, or whether it's semantically important, but it ultimately ought not to be ignored.

5. Given that experiential kinds are qualities in Francez & Koontz-Garboden's (2017) sense, these possessive metaphors ought not to be surprising, since as they demonstrate, the relation of individuals to their qualities cross-linguistically is often encoded using possessive morphology.

While the directional metaphors vary, they often pertain to a portion of an experiential kind being 'aimed at' its stimulus (cf. the other examples in McCloskey & Sells 1988: 181 ff.).[6] As such, '$stim''(w)(x)(y)$' can also equally be read, 'y is directed at x at w' or 'y pertains to x at w.'

The core notions of experiences as individuals that (i) are possibly portions of experiential kinds, (ii) are pre-ordered according to size relative to these kinds, (iii) are possibly associated with experiencers, and (iv) are possibly associated with stimuli, with this entailing association with an experiencer, are now in place, and provide a working ontology to be used in describing the existence of experiences in experiencers. The relations exp'' and $stim''$, which describe the properties of these experiences, may be thought of as encoding eventualities, and in particular experiential states. Thus, to say that $stim''(w)(x)(y)$ and $exp''(w)(z)(y)$, where $fear'(y)$, is to say that there is an eventuality at w, which is a spatial state, and it consists of y, which is a portion of fear, being located in z and directed at x. §2.1.2 clarifies what the existence of such a state entails.

2.1.2 Experiential states and episodes

In the following, much hinges on how being in an experiential state, as described in §2.1.1, is characterized, and in particular on how being in such a state relates to undergoing 'experiential episodes.' These latter are events during which the experiencer actually feels something with the phenomenological character of the relevant kind: the conscious (or at least sentient) undergoing of fear, shock, confusion, etc.[7]

The most obvious way to relate experiential states and episodes is to equate them. Where

6. These directional metaphors seem to be reflected in directional modifiers that nouns denoting experiences can take, unlike nouns denoting entities with no inherent direction: thus, *anger towards Alfonse* may denote anger stimulated by Alfonse, while ?*dog towards Alfonse* is strange. The relation denoted by *towards* may therefore default to $stim''$ in the case of experiencers.

7. The text takes for granted that the crucial feature of these episodes is phenomenological, and does not touch on broader issues about whether and to what extent experiential episodes are to be characterized in terms of their functional role, or by some other criterion. The semantics will be the same, regardless of how the distinction between 'being in an experiential state' and 'having an experiential episode' is cashed out precisely: what matters is that the distinction preserves at least the properties described in the text.

there is an x such that $stim''(w)(y)(x)$ and $exp''(w)(z)(x)$, this means that x is that experience stimulated by y and experienced by z at w, and this can be taken to mean the same as that z is actually undergoing the relevant felt experience at w. But this isn't right, since none of the expressions morphologically related to an experiential kind term enforce this requirement. All of the sentences in (131) assert, or presuppose, that Alfonse is in an experiential state of interest, directed at linguistics.

(131) a. Bethany knows about Alfonse's interest in linguistics.

b. Linguistics is an interest of Alfonse's.

c. Alfonse is interested in linguistics.

d. Linguistics interests Alfonse.

e. Linguistics is interesting to Alfonse.

But none of them, either with a genitive denoting his experience (131a), an experiential nominal (131b), an adjectival passive (131c), an object-experiencer psych verb (131d), or a deverbal psych adjective (131e), require for their truth or definedness conditions that Alfonse actually be undergoing the experience associated with his interest as the time of utterance. They can all be truthful reports, for example, even while Alfonse is asleep, and not consciously directing his interest at anything.

But Alfonse must be disposed in the right way to have experiential episodes of the phenomenological character associated with INTEREST, and directed at linguistics, for these reports to be true. What conditions must occasion these episodes is difficult to articulate, but roughly they must be those in which Alfonse's experiential faculties are suitably focused on linguistics – if Alfonse generally feels interest whose phenomenological 'target' is linguistics in such conditions, then the reports come out true.[8]

8. For an experiencer's 'experiential faculties' to be 'focused' on an individual does not entail that the experiencer is presented with the individual 'in the flesh,' or in causal contact with it: imagination of it, remembrance of it, etc. often suffices. This is especially important with abstract objects like 'linguistics,' as well as non-existent things, which one may nonetheless fear, etc. (cf. §2.2.1). Of course, being in causal contact with something may be a perfect occasion for having these faculties directed at it: if Alfonse is

It is irrelevant to the truth of a report of an experiential state whether this generic disposition (which will be subject to the usual constraints on genericity: cf. §2.3.1) is actualized at the time of the report. In this sense, the linguistic encoding of experiential states is much like that of other psychological states: e.g., Alfonse has a belief in a proposition just in case he is disposed to treat that proposition as true, regardless of whether he is presently actualizing that disposition (again, reports of Alfonse's beliefs can be true even while he is asleep).

Nevertheless, this generic disposition is not sufficient for an experiential state to obtain: it must also be occasioned by an actual experiential episode. Thus, none of the reports in (131) can be true where Alfonse has never actually felt interest in linguistics: it does not suffice that he would feel such interest, if he considered the subject appropriately. This point can be made vividly by the definedness conditions of genitives denoting experiences as in §2.1.1: *Alfonse's fear of heights* denotes nothing unless Alfonse has actually experienced heights (at least in his imagination), and this has resulted in an experiential episode of fear. Prior to having such an episode, this phrase denotes nothing, because then Alfonse does not *have* a fear of heights.[9]

Experiential states are thus characterized as follows. These notions might be formalized using modal machinery modeling dispositions, but as their formal properties are not relevant to what follows, they are left informally paraphrased.

(132) If $y \sqsubseteq x$, and y is a portion of k:

$stim''(w)(z)(x)$ and $exp''(w)(u)(x)$ only if at w, u has undergone an experiential episode of the phenomenological character associated with k and directed at z, in virtue of which u is generically disposed to undergo experiential episodes with that same phenomenological character and direction, when its experiential faculties are

frightened by a movie, then showing that movie to him is likely to make him focus on the movie, and so under the right conditions to feel fear of it.

9. There are some possible exceptions to this observation: a deity that creates a human *ex nihilo* and knows all of its faculties down to the last detail might truly boast, 'I made this one with a fear of heights,' even prior to its having experienced any heights. These won't be addressed in the text.

focused on z.[10]

How these experiences are ranked according to the preroder on qualities will then depend on what sorts of experiential episodes occasioned the experiential state, and what sorts of experiences the individual is generically disposed to have towards the stimulus. This will be determined, in a way not linguistically encoded, by the considerations mentioned in §2.1.1.

This characterization is important, because it follows that attribution of an experiential state to an individual entails having had a certain sort of experience of its stimulus, where there is one. From this it ultimately follows that certain constructions enforce, or fail to enforce, direct experience presuppositions in virtue of how they make reference to experiential states.

2.2 Object-experiencer psych verbs

The target of this section is a group of psych verbs included in the '*preocupare*-class' of the tripartite division in Belletti & Rizzi (1988) (cf. Landau's 2010a 'Class II' verbs), which have a nominative subject and accusative object: these include, per §2.1, *stun*, *frighten*, *interest*, *confuse*, and *horrify*. These verbs are of interest for the reference to experiential kinds that their denotations make, which is often made apparent by their morphological relation to an experiential kind term. Further, since they occur productively with the *-ing* suffix to form experiential adjectives, their analysis is helpful in demonstrating how the semantics of experiential predicates (de)composes.

As noted in §2.1, where five superficial morphological classes of object-experiencer psych

10. If desired, the extension to experiential states with no stimulus (cf. fn. 2) is easy to describe, by removing the language addressing the stimulus and the 'direction' of the experience. The 'in virtue of which' language is meant to enforce a link between the initial experiential episode and the generic disposition. There may be some cognitive reason that experiential states are encoded in this way, with an initial experiential episode being what 'plants' the portion of an experiential kind in the experiencer, and this portion then being able to 'activate' once the experiencer possesses it. It may also be that the natural language metaphysics, in Bach's (1986) sense, is committed to such individuals, that 'appear' in experiencers in virtue of an initial experiential episode, and 'remain' dormant in them for as long as the disposition to undergo further such episodes remains.

verbs were identified in English, these verbs have an inconsistent surface relationship with their experiential kind term counterparts, if they have any. Given this, a number of options present themselves as to how to analyze the composition of each class. In particular, if the surface morphology is taken at face value, then one might e.g. take *confuse* to be verbal in the lexicon, and give a special entry for a nominalizer *-ion*, while taking *fright* to be nominal in the lexicon, and give a special entry for a verbalizer *-en*. However, given the semantic machinery to be provided in what follows, the ability to move between the appropriate nominal and verbal denotations in any direction will be trivially granted, and since all the psych verbs will end up with analogous denotations (relative to the experiential kind associated with them), these morphological questions will not matter at present.

The following therefore adopts the simplifying, unifying assumption that each object-experiencer psych verb is the result of composing a verbalizer with a root, and that each nominal experiential kind term is the result of composing a nominalizer with that same root. The verbalizer and nominalizer will then have a distinct morphophonological realization depending on the superficial class of the psych verb. Thus each verb is treated on the model of the surface behavior of *horrify*, which looks to be the result of composing the verbalizer *-ify* with some root $\sqrt{\text{HORROR}}$, while *horror* is the result of the nominalizer *-or* composing with this same root. This is meant only as an expository convenience, awaiting a more serious morphological analysis, which may or may not be unified this way.

(133)

Root	Nominalizer	Verbalizer
$\sqrt{\text{STUN}}$: *stun*	n/a	-∅
$\sqrt{\text{FRIGHT}}$: *fright*	-∅	*-en*
$\sqrt{\text{INTEREST}}$: *interest*	-∅	-∅
$\sqrt{\text{CONFUSION}}$: *confuse*	*-ion*	-∅
$\sqrt{\text{HORROR}}$: *horr*	*-or*	*-ify*

The roots can be taken to denote the property of being a portion of the relevant kind, and the nominalizers can be taken to be vacuous, preserving the denotations from §2.1.1.

(134) a. $[\![\sqrt{\text{FRIGHT}}]\!]^w = \lambda x_e.fear'(x)$

b. $[\![\text{-}\varnothing]\!]^w = \lambda P_{et}.P$

c. $[\![\text{fright}]\!]^w = \lambda x_e.fear'(x)$

The following section offers a semantics for the verbalizers, and hence the psych verbs. §2.2.1 isolates the class of object-experiencer psych verbs to which this analysis is meant to apply, separating their stative occurrences with genuine stimuli subjects from their related eventive readings with causer subjects. §2.2.2 then provides the analysis of the verbalizers themselves.

2.2.1 Causers versus stimuli

The majority of object-experiencer psych verbs in many languages have two recognizably distinct readings, and it is important for what follows to focus on the reading relevant to the composition of experiential predicates. The distinction can be seen in English with the example of *frighten* in the present tense.

(135) Alfonse frightens Bethany.

(135) can be read either as aspectually stative, or eventive. On the former reading, the verb is construed non-habitually, and what's asserted is that Bethany is in an experiential state of fear directed at Alfonse, of the sort described in §2.1.2. On the latter reading, the verb is eventive, and so due to the present tense habitual: it asserts that Alfonse habitually causes Bethany to experience fear. Aside from this non-habitual/habitual distinction, the classical tests for stativity e.g. in Dowty (1979) can be applied to such sentences, to verify that both stative and eventive readings are allowed. The subject of a stative object-experiencer psych verb is a *stimulus*, while the subject of its eventive counterpart is a *causer*.

A syntactic distinction between stimulus-stative and causative-eventive readings has long been recognized, in that the former, but not the latter, license certain 'psych-effects,' which

are unusual syntactic behaviors pertaining specifically to psych verbs. For example, Anagnostopoulou (1999) shows that in modern Greek, object-experiencer psych verbs embedded in relative clauses sometimes require resumptive pronouns to appear; this effect disappears where the subject is read as the agent-causer of the verb, rather than its stimulus.

Landau (2010a: chs. 3, 5) includes an extensive catalogue of these psych-effects, and demonstrates that in case after case, certain causal-eventive readings of the verbs do not give rise to them (cf. *ibid.* §9.3). On this basis, Landau holds that these two readings of object-experiencer psych verbs have distinct structures: causal-eventive readings are transitive, and include a v-projection and external argument, while stimulus-stative readings are unaccusative, with oblique objects (cf. *ibid.* 7-8, esp. fn. 4 & ex. 12a).[11] Of interest at present are not these syntactic distinctions, but the semantic ones that accompany them. Four of these are salient: a habitual reading in the preset tense (and hence an eventive reading) tracks the behavior of a causer as opposed to a stimulus according to each of these criteria.

The first difference relates to an observation made by Pesetsky (1995: 56ff.), that causers of experiences need not have the caused experiences directed at them, while stimuli (what he calls 'themes' or 'subject matters') must be the target of experiences that they stimulate. Pesetsky claims that subjects of object-experiencer psych verbs always occupy the causer role, but this is apparently mistaken.[12]

Where (135) is read statively and non-habitually, Bethany's fear must be directed at Alfonse. Where it is read eventively and thus habitually, the interpretation is just that Alfonse habitually causes fear to be in Bethany – but he might do this e.g. by playing scary movies for her, and it is these movies, and not Alfonse, towards which Bethany's fear is

11. Landau's (cf. *ibid.* 87-88, exx. 168a-b) typology is actually slightly more complicated than this: he splits the category of causative (non-unaccusative) object-experiencer psych verbs into an agentive category and a non-agentive category. It is the agentive psych verbs which lack the characteristic psych effects and oblique object. This finer difference won't matter for present purposes.

12. And so Pesetsky's (*ibid.*) treatment of object-experiencer psych verbs generally using a null morpheme CAUS, as well as related comments made e.g. by Jackendoff (1990: 262, fn. 4), may conflate the role of causer and stimulus in these constructions.

directed. Thus, *Alfonse frightens Bethany* entails *Bethany is frightened of Alfonse* only on the first reading. It is possible to be a causer of experience, without being its stimulus, and the latter role is tracked only by the stative readings.

The second observation moves in the opposite direction: a stimulus of an experience need not be its cause, and the latter is tracked by the subject of the psych verb only in the eventive reading. This is demonstrated most vividly in cases where there is a stimulus of experience that does not exist, and which *a fortiori* doesn't have the causal powers necessary for causing that experience.

(136) Bigfoot frightens Alfonse.

Assuming that Bigfoot does not exist, a stative and non-habitual reading of (136) is fine, so long as Alfonse imagines or otherwise has his attention directed towards Bigfoot, which generally makes him feel fear toward this imagined object. An eventive and habitual reading, however, says that Bigfoot habitually causes fear in Alfonse, which attributes various causal powers, and so existence, to Bigfoot.[13] Stimuli only need to have an experiencer's experiential faculties directed towards them, and stimulation does not imply causation.

Third, causal-eventive readings of object-experiencer psych verbs do not imply the existence of a corresponding experiential state in the sense of §2.1.2, only the existence of an experiential episode. This is not true of stimulus-stative readings, whose truth requires that such a state obtain. And so (135) on its eventive, habitual reading does entail that Bethany habitually undergoes experiential episodes of fear, caused by the causer; but it does not entail that she is frightened of anything.

Fourth and finally, stimulus-stative readings of these verbs are gradable, while causal-eventive readings are not: this will be addressed in §2.2.2. The contrast between these readings is brought out sharply by the fact that certain object-experiencer psych verbs simply lack a causative-eventive reading: (137a) cannot be read eventively and habitually, and the

13. The imagined scenario is one in which Alfonse is frightened by the thought of Bigfoot, and not e.g. by some fictional depiction of him, which might itself be said to exist and have various causal powers.

above four features characteristic of the stimulus-stative reading always hold of it. There are further cases in which the verb in its causal-eventive reading overtly encodes the distinction between causer and stimulus, in virtue of maintaining a stimulus argument alongside the causer: (137b) must be read with Alfonse as the causer, the project as the stimulus, and Bethany as the experiencer.

(137) a. Alfonse concerns Bethany.

b. Alfonse interested Bethany in the project.

$\rightarrow$ Bethany was interested in the project.

$\not\rightarrow$ Bethany was interested in Alfonse.

The exact relation between stimulus-stative and causal-eventive object-experiencer psych verbs will not be addressed here: it may be that the latter are derived via causative morphology of some sort.[14] What is important is to recognize that the stimulus-stative readings are the ones important for the semantics of experiential predicates. This can be seen from the fact that where these predicates are composed from psych verbs using the suffix *-ing*, all the semantic effects of the stative readings are preserved, as the reader can confirm.

(138) The movie is frightening to Alfonse.

Thus (138) can only be aspectually stative; it requires that Alfonse's fear be directed at the movie; it does not require that the movie be the cause of Alfonse's fear (say, if the movie is not real, but Alfonse is frightened at the thought of it); it requires that an experiential state of fear hold with Alfonse as experiencer; and the associated construction must be gradable. Deverbal psych predicates are thus composed out of stimulus-stative readings of

14. This is suggested by the fact that in many languages, formation of psych verbs occurs using overt causative morphology attached to predicates that relate their subjects to experiential episodes: Akatsuka (1976) demonstrates this for Japanese, and suggests the retrojection of this pattern onto languages like English that show no such overt morphology in these constructions. An account of these facts require figuring out whether constructions with overt causative morphemes only allow for the causal-eventive readings explained in the text: if so, then causative morphology is likely the culprit behind this distinction. If they allow for both sorts of readings, however, then a more nuanced view must be taken of the causative morphology, allowing it to express both stimulus and causer roles in the above sense, and some other source of the difference must be found.

psych verbs, and the relation of experiential predicates to experiential kinds generally mirrors that enforced by these readings of psych verbs.

2.2.2 Composition with psych verbs: stimulus, experiencer, and gradability

With the appropriate target of analysis clarified, a denotation can be given for stative object-experiencer psych verbs. The core idea is simple: these verbs compose with two arguments, an experiencer and a stimulus, and place them in the role of experiencer and stimulus as characterized in §2.1.1.

There are a few key features of these verbs that an analysis in terms of experiential kinds must capture. First, they are gradable, as shown by the fact that they can take degree modifiers like *really* and *a lot* (139), and occur in the comparative, using *more than* and *less than* (140).

(139) a. Mathematics really interests Bethany.

b. The news worries Alfonse a lot.

(140) a. The teacher's behavior shocks us more than the students' does.

b. His opinions confuse Bethany less than they confuse Alfonse.

When read statively, these sentences all receive the appropriate degree interpretation: (139a) means that mathematics interests Bethany to a high degree, (140a) means that the degree to which the teacher's behavior worries the group including the speaker exceeds the degree to which the students' behavior does, etc. Where these verbs are read causally and eventively, these interpretations are unavailable, and the modifiers pertain instead to the frequency of an event: an eventive reading of (139b) means instead that the news causes worry to be in Alfonse frequently, an eventive reading of (140b) means that the referent's opinions cause confusion to be in Bethany less often than they do in Alfonse, etc.

And so there needs to be a way of formalizing the degree to which an individual can be the stimulus of an experience of a certain experiential kind in an experiencer. Using a

simplification of the machinery in Kennedy (1997), each experiential kind can be associated with an ordered set of degrees (of type d), which the verbs will ultimately map experiences to.

(141) For each experiential kind k, there is:

a. a set of degrees Δ_k, where for all k', k'' such that $k' \neq k''$: $\Delta_{k'}$ and $\Delta_{k''}$ are disjoint;

b. a total order $\leq_k$ on Δ_k.

Adapting the Kennedian degree semantics to verbs with two arguments, a psych verb takes an experiencer and a stimulus, and returns the degree to which the stimulus stimulates an experience of the associated experiential kind in the experiencer. That is, psych verbs denote transitive measure functions, from individuals to individuals to degrees (and so are of extensional type $\langle e, \langle e, d\rangle\rangle$).

Experiences can then be 'ranked' according to the orderings on degrees described in (141). These rankings should preserve the preorders on portions of qualities described in §2.1.1: to instantiate an experiential kind 'to a higher degree' is just to be composed of a greater portion of that kind. To formalize this, let 'δ_k' denote a function that operates on an individual (in all interesting cases, an experience), and returns the degree of the experiential kind k that this individual instantiates. This function is defined as follows.

(142) δ_k is a function $\mathcal{D}_e \to \Delta_k$ such that for all x, y:

$\delta_k(x) \leq_k \delta_k(y)$ iff $x \preceq_k y$.

In other words, $\delta_k(x)$ is the degree of k that x instantiates, and the ranking among degrees tracks exactly the rankings among portions of experiential kinds: the greater the portion, the higher the degree. However, since $\leq_k$, unlike $\preceq_k$, is a total order, this means that distinct experiences that are 'tied' according to the size preordering on portions will instantiate the exact same degree.

One last bit of notation will be useful in giving a denotation for object-experiencer psych verbs. Let κ' be a function defined as follows.

(143) If for all w, x: $P(w)(x)$ iff x is a portion of k,
then $\kappa'(P) = k$;
else undefined.

κ' is thus a function that 'fishes out' an experiential kind from its corresponding property: where P is the denotation of *fright*, $\kappa'(P) = \text{FEAR}$, where P is the denotation of *confusion*, $\kappa'(P) = \text{CONFUSION}$, and so on.[15]

With this machinery in place, a denotation for the verbalizer that produces an object-experiencer psych verb can be given. The example used here is *-en*, which attaches to $\sqrt{\text{FRIGHT}}$ to form *frighten*: an exactly analogous denotation will be available for every such psych verbalizer, subject to the morphophonological caveats in §2.2 above.

(144) $[\![\text{-en}]\!]^w = \lambda P_{s,et}.\lambda x_e.\lambda y_e.\delta_{\kappa'(P)}(\sigma z[stim''(w)(y)(z) \wedge exp''(w)(x)(z)])$

-en composes with a term denoting the property of being a portion of an experiential kind (P), and returns a transitive measure function. This function itself takes two arguments, an experiencer (x) and a stimulus (y), and finally returns a degree, viz. the degree of the kind associated with P ($\kappa'(P)$) that the maximal individual (which due to the definition of 'exp''' must be an experience) stimulated by y and experienced by x instantiates. *-en* composes as follows with the root.

(145) $[\![\sqrt{\text{FRIGHT}}]\!]^w = \lambda x_e.fear'(x)$

(146) $[\![\text{frighten}]\!]^w = [\![\text{-en}]\!]^w([\![\sqrt{\text{FRIGHT}}]\!]^w)$
$= \lambda x_e.\lambda y_e.\delta_{\text{FEAR}}(\sigma z[stim''(w)(y)(z) \wedge exp''(w)(x)(z)])$

Since the experiential kind associated with *fright* is FEAR, the result is a transitive measure function that takes an experiencer and a stimulus, and returns the degree of fear in-

15. The uniqueness of k here, which allows κ' to be a function, is guaranteed by the disjointness of qualities: cf. §2.1.1.

stantiated by the maximal individual stimulated by the stimulus in the experiencer. In other words, *frighten* relates a stimulus and an experiencer to the degree of fear produced by the stimulus in the experiencer.

This denotation must be accompanied by some treatment of degree constructions, and for this to happen, one further feature of these verbs needs to be illustrated. Stative object-experiencer psych verbs impose orderings on their associated sets of degrees with a certain formal property. Adopting an analogy with gradable adjectives, and using Yoon's (1996) terminology, these verbs, on composing with their experiencer object, produce 'partial' predicates. In the terminology of Kennedy & McNally (2005), these are 'absolute' gradable predicates, and further they are so in virtue of their associated orderings encoding a 'lower-closed' scalar structure.

This means that in their positive forms (without any overt degree morphology or modifiers), these verbs denote relations that hold between a stimulus and and an experiencer just in case the stimulus stimulates some non-minimal degree of the relevant experiential kind in the experiencer. They do not, as they would if they were relative predicates, denote relations that hold between a stimulus and experiencer just in case the former stimulates a degree above some standard or threshold relative to a comparison class in the experiencer. This occurs in virtue of the fact that each ordering on the set of degrees associated with an experiential kind contains a least element, i.e. a unique 'lowest' or 'minimal' degree.

This is shown by the fact that these verbs occur with modifiers like *slightly / a bit* and *hardly / barely*, and when they do, the interpretation is not that the experience is stimulated to some degree just above a contextually relevant threshold, but rather that the experience is stimulated to some degree just above 'zero.'

(147) a. The movie frightens Alfonse {slightly / a bit}.

b. The movie {hardly / barely} frightens Alfonse.

↔ The movie stimulates fear in Alfonse to a slight degree.

↮ The movie stimulates fear in Alfonse to a degree slightly surpassing some

non-zero standard.

This is not how relative predicates like the adjective *tall* work (148), but it is how absolute lower-closed predicates like *open* work (149).

(148) Alfonse is {slightly / barely / a bit} tall.

⇏ Alfonse has height to a slight degree.

↔ Alfonse has height to a degree slightly surpassing some non-zero standard.

(149) The door is {slightly / barely / hardly} open.

↔ The door is open to a slight degree.

⇏ The door is open to a degree slightly surpassing some non-zero standard.

Further, sentential negation with these verbs creates the interpretation that the experience is produced to a 'zero-degree,' i.e. not at all, in the experiencer, not that the experience is produced in a degree below some standard (cf. Kennedy & McNally 2005: 359, ex. 36, and ff. for some further relevant tests).[16]

(150) The movie doesn't frighten Alfonse.

↔ The movie stimulates no fear in Alfonse.

⇏ The movie stimulates fear in Alfonse to a degree below some standard.

Again, relative predicates do not function this way, but absolute lower-closed predicates do.

(151) Alfonse isn't tall.

⇏ Alfonse has no height.

↔ Alfonse has height to a degree below some non-zero standard.

16. This test with negation is also meant to make follow-ups that explicitly assert that some non-zero degree associated with the predicate is present infelicitous. It is difficult to construct plausible sentences like this using psych verbs, but they would be something like: ?*The movie doesn't frighten Alfonse, but he has some fear of it.* To the extent that this is fine English, the result indeed is semantically anomalous. Of course, there will be borderline cases where what experiential state counts as a non-zero degree e.g. of fear is open to question, and so Alfonse might be reported as not being frightened despite experiencing some negligible experience of fear. This is taken to be either a matter of pragmatic imprecision, or of the granularity of the scale being indeterminate (i.e., there is indeterminacy as to what experience of fear is required in order to have experienced a non-zero degree of it).

(152) The door isn't open.

↔ The door is open to no degree.

↮ The door is open to a degree below some non-zero standard.

The deverbal counterparts of these psych verbs preserve these properties: in §2.3.1, this will be accounted for by having them preserve the scalar structure pertaining to experiential kinds.[17]

(153) a. The movie is {slightly / a bit / barely} frightening (to Alfonse).

b. The movie isn't frightening (to Alfonse).

To capture these facts, let there be a minimal degree of each experiential kind k, symbolized by '0_k,' to be read 'the zero-degree of k.'[18]

(154) For every experiential kind k, there is a least element of Δ_k according to $\leq_k$, such that if x is a minimal element according to $\preceq_k$, then $\delta_k(x) = 0_k$.

From what has been said above, it follows that any individual who instantiates the zero-degree of k is an individual that is composed of no portion of k at all, and vice-versa. To see this, recall from §2.1.1 (cf. (126)) that the minimal individuals according to the preordering on qualities $\preceq_k$ are those that are composed of no sub-parts that are a portion of k. But according to (154), precisely these individuals instantiate the zero-degree of k.

With all this said, degree morphology, and the resulting truth-evaluable sentences containing object-experiencer psych verbs, can be composed. First, let the psych verb form an

17. Insofar as e.g. Glanzberg (2007: 14), Sassoon (2010: 139), Fleischer (2013), and perhaps Lasersohn (2008: 308-309) take for granted that experiential predicates are relative, they appear to be mistaken. The standard tests can be run to the same effect also with deverbal psych adjectives, and their favored adjectives, like *tasty*.

18. It is a further question whether the orderings on these sets of degrees have maximal elements. If they did, this would make the positive form of the verb composed with the adjective (e.g. *frighten(s) Alfonse*) a 'closed predicate.' These are meant to occur with proportional modifiers like *partially* (cf. Kennedy & McNally 2005: 352), which doesn't look to be the case: ?*The movie partially frightens Alfonse.* The text makes no commitment as to the existence of maximal degrees of experiential kinds, but it can be assumed for convenience that they are lacking, making the scales that experiential kinds impose solely lower-closed.

ordinary gradable predicate, denoting a measure function of type $\langle e, d\rangle$, by composing with an experiencer.

(155) $[\![\text{frightens Alfonse}]\!]^w = [\![\text{frighten}]\!]^w([\![\text{Alfonse}]\!]^w)$
$= \lambda y_e.\delta_{\text{FEAR}}(\sigma z[stim''(w)(y)(z) \wedge exp''(w)(a)(z)])$

The result is a function from an individual to the degree of fear stimulated by that individual in Alfonse.

Now the degree morphology operates on the measure function to yield a property. Following Kennedy & McNally (2005: 361), let a lower-closed predicate, i.e. a gradable predicate whose ordering on degrees enforces a least element, make use of this least element as the proper standard of comparison that the degree morphology targets.[19]

This will allow for a standard treatment of degree morphology and comparatives. Relevant for present purposes is the positive form of the predicate, where no overt degree morphology appears: let this be treated, again following Kennedy (1997), by composition with the silent morpheme POS. When composing with a lower-closed absolute predicate, POS takes a measure function G and returns a property, viz. that property true of an individual just in case it instantiates a non-zero (i.e. greater-than-zero) degree according to G.

(156) $[\![\text{POS}]\!]^w = \lambda G_{s,ed}.\lambda x_e.0_G <_G G(w)(x)$

In the case of an object-experiencer psych verb and its experiencer argument, 0_G is the zero-degree of the experiential kind k associated with the verb, and $<_G$ is $<_k$. The resulting positive form of *frightens Alfonse* is then as follows.[20]

19. Why this happens with lower-closed predicates will not be addressed here, but is taken to be an independently established observation, at least for most cases. The denotations for degree morphology can be generalized, to make reference to a variable standard of comparison that depends on the lexical properties of the predicate it composes with, where this will default to the least degree on the scale for lower-closed predicates. For convenience, the lexical entry for POS in the text makes reference to the least degree *simpliciter*, since only lower-closed predicates are presently relevant.

20. More explicitly: for any object G of type $\langle s, \langle e, d\rangle\rangle$ that is denoted by some gradable predicate α, let Δ_G be the range of G relative to any world (i.e., $\Delta_G := \{d : \exists w, x[G(w)(x) = d]\}$), and let there be a unique privileged total ordering on every Δ_G, $\leq_G$, according to the scalar properties of α. Then $0_G :=$ the least element of Δ_G according to $\leq_G$. For the case of psych verbs associated with experiential kind k, the result

(157) $[\![\textsc{pos}\ [\text{frightens Alfonse}]]\!]^w = [\![\textsc{pos}]\!]^w(\lambda w_s.[\![\text{frightens Alfonse}]\!]^w)$
$= \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[stim''(w)(x)(z) \wedge exp''(w)(a)(z)])$

It denotes a property true of individuals that stimulate a non-zero degree of fear in Alfonse. And finally, composing with the stimulus subject, to yield *The movie frightens Alfonse*:

(158) $[\![[\text{the movie}]\ [\textsc{pos}\ [\text{frightens Alfonse}]]]\!]^w = [\![\textsc{pos}\ [\text{frightens Alfonse}]]\!]^w([\![\text{the movie}]\!]^w)$
$= 0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(a)(z)])$

The result is true just in case the movie stimulates a non-zero degree of fear in Alfonse. Given what has been said above, this means that there is some portion of fear in Alfonse directed at the movie, which in turn, given §2.1.2, means that Alfonse has had an experiential episode of fear of the movie, in virtue of which he is generically disposed to do so again when directing his experiential faculties toward it. This is the correct result, and it generalizes to other stative object-experiencer psych verbs, swapping out for the relevant experiential kind.

The above machinery yields a welcome result with respect to how stative object-experiencer psych verbs encode presuppositions of direct experience. Sentences of this form presuppose that the experiencer has had an experience stimulated by the stimulus, at least of the sort relevant to the production of the kind associated with the verb. In the case of a movie stimulating fear, as a shorthand the relevant experience may be taken to be having watched the movie.[21]

---

will be as in (164) so long as the privileged ordering is $\leq_k$, and likewise for every other verb, swapping out for the relevant experiential kind. This machinery suffices for the present, but would have to be complicated to accommodate different predicates that enforce distinct orderings on the same set of degrees, like cross-polar antonyms (*tall* / *short*) (in which case there is no unique privileged ordering on said set of degrees, e.g. height degrees).

21. This limiting of experience to having watched the movie is only to aid the exposition: it isn't reflected in the semantics, as it shouldn't be, since one may report that the movie frightens Alfonse if, e.g. he is scared of its title, or its packaging, even having never watched it. The semantics itself only refers to whether the movie has stimulated experience in Alfonse *simpliciter*, and this looks to be correct, given the wide construal these reports can be given of what kind of experience is necessary (in principle, any whatsoever). In cases where interlocutors seem to make restriction to a certain kind of experience part of the presupposed content, it can be presumed that this happens either due to pragmatic imprecision, or due to domain restriction on the reference to an experience, i.e. on 'σ.'

(159) The movie frightens Alfonse.

↪ Alfonse has watched the movie.

This presupposition projects out of negation (160), can be targeted by the usual discourse moves (161), and is locally accommodated in the usual intensional environments (162).

(160) The movie doesn't frighten Alfonse.

↪ Alfonse has watched the movie.

(161) A: The movie frightens Alfonse.

B: Wait a minute – Alfonse hasn't watched the movie.

(162) The movie will frighten Alfonse.

↪ Alfonse will watch the movie.

This presupposition is already encoded in the semantics provided above, by means of how experiential states are defined. (158) contains the term '$\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(a)(z)]$,' which is defined only if there is a unique maximal individual meeting the condition in the brackets over which 'σz' scopes. But this is so only if there is some individual of which the movie is the stimulus, and Alfonse is the experiencer, i.e. only if an experience is located in him and directed at the movie. And per §2.1.2, this happens only if Alfonse has undergone an experiential episode directed at the movie.

It follows that these psych verbs cannot have merely dispositional readings, but enforce actual experiences: in (158), the σ-term is undefined where there is no such experience, hence δ_{FEAR} operates on no value, and the denotation as a whole is undefined, resulting in presupposition failure. That this presupposition projects out of non-external sentential negation is easy enough to see, for simplicity taking *doesn't* to be a predicate negator.

(163) $[\![\text{doesn't}]\!]^w = \lambda P_{et}.\lambda x_e.\neg P(x)$

(164) $[\![[\text{the movie}]\ [\text{doesn't}\ [\textsc{pos}\ [\text{frighten Alfonse}]]]]\!]^w$

$= [\![\text{doesn't}]\!]^w([\![\textsc{pos}\ [\text{frighten Alfonse}]]\!]^w)([\![\text{the movie}]\!]^w)$

$$= \neg[0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(a)(z)])]$$
$$= 0_{\text{FEAR}} =_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(a)(z)])$$

The resulting denotation again presupposes that the movie has stimulated experience in Alfonse, by the definedness conditions on the σ-term, and where defined it is true just in case the degree of fear stimulated in Alfonse by the movie is the zero-degree. This is so just in case said experience contains no portion of fear, i.e. just in case Alfonse did not have an experiential episode of fear directed at the movie, in virtue of which he is disposed to have such episodes under the relevant conditions. This is right: the sentence doesn't just report that Alfonse lacks fear of the movie, but further reports that it fails to stimulate fear in him.

The language thus draws a three-way distinction using psych verbs with respect to the experience stimulated: there is (i) stimulation of experience of a certain kind (as when watching the movie instills fear), (ii) stimulation of an experience that is not of a certain kind (as when watching the movie fails to instill fear), and (iii) a lack of stimulation of experience (as when one hasn't watched the movie to begin with). This follows from the above treatment of experiential states, combined with the treatment of stative object-experiencer psych verbs as lower-closed predicates, which in particular allows for the distinction between the stimulus of a zero-degree of a kind, and a lack of stimulus altogether.

2.3 Deverbal psych adjectives

With the semantics for object-experiencer psych verbs in place, a denotation for *-ing*, and the resulting experiential predicates, can be provided. Echoing the bare denotation given for *tasty* in §1.3 (cf. (92)), the denotation for a deverbal psych predicate like *frightening* will be an instantiation of the following schema.

(165) $[\![\text{frightening}]\!]^w = \lambda x_e.frighening'(w)(x)$

Here '$frightening'$' is read 'disquotationally,' such that $frightening'(w)(x)$ iff x is frightening at w, and so on for every such adjective. The following section will flesh out this

schematic denotation. The denotation for *tasty*, being sensitive not only to the experiential kinds examined in this chapter, but also to the sensory modality of taste, will have to wait until Chapter 3, where sensory verbs and related constructions are discussed.

§2.3.1 provides a semantics for *-ing* itself, and discusses how it interacts with genericity, such that deverbal psych predicates are true of individuals that are disposed, under the appropriate circumstances, to produce experience of a certain kind. §2.3.2 discusses the treatment of exocentricity in experiential predicates, where said predicates are read as anchored to a contextually relevant experiencer, and shows how the present semantics offers a promising way of incorporating it. Finally, §2.3.3 discusses the treatment of overt experiencers in the form of *to*-headed PP phrases, and summarizes the results of the present chapter for the semantics of experiential predicates generally.

2.3.1 Adjectival -ing *and genericity*

An experiential predicate is completed by composing the suffix *-ing* with an object-experiencer psych verb. Roughly, the resulting predicate is one true of an individual just in case it is disposed to be the stimulus of a non-zero degree of the experiential kind associated with the verb in question. Thus, *frightening* is true of an individual just in case it is disposed to be the stimulus of fear, i.e. just in case it is disposed to frighten.

There are two caveats to keep in mind in providing a denotation for *-ing* that gives this result. First, as concluded in Chapter 1, the relevant sense of being disposed to be a stimulus includes no intrinsic reference to an experiencer: e.g., to be frightening is not to be disposed to produce fear in any particular experiencer, or even just some experiencer or other, but *simpliciter*: what this means will be elucidated in part immediately below, and in part in Chapter 5. Second, the actualization of a disposition requires that certain conditions be met, and since there is no guarantee that they ever are, there is no guarantee that the disposition is ever actualized (as reflected in §1.1.2: experiential predicates don't inherently encode any presuppositions of direct experience).

To get a handle on both these points, note that the use of *-ing* to compose experiential predicates relates to a broader use of this form in English to produce dispositional adjectives. A large class of verbs take a homophonous suffix, resulting in a predicate that is true of an individual roughly just in case that individual is disposed to perform the role indicated by the verb's subject in an eventuality. Thus, with *wash* and *sing*:

(166) a. washing machine

b. singing bird

A washing machine is a machine that is disposed to be the causer of washing events; a singing bird that is disposed to be the agent of singing events. Where the verb is transitive, as with *wash*, and thus encodes reference to arguments besides the external (here, the washer), the result of composing with *-ing* makes no reference to those other arguments (here, the washed thing): a washing machine is disposed to wash *simpliciter*, and the adjective itself doesn't specify what is washed, just as *frightening* makes no reference to what is frightened.[22] This is so even if the washing events in which the machine takes part must involve some washed thing or other.

These adjectives are also applicable where the disposition isn't actualized: thus, a machine can already be a washing machine straight out of the factory, even though it hasn't yet washed anything, and a bird can be a singing bird even when newborn. The same is true of the psych predicate, substituting the role of stimulus for agent or causer: all that is required is that where such-and-such conditions are met, the predicated individual is the stimulus of fear.

What follows won't address whether this general occurrence of *-ing* can be equated with the one attaching regularly to object-experiencer psych verbs, but will just apply the above lessons to the specific case of deverbal psych predicates, and offer a denotation specific

22. Of course, culturally it might be apparent that a washing machine typically washes clothes, or some other stereotypical thing. The point is just that this isn't lexically encoded by the adjectival modifier – a machine that washes dogs can still be called a washing machine.

to these.[23] The role of *-ing* when composing with object-experiencer psych verbs is to 'desaturate' the experiencer, and the relevant dispositional condition then somehow needs to be enforced. This can be done by having *-ing* operate on a transitive measure function, and return an intransitive measure function that preserves the relevant experiential kind and scale structure, while referencing only the stimulus.

(167) $[\![\text{-ing}]\!]^w = \lambda R_{s,\langle e,ed\rangle}.\lambda x_e.\delta_{\kappa''(R)}(\sigma y[stim''(w)(x)(y)])$

κ'' is the transitive counterpart to κ': it applies to the intension of a transitive measure function, and 'fishes out' the relevant experiential kind. It can be defined as follows.

(168) If for all w, x, y such that $R(w)(x)(y)$ is defined, $R(w)(x)(y) \in \Delta_k$, then $\kappa''(R) = k$; else $\kappa''(R)$ is undefined.

Thus where R is the intension of *frighten*, $\kappa''(R) = \text{FEAR}$, and where R is the intension of *confuse*, $\kappa''(R) = \text{CONFUSION}$, and so on.

The result of composing *-ing* with *frighten* is then as follows, by intensional function application, repeating the denotation for *frighten* from (146) as (169).

(169) $[\![\text{frighten}]\!]^w = \lambda x_e.\lambda y_e.\delta_{\text{FEAR}}(\sigma z[exp''(w)(z)(x) \wedge stim''(w)(z)(y)])$

(170) $[\![\text{frightening}]\!]^w = [\![\text{-ing}]\!]^w(\lambda w_s.[\![\text{frighten}]\!]^w)$
$= \lambda x_e.\delta_{\text{FEAR}}(\sigma y[stim''(w)(x)(y)])$

The result is that *frightening* denotes a measure function mapping an individual to the degree of the maximal portion of fright of which that individual is the stimulus. The scalarity

23. There are asymmetries problematic for such an equation. When used with most verbs, dispositional *-ing* is disallowed in predicate position: **The machine is washing* and *The bird is singing* have only progressive readings (meaning that *wash* remains verbal and retains its transitivity). Cf. *The movie is frightening.* This may be related to the non-stativity of the verbs: by contrast, *The parents are understanding* is fine where *understanding* encodes a disposition to understand in the (stative) sense of empathizing. If this were true, then the behavior of stative (cf. §2.2.1) object-experiencer psych verbs would be in line with the generalization (though it is imperfect: cf. *The parents are forgiving*, where *forgive* is eventive, yet the dispositional reading nonetheless survives in predicative position). Most subject-experiencer psych verbs are at least highly degraded with *-ing*, as with **the enjoying customers*, suggesting that *-ing* is sensitive to which role in the eventuality the disposition pertains to (where stimuli are allowed, and experiencers are not). This has to be taken into account in providing a semantics for subject-experiencer psych verbs.

of the psych verb is preserved, as is the lower-closed scale structure, allowing for comparative and degree morphology, and triggering an absolute reading of the positive form. POS then converts the adjective into a property-denoting expression as usual, repeating the denotation for POS from (156) as (171).

(171) $[\![\text{POS}]\!]^w = \lambda G_{s,ed}.\lambda x_e.0_G <_G G(x)$

(172) $[\![\text{POS frightening}]\!]^w = [\![\text{POS}]\!]^w(\lambda w_s.[\![\text{frightening}]\!]^w)$
$= \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y[stim''(w)(x)(y)])$

And so *frightening* in its positive form is true of an individual just in case the degree of the maximal portion of fear that it stimulates is non-zero.

As it stands, this denotation is incomplete: while it makes no explicit reference to a particular experiencer, as desired, it does ensure that there is some experiencer or other when its truth conditions are met, since where there is a stimulus of an experience, there is also an experiencer (cf. §2.1.2). The psych predicate is therefore true of an individual just in case it stimulates some non-zero degree of fear or other, in any experiencer: and so *frightening* turns out to mean roughly *frightening to something*.

Similar proposals for arguably experiential predicates, including *enjoyable* (Lakoff 1970: 127) and *fun* (Chierchia 1984: 405, ex. 3) have been made in passing, but this semantics is both too strong and too weak.[24] It is too strong, because as shown in §1.1.2, experiential predicates do not enforce direct experience effects on their own, and so generally do not require for their application that any experiencer be stimulated by the predicated individual. Too weak, because as Lasersohn (2005: 653) notes, it is possible without semantic incompetence to evaluate such predications as false even where some experiencer or other has the

24. Chierchia's treatment is more complicated than this: the existential quantification over experiencers that *fun* introduces when its internal argument is desaturated interacts with a generic operator, contributed by tense/aspect, such that generic predications of *fun* in the present tense roughly state that in general, there is an experiencer for which the predicated individual is fun, where the generic operator must do the work of specifying the relevant conditions under which there is such an experiencer. Episodic predications don't have this generic reading, leaving the plain existential quantification intact. As will be seen in §2.3.2, this approach is ultimately quite resonant with what will be offered here: cf. Chierchia (1984: 404-406).

relevant experience: *The movie isn't (actually) frightening, even though it's frightening to Alfonse* is not contradictory, while *#The movie isn't (actually) frightening to anyone, even though it's frightening to Alfonse* is.

What's missing is the dispositional component of the predicate mentioned above: experiential predications typically have a generic flavor, such that they hold of the predicated individual just in case that individual becomes a stimulus of the experiential kind 'in the appropriate circumstances.' And so in order for experiential dispositions to be actualized, what Chierchia (1995: 195-196) calls their *felicity conditions* have to be met. For experiential predicates, these conditions involve a number of uncontroversial but only vaguely defined prerequisites: the individual in question must be positioned to make 'experiential contact' with an experiencer, the experiencer must have working sensory or imaginative faculties, the individual and the experiencer must not be defective in any way that relevantly obscures the properties of the potential stimulus (e.g., the movie must be played at the right volume, and not obscured from view, and the potential experiencer must have the relevant knowledge of what's being portrayed), and so on. A typical predication using an experiential predicate is true where the establishment of all these felicity conditions guarantees that the individual is a stimulus of the relevant experiential kind.

The need to interpret experiential predicates generically has been stressed by Snyder (2013) and Pearson (2013), and both independently develop accounts that introduce genericity by means of an operator that scopes over the predicate. However, they both assume that this genericity involves generic quantification over experiencers, and so roughly interpret experiential predicates as being true of an individual just in case that individual is the stimulus of a certain kind of experience 'for experiencers generally.'

They begin with the observation that experiential predicates are individual-level, as diagnosed by the classic tests in Milsark (1977) and Carlson (1977): for instance, like other individual-level predicates, experiential predicates are degraded in the coda of existential *there*-constructions (173), and in predicate position of a small clause beneath a verb of per-

ception (174); they also produce generic, rather than existential, readings when predicated of kinds (175). These observations extend to the deverbal psych predicates discussed here.

(173) a. ?There were people annoying.

b. ?There were people tall.

c. There were people sick.

(174) a. ?I've seen Alfonse frustrating.

b. ?I've seen Alfonse smart.

c. I've seen Alfonse drunk.

(175) a. Hitchcock movies are frightening.

b. Hitchcock movies are well-directed.

c. Hitchcock movies are in the den.

Next, following Chierchia (1995), they propose that individual-level predicates must be licensed by an operator GEN that scopes over them.[25] However, both Snyder and Pearson (cf. Pearson 2013: 122, ex. 45) assume a contextualist denotation for the adjective, which takes an internal experiencer argument. The result is that where this argument is implicit, it enters the composition as a variable which can then be bound by GEN, and the denotation of the resulting predicate is true of an individual just in case it produces the relevant experience in experiencers generally (where the felicity conditions of the generic are met). And so the structure of *The movie is frightening* looks something like the following (reconstructing in part from Pearson 2013: 127, ex. 65b, ignoring the movement of the subject).

25. On Chierchia's analysis, GEN is actually licensed by a null habitual aspectual morpheme Hab, which requires a generic operator in its specifier due to spec-head agreement: cf. Cherchia (1995: 197-198, esp. ex. 53). He ultimately decides that individual-level predicates, including individual-level adjectives like those examined here, carry a habitual morpheme in the lexicon, and so always require GEN to appear for local licensing (cf. *ibid.* §3.4.2). This creates the result that experiential predicates, insofar as they are individual-level, cannot appear absent GEN – this appears to be problematic, as §2.3.2 will address. In the above this is temporarily ignored, and it is taken for granted that GEN appears only when needed, regardless of how this happens.

(176)

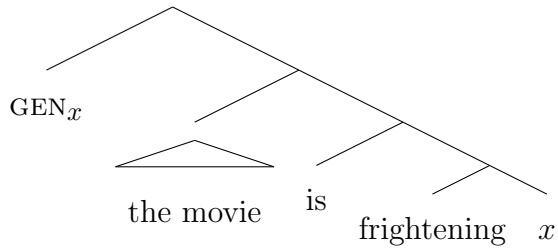

Adopting a standard contextualist denotation for *frightening* as in (177), where $frightening''(w)(x)(y)$ is true just in case y frightens x at w, and '$\mathfrak{Ↄ}$' is the generic quantifier, the resulting truth conditions are roughly as in (178).

(177) $[\![\text{frightening}]\!]^w = \lambda x_e.\lambda y_e.frightening''(w)(x)(y)$

(178) $\mathfrak{Ↄ}[x, w][frightening''(w)(x)(\iota y[movie'(w)(y)])]$

(178) is to be read as 'generally, given a world w and experiencer x, the movie at w frightens x at w.' The quantification over both worlds and experiencers can be read as enforcing the relevant felicity conditions: thus, the denotation is concerned with all worlds and individuals within some relevantly constrained set. This truth condition is then meant to be that expressed by a non-exocentric predication using the experiential predicate.

The issue with quantifying over experiencers in this way is that it does not describe the truth conditions of the predication correctly. There is no plausible way to construe a generic quantification amounting to 'the movie is frightening to experiencers generally' that matches how speakers competently evaluate the non-exocentric denotation of *frightening*. Typically, as noted in Chapter 1 and stressed by Lasersohn (2005: 670) and MacFarlane (2014: 4), speakers tend to evaluate an experiential predicate as true of an individual just in case that individual is disposed to produce the relevant experience in themselves. There are two major exceptions to this tendency: exocentric readings, to be addressed in §2.3.2, and evaluation when taking one's own perceptions to be inaccurate, to be addressed in §4.2.1. But these are not what the proposed reading in (178) is intended to model, and so speakers' behavior in typically evaluating these predicates relative to themselves, and not to experiencers generally, is mysterious.

Snyder (2013: 297), in addressing an objection in this vein from Lasersohn (2005: 653-654), responds that the worry comes from an insufficiently nuanced understanding of genericity. If one adopts a modal notion of generic quantification in the vein Krifka et al. (1995: 52, ex. 86), the truth conditions are meant to come out correctly. What Snyder ends up with is as follows (cf. *ibid.* 2013: 304, ex. 59), where B_w is a modal base determined by w, and $\leq_w$ is an ordering on B_w fixed by an ordering source, where $w' \leq_w w''$ just in case w' is 'at least as normal' according to the generic ordering on B_w as w''.

(179) $[\![\text{GEN}_x \text{ the movie is frightening } x]\!]^w =$
$\forall x, w' \in B_w[\exists w'' \in B_w[w'' \leq_w w' \wedge \forall w''' \in B_w[w''' \leq_w w'' \rightarrow$
$frightening''(w''')(x)(\iota y[movie'(w)(y)])]]]$

The rough idea is as follows. Supposing there is a modal base B_w, and an ordering on B_w that ranks worlds according to how 'normal' they are, to say the movie is frightening is to say that for any experiencer x, the movie frightens x at all worlds within B_w that approach ideal normality. It therefore doesn't follow that every experiencer, or any particular amount of experiencers, are actually frightened by the movie, because ideal conditions may not hold; but experiencers generally do have the relevant experiences insofar as those conditions do hold.

It's isn't clear how this gloss on generic truth conditions helps – conditions on ideal normality across experiencers are also irrelevant to the usual evaluation of the experiential predicate, if this normality is anything like what's usually tracked by generics. Where Alfonse is frightened by a movie, and admits that he is an abnormal experiencer with strange tastes, he is still free to evaluate whether the movie is frightening based solely on his own experiences.[26] Conditions enforced by the generic on experiencers are irrelevant to truth conditions, *except* insofar as those conditions are defined in terms of oneself as experiencer,

26. Of course, he might in some cases be free not to as well, if his abnormality causes him to think of himself as an inaccurate perceiver (cf. §4.2.1). The point, however, is that there is not only nothing *wrong* with his evaluating the predicate based on his own experiences, it represents the usual, competent use of the predicate. A speaker is not semantically incompetent for having an odd opinion.

i.e. unless speakers for some reason conventionally take themselves to be ideally normal experiencers in ideally normal circumstances.

This latter option, which would make ideal normality and one's own experiential reactions coincide, perhaps for some pragmatic reason, is implementable in theory, and Snyder (*ibid.* §6.4) in fact briefly entertains relativizing modal bases and ordering sources to experiencers in a way that might suggest just such a treatment. But in that case, the generic quantification over experiencers does no work: *either* the generic quantification can fail to reflect autocentric evaluation, in which case the truth conditions are wrong, *or* it can reflect it, in which case the result of generic quantification is identical to a relativist treatment, on which the predicate is evaluated purely with respect to oneself, and generic quantification doesn't perform its ordinary function due to some sort of pragmatic interference, so there is no clear motivation for positing the quantification to begin with.

This issue becomes clearer in Pearson's (2013) treatment. Impressed by the data suggesting that speakers evaluate experiential predicates with respect to their own experiences, she posits that the covert variable over which GEN quantifies ranges only over those individuals with which the speaker[27] 'identifies.' The resulting denotation for *The movie is frightening* then turns out to be as follows (cf. Pearson *ibid.* 127, ex. 65), where $I''(w)(x)(y)$ is true iff x 'identifies with' y at w,[28] c is a context of utterance, s_c the speaker in c, R is an accessibility relation on worlds, and $C(x)(y)(w)$ is true iff w is a world at which x and y are 'contextually relevant.'[29]

(180) $[\![\text{GEN}_x \text{ the movie is frightening } x]\!]^{c,w} =$

27. This apparently won't work – presumably, what's wanted is for the variable to range over the individuals with whom the assessor or evaluator of the predicate identifies, not the speaker. Otherwise, all the old contextualist problems addressed in Chapter 1 resurface, as the predicate is anchored to the speaker for its content.

28. Pearson does not relativize the identify-relation to worlds, but it seems that it must be so relativized; it isn't clear what it means for one individual to identify with another *simpliciter*, and presumably identification conditions change counterfactually (as they must if they're to track the speaker's experiences).

29. This notion of contextual relevance looks to be an implementation of felicity conditions by another name.

$$\forall x, w'[[wRw' \wedge C(\iota y[movie'(w)(y)])(x)(w') \wedge I''(w')(s_c)(x)] \rightarrow$$
$$frightening''(w')(x)(\iota y[movie'(w')(y)])]$$

Thus, *the movie is frightening* is true just in case at all accessible worlds subject to various contextual conditions set by C, the movie frightens all individuals with which the speaker identifies. The sensitivity of generic quantification to the speaker's experiences is thus 'baked in:' the only experiencers that 'count' are those that, in addition to meeting the normal felicity conditions on generic quantification, match the speaker in some relevant way.

But now the superfluity of the quantification over experiencers is obvious: for if the truth conditions are correct, the 'identify-with' relation must track the experiential reactions of the speaker insofar as they relate to fear engendered by the movie. But in this case, the identification relation can simply be done away with: it is the same just to say that the speaker is frightened by the movie where the felicity conditions enforced by C are met.[30]

With the above said, it's important to understand that the present treatment will implement genericity in a way that does *not* generically quantify *over experiencers*. This is reflected in the denotation already given in (172): there is no variable representing an experiencer for a generic operator to bind. In adding genericity to this denotation, the result will be a predicate simply true of an individual just in case it is the stimulus of a certain kind of experience whenever the relevant felicity conditions are met.

The revised denotation for *-ing* that implements genericity in this experiencer-independent way then works as follows. *-ing* composes with a transitive measure function R, denoted by a psych verb, and invokes an accessibility relation $\rho_{\kappa''(R)}$, which can be thought of as an accessibility relation on worlds as they pertain to the felicity conditions for the production of $\kappa''(R)$. Thus, where R is the intension of *frighten*, $\rho_{\kappa''(R)}$ is ρ_{FEAR}, and $w\rho_{\text{FEAR}}w'$ just in case w' is accessible from w according to the fear-production accessibility relation. This relation is in turn constrained as follows.

30. For Pearson, this is actually not quite true: this is because she allows that the speaker may sometimes be excluded from the relevance conditions, which leads to exocentric readings of experiential predicates. See Pearson (2013: 142-143).

(181) $w\rho_k w'$ iff:

a. w' preserves all intrinsic properties holding at w of potential stimuli relevant for their production of k;

b. w''s conditions, given these properties, are ideally normal for the production of k by potential stimuli.

If one likes, one can think of these two clauses as reflecting the operation of a modal base and an ordering source respectively, which work to determine a set of accessible worlds w' from each world of evaluation w. (181a) sets the modal base, as containing all worlds w' that are compatible with the intrinsic properties of stimuli at w relevant to the production of a kind. This would mean, in the case of movies and how frightening they are, that in all such worlds the movies themselves do not change their content across these worlds – they are, in whatever relevant respect, the 'same movies.'[31]

(181b) then reflects an ordering source, such that the worlds in this modal base that are quantified over are those that are ideally normal, given the intrinsic properties of the stimuli, for the production of the experiential kind. With respect to the movies again, this will mean that given that they are the 'same' movies, the appropriate conditions for viewing those movies to produce fear have been actualized, as described above with respect to felicity conditions. The worlds that are quantified over in expressing the generic flavor of *-ing* are thus those worlds in which the felicity conditions for experiencing an experiential kind from stimuli, such as they actually are, are met.[32]

31. The restrictions of the modal base and ordering source are necessary to enforce a distinction between two types of counterfactual reasoning: first, that regarding what experiences an individual is disposed to produce given that only surrounding conditions are different (this is the type of reasoning tracked by the application of the psych predicate relative to the world of evaluation), and second, that regarding what experiences an individual would be disposed to produce if it changed in the relevant intrinsic way, like a movie changing its directing, or a food changing its flavor (this is the type of reasoning tracked by counterfactual assessments of experiential predicates, like *If the movie were frightening..*).

32. More formally, and using the language of Kratzer (1986): Let $f_{w,k}$, $g_{w,k}$ be conversational backgrounds, i.e. sets of propositions (objects of type $\langle s,t\rangle$) determined at w. $f_{w,k}$ is the set of propositions true at w pertaining the the properties true of potential stimuli relevant to their production of k. Then where $f_{w,k,S}$ is the set of characteristic sets of elements of $f_{w,k}$, $B_{w,k} = \bigcap f_{w,k,S}$ is the modal base relevant to k at w. $g_{w,k}$ is the set of propositions true at w pertaining to the conditions relevant for the production of k by stimuli at

One final stipulation: it is taken for granted that there is some one degree of the relevant experiential kind that is produced under ideal conditions. To speak in terms of the ordering source, those worlds ranked highest with respect to the felicity conditions for the production of the kind converge on how much of that kind is produced. With that said, *-ing* composes with the psych verb to return a measure function, which takes an individual and returns that very degree d produced under said ideal conditions.

(182) $[\![\text{-ing}]\!]^w = \lambda R_{s,\langle e,ed\rangle}.\lambda x_e.\iota d[\forall w' : w\rho_R w'[d = \delta_{\kappa''(R)}(\sigma y[stim''(w')(x)(y)])]]$

And the denotation for *frightening*, and its positive form, is as follows.

(183) $[\![\text{frightening}]\!]^w = [\![\text{-ing}]\!]^w(\lambda w_s.[\![\text{frighten}]\!]^w)$
$= \lambda x_e.\iota d[\forall w' : w\rho_{\text{FEAR}} w'[d = \delta_{\text{FEAR}}(\sigma y[stim''(w')(x)(y)])]]$

(184) $[\![\text{POS frightening}]\!]^w = [\![\text{POS}]\!]^w(\lambda w_s.[\![\text{frightening}]\!]^w)$
$= \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \iota d[\forall w' : w\rho_R w'[d = \delta_{\text{FEAR}}(\sigma y[stim''(w')(x)(y)])]]$

The positive form of *frightening* is therefore true of an individual just in case in all worlds relevantly similar to the world of evaluation in which the felicity conditions relevant to causing fear are met, that individual is the stimulus of a non-zero degree of fear (which is by definition fear in some experiencer or other). In short: a frightening individual is one that is disposed to produce fear *simpliciter* – and so on for every deverbal psych predicate, swapping out for the relevant experiential kind.

This treatment anchors the genericity involved in the predicate intrinsically to the adjective, so that all uses of experiential predicates encode this genericity. §2.3.2 will give some reasons to believe this is a simplification, and suggest that genericity is in fact contributed by a distinct element, as Snyder and Pearson maintain.

Since the denotation in (184) appeals to a quantification over worlds in which some felicity conditions are met, and there is no guarantee that those conditions are met at the

---

w given $f_{w,k}$. $\preceq_{g_{w,k}}$ is that partial order on $B_{w,k}$ determined by $g_{w,k}$ such that for any w', w'': $w'' \preceq_{g_{w,k}} w'$ iff $\forall p \in g_{w,k}[p(w') \rightarrow p(w'')]$. Adopting the limit assumption, $\text{BEST}_{B_{w,k}} := \{w' : \neg\exists w''[w'' \prec_{g_{w,k}} w']\}$. Then $w\rho_k w'$ iff $w' \in \text{BEST}_{B_{w,k}}$.

world of evaluation, it follows that the truth of the predicate does not hinge on the predicated individual actually stimulating the relevant experience in that world. Nor, even when that experience is stimulated, does a predication using *frightening* merely report this actual stimulation. The reading is both weaker and stronger than this: weaker, because the predicate can hold of an individual even where it has never stimulated any experience at all, and stronger, because it holds of an individual not just when it happens to produce the relevant experience, but where it counterfactually would produce that experience *whenever* the felicity conditions are met.

This means that, as noted in §1.1.2, positive forms of experiential predicates do not give rise to presuppositions of direct experience, unlike psych verbs. To illustrate briefly, the interpretation of (185a) is as in (186), taking an interpretation of *wonder* as in (185b), where '?' is a polar question operator.

(185) a. Alfonse wonders whether the movie is frightening.

b. $[\![\text{wonder}]\!]^w = \lambda p_{st}.\lambda x_e.wonder''(w)(?p)(x)$

(186) $[\![\text{Alfonse [wonders [whether [the movie] [is [POS [frightening]]]]]}]\!]^w =$
$wonder''(w)(?\lambda w'_s.0_{\text{FEAR}} <_{\text{FEAR}} \iota d[\forall w'' : w' \rho_{\text{FEAR}} w''$
$[d = \delta_{\text{FEAR}}(\sigma y[stim''(w'')(x)(y)])]])(a)$

Thus, what Alfonse wonders about is the truth of the proposition that the movie is disposed, under the relevant felicity conditions, to be the stimulus of a non-zero degree of fear. He is not reported as wondering whether the movie is frightening to any experiencer, and so he can coherently wonder whether the movie is frightening while committed to thinking either that no one has ever seen the movie, or that someone has. The same will hold of all the relevant intensional contexts that give rise to no direct experience presuppositions. In fact, even the unembedded *The movie is frightening* gives rise to no such presupposition, although it triggers a speaker-oriented acquaintance inference, for reasons to be explored in §4.1.

What are the felicity conditions relevant to the production of an experiential kind? Since they are only loosely defined, the linguistic conventions don't settle exactly what they are. Speakers thus have diverging opinions about them, which result in diverging evaluations of the generic predicate, even where there is agreement on the state of the world. And so even where Alfonse and Bethany are fully informed about the same horror movie, and take it to produce fear only at night when the viewer is alone, Alfonse might take the movie not to be frightening, since he takes the felicity conditions for being frightening to be met in other circumstances as well, where no fear is produced (e.g., when not alone). But Bethany might then take it to be frightening, since she takes the relevant felicity conditions for the movie producing fear to involve the viewer being alone at night.

Most of these felicity conditions pertain to any sort of genericity, and so are orthogonal to the present concern with experiential semantics, but there is one crucial exception: these conditions are also sensitive to which *experiencers* are stimulated by the individual in the relevant conditions, and so part of the felicity conditions being met will be that the stimulus comes into contact with the 'right' sort of experiencer. Chapter 5 deals with this issue, and shows that the interaction of experiential predicates with direct evidentiality causes speakers to treat their own experiences as those tracking the felicity conditions of the predicate, resulting in autocentric evaluation.

2.3.2 Exocentricity, aspect, and definiteness

Exocentric readings of experiential predicates are those that are contextually anchored to some relevant experiencer, despite not occurring with any overt marker of this experiencer.[33]

33. Some confusion has arisen surrounding the terms 'autocentric' and 'exocentric' since their introduction to the relativist literature in Lasersohn (2005). A number of researchers have taken autocentric uses of experiential predicates to be anchored to the speaker for evaluation of truth, and this in turn has led to some misguided criticisms of the pragmatics of relativism (e.g. Wolf 2016: §2.2.1). Autocentricity has nothing to do with the speaker: it refers to the evaluation of predicates with oneself as the relevant experiencer, whether one is the speaker or not. Exocentricity is then contra-defined by Lasersohn as evaluation of the predicate relative to someone other than oneself (or at a time other than the assessment time: cf. Lasersohn 2017: 140-141), and has nothing to do with the speaker either. In the present text, autocentric evaluation is taken as the default way to evaluate 'bare' uses of experiential predicates, while exocentric uses of experiential

They do not create the sorts of evaluative, merely dispositional readings accounted for in §2.3.1. Rather, they simply report that the experiencer in question has the relevant experience, directed toward the relevant stimulus.

Exocentric readings occur with deverbal psych predicates, in all the environments expected of experiential predicates generally. Sometimes, strong contextual cues are needed to make an exocentric reading felicitous, as with (187a), which might be anchored to Alfonse's experiences if uttered while looking at him watching a horror movie, clearly frightened. In other contexts, such as (187b) and (187c), an exocentric reading is more easily accessible, and may even be the default.

(187) a. The movie is frightening.
↪ Alfonse has experienced fear of the movie.

b. Alfonse went to a funeral service. It was distressing.
↪ Alfonse has experienced distress at the funeral service.

c. Alfonse's boss gave him a boring speech.
↪ Alfonse has experienced boredom at his boss' speech.

Exocentric readings also give rise to the presuppositions of direct experience indicated by the hooked arrows in (187); these are the same that arise where an experiencer is overtly marked, as with the objects of object-experiencer psych verbs (cf. §2.2.2).

These presuppositions also display the normal projective and accommodative behavior. Thus for instance, (188) on its exocentric reading, where the speaker wonders only about how the movie affected Alfonse, locally accommodates the presupposition that Alfonse experienced the movie in a way appropriate to engendering fear; what the speaker wonders must include whether Alfonse actually saw the movie and was frightened by it, and not just whether the movie would have frightened him, had he seen it.

(188) Alfonse went to see a movie. I wonder if it was frightening.

---

predicates are those that anchor to a contextually relevant experiencer for evaluation, whether this is the speaker or not.

The same is true for any of the various intensional contexts reviewed in §1.1.2: the behavior of the experiential presupposition is as expected where reference to an experiencer is made by the embedded content. Exocentric readings therefore behave 'as if' an overt experiencer were present: they are not evaluated autocentrically, but rather relative to the experiences of the contextually relevant individual; they do not give rise to acquaintance inferences;[34] and they report an actualized experiential state, rather than reporting the experience a stimulus would engender under the right felicity conditions.

The naïve contextualism criticized in §1.1 can be summarized as the position that all occurrences of experiential predicates with no overt experiencer are exocentric. Chapter 1 showed that this is not correct, and that a solely exocentric account of experiential predicates lacks the expressive power to capture all of their readings. Among those views, like the present, that allow for non-exocentric readings, the systematic relation between readings in which exocentricity is present and absent needs to be explicated.

Here there is a fundamental choice point as to how exocentricity should be treated. The above suggests that exocentricity is a matter of semantic content, and therefore that the content of experiential predicates differs between exocentric and non-exocentric readings: this is the view espoused by Stephenson (2007a: §7.2) and MacFarlane (2014: §7.2.6), both of whom suggest roughly that an exocentric occurrence e.g. of *frightening* is equivalent in content to *frightening to* x, for some contextually salient x.

But Lasersohn (2005: §6.1; 2017: ch. 7) offers an alternative view, according to which exocentricity is not a matter of semantic content, but rather pertains to the pragmatics of truth assessment. According to Lasersohn, the semantic content of a sentence like (187a) is invariant across exocentric and non-exocentric readings: both state that the movie is frightening *simpliciter*, which according to a relativist semantics (cf. §1.2) requires that

34. It is of course possible that an exocentric reading be anchored to the speaker, where the speaker's experiences are for some reason the topic of discussion. In such a case, a presupposition of direct experience anchored to the speaker, similar to that seen in *frightening to me*, arises, rather than an acquaintance inference.

an experiencer be supplied during truth assessment in order for the content in question to yield a truth value. It is then a pragmatic matter which experiencer is provided; and while autocentric evaluation, for which assessors take themselves to be the relevant experiencers, is a robust default, under certain circumstances another contextually relevant experiencer might be supplied as well, and this yields the exocentric interpretation.

Thus the predicate *frightening* has the denotation in (189a) for all its occurrences, where the value for the experiencer intensional parameter x is to be supplied during truth assessment, and the extension of (187a) is as in (189b). A supplementary pragmatic norm as in (190) then allows for the value of the experiencer to shift depending on the pragmatic purposes of assessment (and *mutatis mutandis* for falsity).

(189) a. $[\![\text{frightening}]\!]^{w,x}$
$= \lambda y_e.\delta_{\text{FEAR}}(\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(x)(z)])$

b. $[\![[\text{the movie}]\ [\text{is}\ [\text{POS frightening}]]]\!]^{w,x}$
$= 0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[stim''(w)(\iota y[movie'(w)(y)])(z) \wedge exp''(w)(x)(z)])$

(190) a. ϕ is properly assessed as true in c' iff $\phi(w_{c'})(\epsilon_{c'})$, where $\epsilon_{c'}$ is the experiencer relevant for truth assessment in c'
(and is improperly assessed as true otherwise).

b. In the absence of explicit contextual cues to the contrary, $\epsilon_{c'} = a_{c'}$.

The result is that (187a) expresses a constant relativist-proposition, which will be assessed properly as true or false depending on who the contextually relevant experiencer is: where it is Alfonse's experiences under discussion, the truth or falsity of the expressed proposition will track whether or not the movie frightens Alfonse.

This pragmatic treatment of exocentricity ought to be rejected, for three reasons: the first is a theoretical concern, and the other two are empirical. The first problem has to do with the fact that exocentricity survives embedding in intensional contexts, including attitude reports, as shown by (188). A belief report such as (191) can also be read exocentrically,

where Alfonse is reported as interested only in Bethany's reaction to the movie, and so what he thinks is roughly that the movie was frightening *to Bethany*.

(191) Alfonse thinks that the movie was frightening.

The issue is that whether (191) is true does not depend on the value supplied by the experiencer parameter: this is because *think*, being an intensional verb, operates on the content of its embedded clause. The sentence is thus true when Alfonse believes that content, and false when he doesn't believe it, and the experiencer parameter is idle in the evaluation of the truth of (191). This is as it should be, since the truth of the belief ascription is not experiencer-sensitive, but depends only on the non-experiential matter of what Alfonse believes.

But if the embedded clause *the movie was frightening* has the same content regardless of whether its use is exocentric or not, then what Alfonse is reported as believing in both the exocentric and non-exocentric case must be the same, and so the truth value of the belief ascription cannot change across the two cases. In other words, the pragmatic treatment of exocentricity, so long as *think* takes a propositional argument as normal, predicts that exocentricity ought not to appear in attitude reports at all, contrary to fact.

Aware of this problem, Lasersohn (2005: §6.2) alters the semantics of attitude verbs, effectively to include an extra argument for an experiencer. Thus, an individual does not believe a proposition *simpliciter*, but only a proposition relative to an experiencer. He later updates the denotation to simply existentially quantify over experiencers, so that *think* would be as follows, where '$believe'''$' denotes a triadic relation relative to worlds, such that $believe'''(w)(p)(x)(y)$ is true iff y believes p at w relative to experiencer x (adapted from Lasersohn 2017: 158, ex. 188).[35]

(192) $[\![\text{think}]\!]^{w,x} = \lambda p_{s,et}.\lambda y_e.\exists z[believe'''(w)(p)(z)(y)]$

35. This denotation would appear to make *Alfonse thinks that the movie is frightening* mean that Alfonse thinks the movie is frightening to some experiencer or other. Presumably, some domain restriction on the existential quantification is meant to narrow down the options so that the report is anchored to the contextually relevant experiencer.

The idea is therefore that when (191) has a reading relevant to Bethany's experiences, this is encoded in the content of the belief report: Alfonse believes the proposition denoted by *the movie was frightening*, relative to Bethany as experiencer, even if he doesn't believe it relative to some other experiencer, or *simpliciter* as in the default *de se* case.

As shown in §1.2.1, however, this cannot work as a generalized account of attitude reports, because it does not track factivity and (dis)agreement involving non-exocentric belief, where belief tracks individuals' own experiences in a way that allows them to take the believed propositions to be true *simpliciter*. To remedy this, a *de se* denotation that tracks the notion of belief *simpliciter* was introduced (cf. (90b)). The denotation in (192) cannot perform this function, since it only allows reports as to whether an agent believes that individuals are frightening, etc. to some experiencer or other, reviving all the old contextualist problems (cf. Ch. 1, fn. 41).

The *de se* notion of belief for non-exocentric belief reports is thus needed as well. Lasersohn recognizes this (cf. Lasersohn 2017: §8.6), and so offers a notion of properly 'autocentric' belief, which corresponds to what has been referred to in the present work as belief *simpliciter*. But this notion of belief, insofar as it is tracked by attitude reports, would have to be introduced by a separate lexical item that lacks an experiencer argument, or has no existential quantification over experiencers (as entertained in Lasersohn 2009: 366-367). The result is the multiplication of attitude expressions, and the effective positing of a special lexical item like (192) meant specifically to take care of exocentric *think*-reports, alongside the ordinary non-exocentric item. The same move would have to be made for every sort of attitude predicate that embeds exocentric reports, and so there is an explosion of homophonous attitude predicates.

Further, this doubling of attitude predicates doesn't increase the expressive power of the semantics. This is because Lasersohn defines belief in a proposition relative to an experiencer as identical to autocentric (*simpliciter*) belief in a proposition where the content of the

experiential predicate is itself relativized to an experiencer.[36] Thus, to believe *simpliciter* that the movie is frightening *to Alfonse* is exactly the same as to believe, relative to Alfonse as experiencer, that the movie is frightening *simpliciter*. If implicit relativizations of the predicate to an experiencer were allowed, then the single notion of belief *simpliciter* would therefore suffice. The refusal to encode exocentricity in the content of the predicate thus results in a purposeless proliferation of attitude predicates.

Second, a semantic concern: as noted above, exocentric occurrences of predicates give rise to direct experience presuppositions, while non-exocentric occurrences do not. If an identical proposition were to be assessed in both cases, and the only difference between the two assessments is the value contextually supplied for the experiencer parameter, this difference is inexplicable.[37] But it will be made sense of in what's to come, where exocentricity is cast in terms of content.

Third, a pragmatic concern: this account predicts that where a speaker utters a sentence like (187a), intending an exocentric use of the predicate, that speaker ought to be competently taken to have the truth value of what they said change as the pragmatic concerns of truth assessment change at a different context of assessment. Whenever the relevant experiencer for exocentric evaluation is someone to whom the movie wasn't frightening, one ought to think that the speaker said something false, even where the speaker at the context of utterance intended the utterance to talk about Alfonse's experiences, all interlocutors at the context of utterance understood and were mutually agreed on this, and the movie indeed was frightening to Alfonse at that time.

In other words, where the target of exocentric assessment is not fixed by the context of utterance, but varies with the context of assessment, one ought to find that the truth of

36. Cf. *ibid.* 156 for the definition of belief relative to an experiencer, cast in terms of the notion of 'perspectives,' *ibid.* 190-191 for the definition of autocentric belief, and *ibid.* §5.5 for the semantics of relativization of predicates to overt experiencers.

37. Cf. (189a), which as written must invoke a direct experience presupposition, regardless of the value supplied by the parameter. If the denotation were changed to void the presupposition, then it would fail to invoke such a presupposition, again regardless of the value supplied for the parameter. The pattern that cannot be captured is the presence of the presupposition in one case but not the other.

exocentric utterances varies from context to context. But this apparently does not happen: where (187a) originally speaks about Alfonse's experience, the speaker has spoken truly so long as Alfonse experienced the relevant fear, and cannot 'come to have' spoken falsely in other circumstances, in virtue of the pragmatic concerns of the competent assessor.[38]

Of course, the pragmatic account of exocentricity can always appeal to norms that govern when it is or is not appropriate to assess for truth relative to a certain experiencer (cf. Lasersohn 2017: 143), and so give some pragmatic story about why an utterance like (187a), when initially anchored to Alfonse, will typically, or maybe invariantly, have to be assessed relative to Alfonse as experiencer.[39] But to do this is just to say that the norms of assessment for exocentrically evaluated statements track the circumstances of the context of utterance – in other words, that exocentric evaluation happens as if it were not governed by the context of assessment at all. But this is just what one would expect anyway if the context of utterance fixed what experiencer the predicate was relativized to, and so the content of the sentence itself encoded the relevant experiencer.

For these reasons, the present approach holds that exocentricity is a matter of the semantic content of the experiential predicate, and not a feature of the pragmatics of truth

38. This apparently differs from 'bare' experiential predicates meant to be assessed autocentrically: one can indeed, in virtue of changing one's own experiential dispositions, come to think that such a proposition, which one once thought was true, is instead false on further reflection. And so speakers can 'change their minds' about truth evaluation of experiential predicates, based solely on a change in the experiences relative to which the predicate is assessed, where non-exocentric readings are at stake: cf. MacFarlane (2007: 20-21), on changing one's mind about whether fish sticks are delicious. Stephenson (2007a: 521-522) notes in a similar vein that exocentric readings, because their interpretation is fixed to the context of utterance like an indexically anchored value, tend not to license so-called 'faultless disagreement' (cf. Ch. 1, fn. 38). Lasersohn (2017: 149-150) claims that they actually do allow for fautlless disagreement, and that interlocutors can, even in an exocentrically-minded context, refuse to coordinate their evaluations; but his examples show no more than that speakers can refuse to be cooperative, and contest how the parameters fixing content at a context of utterance are to be set, and this a feature of context-sensitive language with covert material generally.

39. Even this latter option will probably not work, for it is the *proposition* that is assessed for truth, not the sentence or speech act. It cannot be that the proposition *that the movie was frightening* becomes permanently anchored to Alfonse for evaluation: so the pragmatic theory of exocentricity must appeal to yet finer-grained notions of assessment, whereby one is obliged to assess a certain proposition relative to a certain experiencer only when one is considering how that proposition was expressed by a certain speech act. The idea bears some resemblance to MacFarlane's (2009) notion of 'nonindexical contextualism,' which in principle distinguishes truth of utterances from assessed truth of propositions those utterances express – but truth assessment apparently does not work this way.

assessment. Having said that, there is still the question of how to implement the semantic relativization of an experiential predicate to an implicit experiencer. One possibility is simply to posit silent experiential modifiers, which re-introduce the experiencer desaturated by *-ing*: this was entertained in §1.1 and §1.2 in sketching a contextualist and relativist semantics for exocentricity, respectively. §2.3.3 will demonstrate that the present treatment can handle such experiential modifiers when they appear overtly, as with *to*-headed PPs, and so it is formally trivial to offer the same treatment to covert experiencers. Some story would then need to be told about the distribution of these covert experiencers, and the pragmatic conditions under which experiential predicates without overt experiencers are interpreted as having a covert experiencer, rather than as bare.

However, doing this would miss a generalization regarding the behavior of experiential predicates as special cases of individual-level predicates (cf. §2.3.1). The alternation between non-exocentric and exocentric interpretations of experiential predicates is apparently just one instantiation of a wider phenomenon: a huge number of individual-level predicates occur both with dispositional, generic-flavored interpretations, as well as episodic interpretations that report an actualization of this disposition. Where this distinction is implemented with respect to experiential predicates specifically, the result is exactly the non-exocentric/exocentric split: the former are dispositional and generic with respect to the conditions under which the stimulation of an experience occurs, and the latter report the stimulation of an experience.

On this view, exocentricity is not a phenomenon dependent on the behavior of covert experiencers, but rather arises out of the genericity of individual-level predicates and experiential semantics. To illustrate this behavior of individual-level predicates, consider adjectives like *nice*, *smart*, and and *helpful*. To be nice is to be disposed under the right felicity conditions to nice behavior; to be smart is to be disposed under such conditions to intelligent behavior; to be helpful is to be disposed under such conditions to help. These dispositional readings of the predicate are typical, but they all have counterparts that report an actual-

ization of that disposition: thus, *Alfonse was smart*, read episodically, can report that there was an event of Alfonse displaying intelligent behavior (even if Alfonse himself is not smart), and so on for the other predicates.

Since on the present treatment, experiential adjectives have no internal argument, it is useful to look at a non-experiential example that also plausibly has no such argument in order to make the parallelism explicit. This can be done with the exemplar *fast*, taken to mean 'disposed to move at a high speed.'[40] This dispositional reading is the default with ordinary predications in the present tense.

(193) Alfonse is fast.

This sort of reading does not require that Alfonse's disposition to move quickly has ever been actualized. This is made more apparent with the right embedding environments, as beneath *wonder*: (194) might be uttered by a track coach scouting for a sprinter, even where this coach knows that Alfonse has never run seriously before, and so has never moved quickly in the relevant sense. What is wondered is whether, under the right felicity conditions, Alfonse would or could move at a high speed.

(194) I wonder whether Alfonse is fast.

But in the past tense, or with predictive *will*, there is no such default to the dispositional reading: both this and the episodic, actualized reading are available. Thus (195a) can mean either that in the past (say, when he was younger or healthier) Alfonse had an ideal disposition to move at a high speed, or that there was some single event during which Alfonse moved at a high speed. Likewise for (195b), with a future orientation: it might mean that Alfonse will be ideally disposed to move at a high speed at some future time, or that during some future event (say, an upcoming race), Alfonse will move at a high speed.

(195) a. Alfonse was fast.

40. This is as opposed to *fast* read with respect to the time taken to finish accomplishments, which plausibly takes an internal argument manifesting as an *at*-phrase, e.g. *fast at trimming trees*. *Nice*, *smart*, and *helpful* also plausibly have such arguments, as shown in e.g. *nice to Alfonse*, *smart to leave early*, and *helpful to me*.

b. Alfonse will be fast.

Surrounding context can also make the episodic, actualized readings more salient, or even default.

(196) a. There was a race yesterday. Alfonse was fast.

b. There is a race tomorrow. Alfonse will be fast.

The second clauses in each of these examples can be true, even where Alfonse himself is not 'generically' fast at these times (and so where conditions are for some reason abnormal or non-ideal): all that matters is that Alfonse moves at a high speed during the relevant event.

In general, where the past and future tense are read perfectively, the actualized, episodic reading occurs. Where they are read imperfectively, the idealized, dispositional reading is the default, and occurs in the absence of strong contextual clues to the contrary. This, combined with the classical view that the English present tense is imperfective, explains why the present tense in particular discourages the actualized reading without strong contextual cues. This is granted further support by the fact that it is precisely where the English present tense can be read perfectively, i.e. in the narrative or historical present, that the actualized reading reappears as the default. Thus where (197) is narrating an event, rather that reporting on Alfonse's habits, the second clause again reports that Alfonse moves at a high speed during the presently narrated event.

(197) Alfonse takes off from the starting line. He's fast.

This pattern of behavior is replicated exactly with experiential predicates, both with deverbal psych adjectives, and adjectives like *tasty*. Thus (198a) takes a non-exocentric reading as a strong default, but can be steered toward an exocentric reading where the contextual clues supporting this reading are overwhelming. (198b) and (198c) by contrast more freely take either exocentric or non-exocentric readings depending on the context: they

can either be read as saying that at some past or future time, the movie had or will have the ideal disposition to stimulate fear, or that there was or will be some eventuality of the movie actually stimulating fear. Finally, contexts that coerce a preterite reading of the copula, as in (198d) and (198e), suggest the exocentric reading as the default. The reader can confirm that the same patterns hold for any experiential predicate, including e.g. *tasty*.

(198) a. The movie is frightening.

b. The movie was frightening.

c. The movie will be frightening.

d. Alfonse saw a movie. It was frightening.

e. Alfonse is going to see a movie. It will be frightening.

The narrative present further allows exocentricity to surface as a default even in the present tense, and an explicit imperfective marker like *used to* forces the exocentric reading to require strong contextual clues once again: (199b)'s most natural reading is one on which at some time in the past, the movie had a disposition to produce fear (which by implicature it no longer has).

(199) a. Alfonse watches a movie. It's frightening.

b. The movie used to be frightening.

A full explanation of where and why these dispositional/actual differences occur among individual-level predicates would require an in-depth examination of the aspectual system of English, which is beyond the scope of this treatment. The claim is, however, that exocentricity in experiential predicates and actualized readings of individual-level predicates that have corresponding dispositional readings pattern alike in their behavior, as to where they are the default, and where they require contextual cues to be felicitous.[41]

41. Lasersohn (2005: 671) suggests also that exocentric readings become increasingly felicitous in free indirect discourse. The examples he provides are questionable: free indirect discourse generally is capable of reporting according to a shifted narrator's opinion, and it may be this fact, that said opinions are from this

At present, this distinction can be implemented by holding that these individual-level predicates can occur with or without a generic operator GEN internal to the adjective. The presence of this operator results in the dispositional reading, and its absence results in the actualized reading.[42] Thus the denotation of (the positive form of) *fast* is as follows, where $fast'(w)(x)$ is true iff x moves at a high speed at w.

(200) $[\![\text{fast}]\!]^w = \lambda x_e.fast'(w)(x)$

GEN then operates on the predicate and quantifies over worlds in which the relevant felicity conditions are met (201). The accessibility relation ρ_P is that relation relevant to the felicity conditions of P: in the case of *fast*, this tracks the relevant conditions for moving at a high speed. The generic reading of *fast* is then as in (202): it denotes the property of being disposed under the relevant felicity conditions to move at a high speed.

(201) $[\![\text{GEN}]\!]^w = \lambda P_{s,et}.\lambda x_e.\forall w' : w\rho_P w'[P(w')(x)]$

(202) $[\![\text{GEN fast}]\!]^w = [\![\text{GEN}]\!]^w(\lambda w_s.[\![\text{fast}]\!]^w)$
$= \lambda x_e.\forall w' : w\rho_{\lambda w''_s.\lambda y_e.fast'(w'')(y)}w'[fast'(w')(x)]$

Importing this treatment to an experiential predicate yields the desired results. *Frightening* in its positive form then just denotes the property of being the stimulus of a non-zero degree of fear ((203), repeated from (172)). GEN composes with this to create the denotation proposed in §2.3.1 (cf. (184), which is equivalent to (204) given the assumption that the accessible worlds converge on a single degree of the experiential kind), assuming that

shifted narrator's 'point of view' generally, that gives the illusion of exocentricity in the case of experiential predicates: cf. §4.1.4 on the connection between evaluative beliefs and experiences. As such the present account ignores these cases.

42. This treatment differs from that in Chierchia (1995), and discussed in §2.3.1, in that the generic operator is not a tense/aspect element merely licensed by the adjective. The reason for this is that the appearance of generic readings of adjectives is apparently not straightforwardly determined by the aspectual features of a clause. Where individual-level predicates appear in attributive position, they can have dispositional or actualized readings, even where the aspect of the clause as a whole is episodic, as in *Alfonse chased the fast runner*. The same thing happens with experiential predicates: *Alfonse ordered a tasty meal* has both an exocentric and non-exocentric reading of *tasty* available. But as noted above, an episodic clause does apparently require an exocentric reading for an individual-level predicate in predicate position.

for any experiential property P, the relevant accessibility relation is that governed by the corresponding experiential kind.

(203) $[\![\textsc{pos}\ \text{frightening}]\!]^w = \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y[stim''(w)(x)(y)])$

(204) $[\![\textsc{gen}\ [\textsc{pos}\ \text{frightening}]]\!]^w = [\![\textsc{gen}]\!]^w(\lambda w_s.[\![\textsc{pos}\ \text{frightening}]\!]^w)$
$= \lambda x_e.\forall w' : w\rho_{\text{FEAR}}w'[0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y[stim''(w')(x)(y)])]$

(203) is then the denotation for exocentric *frightening*, while (204) is the denotation for non-exocentric *frightening*, and so on for any experiential predicate. §2.3.1 showed how the non-exocentric denotation results in no presuppositions of direct experience. That the exocentric interpretation does carry such a presupposition can be seen from the fact that the term '$\sigma y[stim''(w)(y)(x)]$' is defined only where x, the stimulus, is indeed the stimulus of some experience at the world of evaluation w, and that where defined, '$\delta_{\text{FEAR}}(\sigma y[stim''(w)(y)(x)])$' denotes the degree of fear stimulated by this experience. But where there is a stimulus, there must be an experiencer, due to the nature of experiential states (cf. §2.2.2); and where there is a stimulus of fear, there is an experiencer that has experienced fear.

The exocentric predicate simply holds where some relevant experiential state holds – there is no explicit reference in the denotation to what the relevant experiencer is. This means that where exocentric reports are intended to target specific experiencers, speakers must coordinate to figure out which experiencer is relevant – the harder this is to do, the less plausible an exocentric interpretation becomes. Since the exocentric interpretation in effect enforces an existential generalization over experiencers (cf. the discussion of (172) in §2.3.1), narrowing down the interpretation to the relevant experiencer can be implemented via domain restriction on σ, which is anchored to a parameter fixed at the context of utterance, plausibly like all domain restriction. Thus (198a)'s interpretation, where Alfonse's experiences are relevant, would have the predicate's denotation as follows, where r_c is the restriction relevant in the context of utterance c.

(205) $[\![\textsc{pos}\ \text{frightening}]\!]^{c,w} = \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y : r_c(w)(y)[stim''(w)(x)(y)])$

(206) $[\![$[the movie] [is [POS frightening]]$]\!]^{c,w}$

$= 0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y : r_c(w)(y)[stim''(w)(\iota x[movie'(w)(x)])(y)])$

The extension will target Alfonse's experiences where r_c is any relevant property of experiences that limits the domain of quantification to Alfonse's experiences (this could be anything, such as: $r_c(w)(x)$ is true iff x is an experience of someone watching the movie at w, given that only Alfonse is watching the movie at w). In this case, it will be presupposed that Alfonse has experienced the subject of predication, and asserted that this subject is the stimulus of a non-zero degree of fear in him.[43]

The present treatment, on which exocentric predications simply assert that some experiential state or other holds, and it is up to the contextual coordination of domain restriction to narrow down to more determinate reports regarding specific experiencers, has a couple of beneficial results. First, as Hirvonen (2014: 99 ff.) notes, exocentric experiential predicates display some behavior with respect to negation and contradiction that is surprising if these predicates simply encode some definite experiencer in their semantic content. For instance, (207) reads as a contradiction, regardless of the intended exocentric targets.

(207) #The movie is frightening, but it isn't frightening.

This is not a feature of overt reference to experiencers, which uncontroversially do specify some definite experiencer in their content: (208) is not a contradiction.

(208) The movie is frightening to Alfonse, but it isn't frightening to Bethany.

43. Speakers' intensional reasoning may reveal what sort of restriction they have in mind. So an exocentric reading of *If the movie were frightening...*might consider counterfactuals where the movie is frightening to whoever watches it, whether this is Alfonse or not, or it might consider just counterfactuals where the movie frightened Alfonse. Where Alfonse is the only one actually watching the movie, the two exocentric readings will be extensionally equivalent. Note also that exocentric readings will allow binding, so long as the domain restriction includes a free variable: thus where $r_c = \lambda w'_s.\lambda z_e.exp''(w')(\iota v[see''(w)(x)(v)])(z)$, '$x$' being the individual variable bound by 'λx_e' in the denotation of POS *frightening*, an exocentric reading of *Everyone saw a frightening movie* will be true iff everyone saw a movie frightening to the one who saw it. If this is right, then bound readings of experiential predicates require setting a special domain restriciton, rather than just binding a pre-existent variable, which may predict that bound readings of these predicates are more difficult to access than those arising with adjectives that more plausibly project variables, like *local*. This would need to be tested, but the present treatment is welcome if there really is such a difference.

But if an exocentric reading of an experiential predicate simply specifies some experiencer or other in each case, there is no *a priori* reason why (207) ought not to have a non-contradictory reading, where the experiencer is Alfonse in the first clause, and Bethany in the second.

This behavior is straightforwardly predicted if exocentric predications merely assert that some experiential state holds, and narrowing down to a relevant experiencer happens by coordination on the domain restriction. In this case, where the two clauses of (207) must have their contents determined with respect to the same context of utterance, a contradiction results no matter how the interlocutors attempt to set the domain restriction. The denotation would be as follows, where no value of r_c salvages the contradiction.[44]

(209) $[\![207]\!]^{c,w} = 0_{\text{FEAR}} < \delta_{\text{FEAR}}(\sigma y : r_c(w)(y)[stim''(w)(\iota x[movie'(w)(x)])(y)]) \wedge$
$\neg 0_{\text{FEAR}} < \delta_{\text{FEAR}}(\sigma y : r_c(w)(y)[stim''(w)(\iota x[movie'(w)(x)])(y)])$

Second, the present account predicts that where the interlocutors are not interested in making reference to a specific experiencer, the domain restriction can be omitted or left vacuous, resulting in genuinely 'existential' readings of exocentric predicates, which simply report that there is an experiential state. This would entail that the state holds with respect to some experiencer or other, but there would be no specification at all in the semantic content of who this is, resulting in exocentric readings of *frightening* roughly equivalent to *frightening to someone or other*.

There do in fact appear to be such readings of experiential predicates. Where Alfonse hears that a horror movie has caused hysteria in some of its watchers, but has no idea who they are, he might say the following to report on an experiential state.

44. One needs to be cautious here, since this is apparently also a feature of adjectives with definite implicit arguments: *#The concert is local, but it isn't local.* It may ultimately be, then, that even definite implicit arguments do not simply anchor to individuals straightforwardly, but must do so mediately by some contextual parameter, resulting in contradiction wherever two clauses such as these are read with respect to the same context. The treatments of covert experiencers in §1.1 and §1.2 actually capture this pattern as well, since they receive their extension relative to a contextually supplied assignment function, and the above piece of data, while compatible with the present approach, does not yet empirically distinguish it from an approach on which exocentric predicates have covert definite experiencers.

(210) The movie was frightening.

That Alfonse is not reporting that the movie was frightening to any experiencer in particular is shown by the fact that he can felicitously avow ignorance of the experiencer, or question it (using a sluice for naturalness: cf. Schaffer 2011: 200, ex. 45), which is not possible given even a minimal identifying description of the experiencer, as with *to whoever saw it*.

(211) a. The movie was frightening, but I don't know to who.

b. The movie was frightening, but to who?

(212) a. ?The movie was frightening to whoever saw it, but I don't know to who.

b. ?The movie was frightening to whoever saw it, but to who?

This is not a feature of implicit definite material of predicates generally. Thus there appears to be no reading of *local* without an overt location argument that is read as *local to somewhere*. Supposing that Alfonse hears that there was a concert somewhere, but he has no idea where, (213) remains strange, and ignorance of or questioning of the location is strange whether or not a minimally identifying location is provided.

(213) The concert was local.

(214) a. ?The concert was local, but I don't know to where.

b. ?The concert was local, but to where?

(215) a. ?The concert was local to wherever it took place, but I don't know to where.

b. ?The concert was local to wherever it took place, but to where?

Where *local*'s denotation specifies a location argument, only a definite reading of this argument, however thinly interpreted, is possible, and the predicate must be read as *local to* x for some determinate x.

Experiential predicates thus do not specify an experiencer *simpliciter* even on their exocentric readings, but only in virtue of coordination on the restriction of the exocentric report.

In this they differ from predicates that genuinely specify a definite individual as part of their semantic content.

2.3.3 Overt experiencers: to*-headed PPs*

It's now possible to clarify what the semantic role is of overt experiencers that occur alongside experiential predicates as overt *to*-headed PPs. Some examples of the relevant phenomenon, as it occurs with deverbal psych predicates, are as follows.

(216) a. The movie is frightening to Alfonse.

b. Walking on the pier is relaxing to Bethany.

c. Every piece of furniture irritating to him was removed from the house.

Before offering an account of these, it's worth noting that these constructions may be instantiations of a wider phenomenon, in two senses. *To*-headed PPs of exactly this sort look to have a role that extends beyond their occurrence as overt experiencers. There are non-experiential predicates with which these phrases occur, and there they seem to have the same semantic effect: presumably, a more general account of their behavior ought to derive their behavior as overt experiencers as a special case.

The predicates in question are those that are broadly 'response-dependent,' true of an individual just in case that individual is ideally disposed to produce a certain reaction in an agent. Some of these predicates have an epistemic flavor, and include *persuasive*, *obvious*, and *clear*. To take *persuasive* as the exemplar, its similarity to experiential predicates is remarkable. Not only is its interpretation in ordinary predications paraphrasable in a way similar to that of experiential predicates in the same environment, but the evaluation of the predicate by speakers is also 'autocentric' with respect to what disposition the predicate actualizes.

Thus *persuasive* is true of an individual just in case it is ideally disposed to produce belief in some proposition by a process of reasoning (217a), and a speaker tends to think

that something is persuasive where that speaker is (disposed to be) persuaded by it. Further, the adjective relates to a verb that takes an agent as object, and asserts the actualization of the relevant disposition with respect to that agent (217b). Finally, the occurrence of a *to*-headed PP 'restores' the individual with respect to whom the disposition is actualized (217c).[45]

(217) a. The argument is persuasive.

b. The argument persuaded Alfonse.

c. The argument is persuasive to Alfonse.

Similar comments can be made with respect to the behavior of these PPs with the other adjectives: note the relation between *obvious* and *obvious to Alfonse*, and *clear* and *clear to Alfonse*. As with *-ing* above, a denotation for *to* will be provided here that only deals with its behavior as it pertains to the special case of experiencers: but §5.3.1 will expand this notion slightly in dealing with the behavior of *to* as it relates to direct evidence. Nevertheless, the cue can be taken from the epistemic response-dependent predicates: the PPs actualize a disposition to produce a state with respect to some individual – here, an experiencer – and the PP plays the role that the experiencer object of the psych verb played prior to being desaturated.

Second, it is unclear to what extent these phrases are or are not part of the broader phenomenon traditionally labeled *dativus iuidicatnis*, viz. the use of dative marking to relativize the content of predicates or propositions to individuals. In English, the use of *to*- and *for*-headed PPs to express individuals' evaluative opinions is widespread, and not restricted to experiential predicates.

(218) a. If there is alcohol, that's a party to me.

b. For Alfonse, only one Hitchcock movie is worth seeing.

45. These epistemic response-dependent adjectives even have exocentric readings, which appear in exactly the same circumstances as their experiential counterparts. Thus, on a preterite reading: *Alfonse heard the argument. It was persuasive.*

The use of dative marking for this purpose is widespread and poorly-understood – §5.3.1 offers some brief comments on it, in relation to the behavior of subjective attitude verbs. What is at stake for experiential semantics is whether the overt experiencers appearing with experiential predicates can be semantically assimilated to whatever more general mechanism *dativus iudicantis* marks. This ultimately requires a broad cross-linguistic survey of the behavior of free dative marking, to determine whether or not languages have markers specialized for the expression of experiential modifiers, or whether, even if they do not, the appearance of datives with experiential predicates has a semantic effect not reducible to a more general 'relativizing' mechanism.[46]

With that said, the denotation of *to* can be cast as introducing an experiencer as follows. Recall from §1.3.1, where it was demonstrated that a bare account of experiential predicates could relativize predicates to an experiencer by quantifying over 'experiential alternatives,' that the following denotation was given (repeated from (95a)).

(219) $[\![\text{to}]\!]^w = \lambda x_e.\lambda P_{s,et}.\lambda y_e.\forall w' \in Exp_{x,w}[P(w')(y)]$

With the machinery put in place this chapter, the notion of experiential alternatives can be given content. Some abbreviations will make the exposition more readable. The foregoing has appealed to two sorts of degrees of experiential kinds produced by stimuli. The first is a dyadic notion, invoked for the denotation of psych verbs in §2.2.2: this is the degree of an experiential kind produced at a world by a stimulus in an experiencer. The second is a monadic notion, invoked for the denotation of deverbal psych predicates in §2.3.1: this

46. If it is true that overt experiencers are a sub-class of *dativus iudicantis*, then the contextualist idea that they are manifestations of an argument that a transitive predicate selects for (cf. §1.1) is yet more dubious, since *dativus iudicantis* in the above sense is uncontroversially adjunctive in character. Stephenson (2007a: 520, ex. 85) has claimed that there is evidence for overt experiencers being internal arguments to the adjective, from the idiosyncratic lexical requirements as to which preposition must head the PP; however, the data are extremely messy, far more so than one expects for lexical requirements typical of the selectional behavior of adjectives, and all of Stephenson's judgments are at least questionable. *To*-headed PPs are examined in the text, because their occurrence with experiential predicates, especially those derived by compositional processes encoding experience-sensitivty, is extremely regular. McNally & Stojanovic (2017) and Bylinina (2017) make the case that certain predicates do occur with markers that have a narrowly experiential function, in part based on the uneven distribution of experiential dative markers with predicates cross-linguistically.

is the degree of an experiential kind of that an individual is disposed to be the stimulus of when certain felicity conditions are met. These two sorts of degrees can be abbreviated as '$\delta''_k(w)(x)(y)$' and '$\delta'_k(w)(x)$,' respectively.

(220) a. $\delta''_k(w)(x)(y) := \delta_k(\sigma z[exp''(w)(x)(z) \wedge stim''(w)(y)(z)])$

b. $\delta'_k(w)(x) := \iota d[\forall w' : w\rho_k w'[d = \delta_k(\sigma y[stim''(w')(x)(y)])]]$

'$\delta''_k(w)(x)(y)$' is thus read: 'the degree of k produced in x by y at w,' while '$\delta'_k(w)(x)$' is read: 'the degree of k that x is disposed to produce at w (*simpliciter*).'

The experiential alternatives of an experiencer are all the live options for the way the world is 'according to the experiences of the experiencer.' In other words, at w, x's set of experiential alternatives contains all and only those worlds at which each individual is disposed to produce the degree of every experiential kind that is (actually) produced by that individual in x at w. At each experiential alternative, then, each individual produces *simpliciter* that experience that the experiencer actually has at the anchor-world, meaning that the experiential alternatives track what the world would be like if the experiencer's experiences were all 'accurate,' or correctly tracked the sorts of experiences that stimuli are 'really' disposed to produce.

(221) $Exp_{x,w} := \{w' : \forall y, k[\delta'_k(w')(y) = \delta''_k(w)(x)(y)]\}$

And now the denotation of *to* in (219) works as intended. It composes with an individual (experiencer), then with the intension of a property, and returns a property true of an individual just in case the first property is true of that individual in all the experiencer's experiential alternatives. The composition of (216a) proceeds as follows, making use of the new notation, taking the non-exocentric *frightening*, in line with §2.3.2, to be in its generic positive form, and taking the experiencer phrase to adjoin after composition with GEN.[47]

47. Where GEN is internal to the adjective, this makes the PP an adjectival adjunct. I take this to be the right structural analysis, given constructions like (216c), and given the fact that, as Schaffer (2011: 197, exx. 29-32) points out, experiencer phrases are degraded when following verbal adjuncts: compare *Hitchcock movies are frightening to Alfonse at night* versus *?Hitchcock movies are frightening at night to*

(222) a. $[\![\text{GEN [POS [frightening]]}]\!]^w = \lambda x_e.0_{\text{FEAR}} < \delta'_{\text{FEAR}}(w)(x)$

b. $[\![\text{to Alfonse}]\!]^w = [\![\text{to}]\!]^w([\![\text{Alfonse}]\!]^w)$

$= \lambda P_{s,et}.\lambda y_e.\forall w' \in Exp_{a,w}[P(w')(x)]$

c. $[\![[\text{GEN [POS [frightening]]}]\text{ [to Alfonse}]]\!]^w$

$= [\![\text{to Alfonse}]\!]^w(\lambda w_s.[\![\text{GEN [POS [frightening]]}]\!]^w)$

$= \lambda y_e.\forall w' \in Exp_{a,w}[0_{\text{FEAR}} < \delta'_{\text{FEAR}}(w')(y)]$

d. $[\![[\text{the movie}]\text{ [is [[GEN [POS [frightening]]}]\text{ [to Alfonse}]]]]\!]^w$

$= [\![[\text{GEN [POS [frightening]]}]\text{ [to Alfonse}]]\!]^w([\![\text{the movie}]\!]^w)$

$= \forall w' \in Exp_{a,w}[0_{\text{FEAR}} <_{\text{FEAR}} \delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])]$

Thus, *the movie is frightening to Alfonse* is true just in case in all of Alfonse's experiential alternatives, the movie is disposed to produce a non-zero degree of fear. Due to the way experiential alternatives are defined in (221), it must be that the value of $\delta'_{\text{FEAR}}(w')(x)$ is uniform for all x and $w' \in Exp_{a,w}$ – and this value must match $\delta''_{\text{FEAR}}(w')(a)(x)$. The denotation in (222d) can therefore be simplified as follows.

(223) $0_{\text{FEAR}} <_{\text{FEAR}} \delta''_{\text{FEAR}}(w)(a)(\iota x[movie'(w)(x)])$

Thus, it is equivalent to say that *The movie is frightening to Alfonse* is true just in case the movie stimulates a non-zero degree of fear in Alfonse, and the result is as desired. Furthermore, this is identical to the denotation of *The movie frightens Alfonse*, which confirms that the overt experiencer fills the same role as the object of the object-experiencer psych verb, and simply 'resaturates' the experiencer. Recall the denotation for this latter sentence from §2.2.2 (repeated from (158)).

Alfonse. Schaffer takes this fact to be evidence that the experiencer phrase is an argument, that cannot have an adjunct interposed between it and its selector, but this reasoning only works against an adjunct analysis if both the experiencer and temporal phrases are taken adjoin to the same projection. Since it is independently implausible that *at night* can function here as an adjectival rather than verbal adjunct (cf. ?*the movies frightening at night*), and on the present analysis the experiencer phrase is not a verbal adjunct, then the ordering constraint is exactly what is predicted, since the verbal modifier adjoins higher. While Lasersohn (2005: 666, ex. d) and MacFarlane (2014: 153) don't discuss the syntactic status of these phrases, their semantic analysis of them as property modifiers is compatible with this account, and Collins (2013: 91, ex. 62) proposes an AP-adjunct analysis explicitly. Cf. also Lasersohn (2017: §5.5).

(224) $[\![$[the movie] [POS [frightens Alfonse]]$]\!]^w$

$= 0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma z[exp''(w)(a)(z) \wedge stim''(w)(\iota x[movie'(w)(x)])(z)])$

But given the definition in (220a), (224) is identical to (223). This further means that overt experiencers inherit all the presuppositional properties of object-experiencer psych verbs: they require that the experiencer be the possessor of an experience of the relevant kind, which for the reasons familiar from §2.1.2 requires that a certain experiential episode has been had. This is the right result, since experiencer PPs do in fact enforce these presuppositions, and they project out of negation in the expected way.

(225) a. The movie is frightening to Alfonse.

↪ Alfonse has experienced fear of the movie.

b. The movie isn't frightening to Alfonse.

↪ Alfonse has experienced the movie, and experienced no fear of it.

These is in turn are just the sort of experiential presuppositions predicted to be enforced, by Alfonse's possessing a portion of fear stimulated by the movie, or in virtue of his having been stimulated by the movie but failing to possess any such portion of fear.

This result follows from the equivalence of the denotations shown above, but it's worth briefly illustrating. In the case of a positive predication, as in *The movie is frightening*, the denotation is as in (222d). Because '$\delta''_{\text{FEAR}}(w)(a)(\iota x[movie'(w)(x)])$' denotes a single value, due to the definition of experiential alternatives in (221), there must similarly be a single value for '$\delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])$' for all Alfonse's experiential alternatives w', and the two values must be equivalent. There are three options as to what this value could be.

First, the value may be a non-zero degree: by definition, this is just where (i) the movie has been the stimulus of an experience in Alfonse, and (ii) Alfonse is an experiencer of a portion of fear whose stimulus is the movie. Here, the presuppositions have been met, and the sentence is true. Second, the value may be a zero-degree. In this case, by definition, (i) the movie has been the stimulus of an experience in Alfonse, and (ii) Alfonse is not the

experiencer of a portion of fear whose stimulus is the movie. Here, the presuppositions have been met, and the sentence is false. Finally, it may be undefined, where the movie is not the stimulus of an experience in Alfonse. In this case, since $\delta''_{\text{FEAR}}(w)(a)(\iota x[movie'(w)(x)])$ is undefined, so is $\delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])$ for every w' in Alfonse's experiential alternatives; but then, the denotation in (222d) is also undefined, since the universal quantifier scopes over values undefined on all the relevant assignments. Here, there is presupposition failure.

These presuppositions then project out of non-external negation in the desired way. Thus, where negation of the copula is just treated as predicate negation:

(226) $[\![\text{isn't}]\!]^w = \lambda P_{et}.\lambda x_e.\neg P(x)$

(227) $[\![[\text{the movie}]\ [\text{isn't}\ [[\textsc{gen}\ [\textsc{pos}\ [\text{frightening}]]]\ [\text{to Alfonse}]]]]\!]^w$
$= \neg\forall w' \in Exp_{a,w}[0_{\text{FEAR}} < \delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])]$

The result is true just in case $\delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])$ is not a non-zero degree of fear in all of Alfonse's experiential alternatives w'. Since the value is uniform across the alternatives, this is the same as to say that it is the zero-degree of fear in all these alternatives. Thus again, where $\delta''_{\text{FEAR}}(w)(a)(\iota x[movie'(w)(x)])$ is undefined, so is $\delta'_{\text{FEAR}}(w')(\iota x[movie'(w)(x)])$ for all w', and the denotation is undefined, since again the universal quantifier operates on undefined values, and thus so does '$\neg$.' This results in an undefined value for the denotation as a whole, hence presupposition failure, where the movie has not been the stimulus of an experience in Alfonse. Where the term is defined, due to Alfonse having an experience of the movie, truth and falsehood are just reversed from the positive case, as desired.

While the above treatment captures the core truth-conditional contribution of overt experiencers when they appear, including their presuppositional constraints, it has nothing to say about their limited distribution. Experiencer PPs do not stack (228a), they typically do not occur with adjectival predicates that are non-experiential and non-evaluative (228b), and they are infelicitous even with experiential predicates in predicate position of a small clause beneath *find* (228c) (cf. §1.1.1).

(228) a. #The movie is frightening to Alfonse to Bethany.

b. #The table is wooden to Alfonse.

c. #Alfonse finds the movie frightening to {Bethany / him(self)}.

The semantics provided above, in of itself, offers no reason why this should be. It's also important that overt experiencers not attach to exocentric occurrences of experiential predicates, which according to §2.3.2 simply lack an internal GEN operator, and so according to the above should allow experiencer PPs equally to adjoin to them.

What all these cases of disallowed overt experiencers have in common is that they are apparently not attaching to experiential predicates in the proper, evaluative sense that this work has been examining. *Frightening to Alfonse* is a predicate that simply reports Alfonse's experiences, and so has no evaluative component; the same is true of exocentric readings of *frightening*. Meanwhile, *wooden* has no experiential component at all, while embedding *frightening to* x beneath *find*, for any x, seems to 'doubly saturate' the experiencer, resulting in *find* scoping over a predicate that is already non-evaluative. The reason for these restrictions will be elucidated in §5.3.1, where *to* is discussed in the context of evaluativity and direct evidence.

With that, there is a working account of the semantics of experiential predicates and their relation to experiencers. Before ending this chapter, it's worth considering to what extent the above allows the revitalization of a contextualism of a certain sophisticated sort, that the criticisms leveled in §1.1 don't directly apply to. According to the above, the non-exocentric reading of experiential predicates, by which they gain their evaluative character, is the result of an operator GEN, and the absence of this operator results in exocentricity. This view has an obvious affinity with the accounts provided in Snyder (2013) and Pearson (2013), which were surveyed in §2.3.1. Why, then, might not a version of one of these views survive?

In that section, it was pointed out that there is no truth-conditional reason to believe in generic quantification over experiencers, since the pragmatics would have to deflate this generic quantification into what is effectively a relativist, autocentric view. But this in of

itself does not count against the position, if there in fact is some independent pragamatic reason that generic quantification over experiencers tracks autocentric evaluation. Then said quantification is empirically harmless, and so the lack of positive motivation would only show that the two accounts – that experiential predicates have an internal experiencer argument, and that they don't – are indifferently adoptable on semantic grounds.[48]

Here is a sketch of such a contextualist account adopting the present machinery. First, the deverbal psych predicate receives a dyadic interpretation (229), and the preposition is vacuous (230) (cf. (4a)).

(229) $[\![\text{frightening}]\!]^{c,w,g} = \lambda x_e.\lambda y_e.\delta''_{\text{FEAR}}(w)(x)(y)$

(230) $[\![\text{to}]\!]^{c,w,g} = \lambda x_e.x$

POS then works as before (231), as does GEN, except that it comes with an index, call it α, for an individual variable it can bind, and quantifies over both worlds and individuals (232). It further asserts felicity conditions on the individual as well as the world, represented by the operator F: presumably, some pragmatic mechanism will make it such that the only 'felicitous' individuals to which GEN is sensitive when composing with experiential predi-

48. Of course, this does nothing to show that the two accounts are indifferently adoptable on non-semantic grounds. Because the present study is concerned narrowly with experiencer semantics, the text will not address these, but there are several points to consider in the ultimate evaluation of the merits of these views. First, a contextualist in the present predicament will effectively be committed to thinking that experiencer phrases are overt realizations of an internal argument – this is because if they are not, since there are no other plausible candidates as to what does realize the argument overtly, this will result in the highly unpalatable claim that the predicates have internal arguments that are not overtly realizable (unless, as entertained in §1.1.2 and §1.1.5, the predicate encodes an experiencer intrinsically, and the PP is an assignment-modifier: but this appears to make the contextualist and bare views all but indistinguishable in the present machinery). But if they are selected for by the predicate, then there ought to be independent syntactic evidence that said experiencer phrases behave as arguments rather than adjuncts, and the evidence here is *prima facie* worrying for the contextualist: cf. Collins (2013: §4) for a summary of some basic tests (some questionable) that point to adjuncthood, and Rákosi (2006) for another approach, possibly congenial to the contextualist, on which certain experiencers are 'thematic adjuncts.' Second, if *-ing* really does operate by a unified mechanism for object-experiencer psych verbs and other sorts of verbs (cf. §2.3.1), then the contextualist will have to explain why it desaturates the other arguments of the verb in the non-experiential cases, but retains the transitivity of psych verbs. Finally, if these phrases are arguments, insofar as contextualism is a general thesis about experiential predicates cross-linguistically, all experiential predicates ought to be able to occur with them. There is some preliminary evidence that this is not so: cf. Vardomskaya (2018: 10, ex. 16). On the other hand, an adjunct analysis allows in principle that particular languages may simply lack the relevant modifiers.

cates are the evaluators of the predicate, i.e. oneself (or those that one shares the relevant experiential dispositions with).[49] Intensions now include reference to variable assignments, of type a.

(231) $[\![\text{POS}]\!]^{c,w,g} = \lambda G_{a,\langle s,et\rangle}.\lambda x_e.0_G <_G G(x)$

(232) $[\![\text{GEN}_\alpha]\!]^{c,w,g} = \lambda P_{a,\langle s,et\rangle}.\lambda x_e.\forall w' : w\rho_P w', y : F(w')(y)[P(g[\alpha \to y])(w')(x)]$

Finally, as in §1.1, there is a silent variable for which the adjective can select, which can be coindexed with GEN, and whose interpretation when unbound is appropriately relativized to the context of utterance c (cf. (12), (13)).

(233) $[\![x_\alpha]\!]^{c,w,g} = g(\alpha)$

Then three options for the interpretation of the experiential predicate occur as follows. First, the predicate may select for a *to*-headed PP, and not take the GEN operator, resulting in an overt experiencer with the right interpretation.

(234) $[\![\text{POS [frightening [to Alfonse]]}]\!]^{c,w,g} = \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta''_{\text{FEAR}}(w)(a)(x)$

Second, the predicate may select for x_α, and not take the GEN operator, resulting in an exocentric interpretation, where the experiencer is contextually specified by the variable assignment determined in c.

(235) $[\![\text{POS [frightening } x_\alpha]]\!]^{c,w,g} = \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta''_{\text{FEAR}}(w)(g(\alpha))(x)$

And third, the predicate may again select for x_α, but occur with a coindexed GEN, binding off the variable and yielding the bare, non-exocentric interpretation with the flavor of a predicate disposed to produce experience of the relevant sort when the felicity conditions are met for experiencers generally.

49. Care is needed here: insofar as GEN will operate on non-experiential predicates, generic quantification over individuals should *not* behave autocentrically in these cases. The contextualist's pragmatic mechanism will need to rein in the semantics of GEN appropriately, so that e.g. the extension of *Alfonse is nice* on a generic reading does not reduce for the assessor to 'Alfonse is nice to oneself.' That generic quantification and autocentric evaluation do not track each other in this way is *prima facie* reason not to assimilate them, until some independent reason to the contrary is discovered.

(236) $[\![\text{GEN}_\alpha\ [\text{POS}\ [\text{frightening}\ x_\alpha]]]\!]^{c,w,g}$
$= \lambda x_e.\forall w' : w\rho_P w', y : F(w')(y)[0_{\text{FEAR}} <_{\text{FEAR}} \delta''_{\text{FEAR}}(w')(y)(x)]$

This denotation gets the right results, and avoids the problems that plagued the contextualist in §1.1. It predicts that presuppositions of direct experience are voided in non-exocentric cases, and has the ability to incorporate autocentric evaluation, due to the appeal to felicity conditions. It further allows for a notion of something being 'frightening *simpliciter*,' and so on for any experiential predicate, which is the crucial reading the contextualist was missing.

However, even this treatment ultimately faces problems. The updated contextualist semantics here predicts yet a fourth reading of experiential predicates: one on which an overt experiencer is selected for, *and* GEN operates on the predicate. The result would be as follows.

(237) $[\![\text{GEN}_\alpha\ [\text{POS}\ [\text{frightening}\ [\text{to Alfonse}]]]]\!]^{c,w,g}$
$= \lambda x_e.\forall w' : w\rho_P w', y : F(w')(y)[0_{\text{FEAR}} < \delta''_{\text{FEAR}}(w')(a)(x)]$

The intended reading is that *frightening* denotes a property true of individuals that, under the appropriate felicity conditions for stimulating fear, do so *in Alfonse*. This interpretation would allow a merely dispositional reading, with no direct experience presupposition, with an overt experiencer. There is no such reading of the experiential predicate: *The movie is frightening to Alfonse* simply cannot mean that Alfonse would be frightened by the movie, were he to watch it (in ideal conditions).[50]

The logic of the predicament is this. In order to get an appropriate non-exocentric reading, where *frightening* takes an internal argument, GEN must scope over the predicate after this argument has been selected for. Otherwise, it cannot bind the experiencer variable. However, the very same operator that does this must also void the direct experience

50. Anthony (2016: 697-698) seems to propose exactly this deficient treatment of generics interacting with overt experiencers. A similar non-attested interpretation appears where GEN operates on an experiential predicate that selects for a null experiencer not coindexed with it: GEN_α [POS [frightening x_β]]. On this interpretation, there is again a dispositional reading with no direct experience requirement, which is centered on whoever the contextually relevant individual that β maps to is. To avoid this, presumably the contextualist must claim that GEN is an unselective binder.

presupposition, or else the fact that the merely dispositional reading arises precisely when interpretations are non-exocentric cannot be captured. But then, it follows that the operator must scope also over overt arguments, and so must be able to void the direct experience requirement there as well. Hence, the contextualist overgenerates, unless some principled reason can be found why GEN only appears if it has a variable to bind.[51]

By contrast, the three extant interpretations of the predicate, and the missing fourth, are naturally explained on a non-contextualist approach. GEN either appears or does not, subject to the aspectual and pragmatic restrictions discussed in §2.3.2, resulting either in a genuinely 'experiential' interpretation, in the interesting evaluative sense, or an exocentric one. In the former case, it is also possible for an overt experiencer to adjoin to the adjective over GEN: and reasons that it cannot do so below GEN will be further explored in §5.3.1. The resulting structures for the positive form are as follows.

(238) a. Genuinely 'experiential' ('evaluative') *frightening*

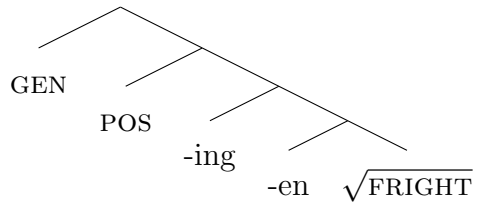

b. Exocentric *frightening*

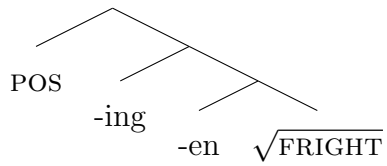

c. Overt experiencer *frightening*

51. Insofar as the data on individual-level adjectives that may have genuine internal arguments bear on this question, the contextualist may have grounds for hope. Thus *Alfonse is helpful to Bethany* indeed does seem to require that the disposition of Alfonse to help has been actualized with respect to Bethany: he must have helped her. However, §5.3.1 will show that these sorts of *to*-phrases are plausibly analyzed as adjuncts pertaining to direct evidence of a certain sort. If this is correct, then these phrases are not arguments to the adjective, either, but behave exactly like overt experiencers.

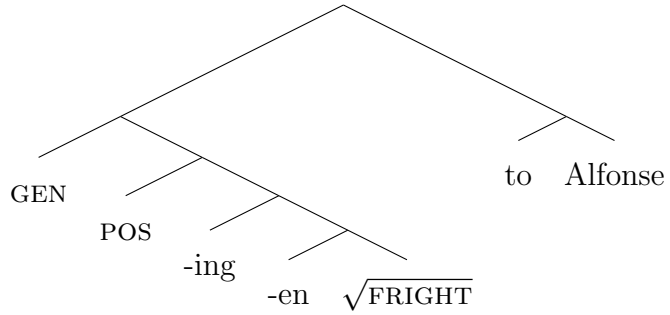

Finally, although the present chapter has given a decompositional semantics for deverbal experiential predicates, the same results can be applied to those predicates that apparently cannot be so decomposed. Where a predicate in the lexicon has all the same semantic features e.g. of a deverbal psych predicate, it can be assigned the appropriate experiential semantics just by plugging in the right experiential kind. Thus for *funny*, what is relevant is humorous experiences, and in the lexicon it can appear as follows.

(239) $[\![\text{funny}]\!]^w = \lambda x_e.\delta_{\text{HUMOR}}(\sigma y[stim''(w)(x)(y)])$

It will then interact with GEN and POS in the familiar way, to give the relevant results. This treatment can further trivially apply to experiential nominals such as *(a) bore* and *(a) horror*.

CHAPTER 3

SENSORY EXPERIENTIAL PREDICATES

This brief chapter expands on the experiential semantics provided in Chapter 2, to demonstrate how it can treat experiential predicates that encode reference to specific sensory modalities. *Sensory experiential predicates* are, like deverbal psych predicates, true of individuals disposed to produce a certain sort of experience – but they additionally enforce a causal requirement, related to a specific modality through which this experience must be stimulated.

Just as deverbal psych predicates systematically encode the kind of experience they relate to, by deriving psych verbs related to specific experiential kinds, sensory predicates are often predictable due to their phrasal composition, both with respect to the sensory modality to which they're sensitive, and with respect to the quality of the experience stimulated through that modality. This latter feature of sensory predicates may in turn relate to experiential kinds produced through the specific modality, or not, depending on the construction.

§3.1 addresses sensory verbs in construction with adjectives, e.g. *look tired* and *sound Bulgarian*. The resulting experiential predicates do not encode experiential kinds, but are instead *evidence-related experiential predicates*, and pertain to the evidence produced for a certain proposition by an experience causally linked to the sensory modality encoded by the verb.

§3.2 demonstrates that these same sensory verbs participate in special constructions with adjectives that already lie at the root of experiential predicates themselves, giving them a unique 'doubly experiential' reading, as with *sound funny* or *look disgusting*. These are *hybrid experiential predicates*, in the sense that they simultaneously encode sensitivity to a sensory modality, and an experiential kind causally stimulated through that modality. Hybrid experiential predicates are then shown to exist outside constructions with sensory verbs, in adjectives like the perennial *tasty*.

3.1 Evidence-related experiential predicates

Sensory experiential predicates occur in the construction *x V ADJ*, where *x* denotes a stimulus, *V* is a sensory verb – including *look*, *sound*, *feel*, *smell*, and *taste* – and *ADJ* is an adjective. In this construction, *V ADJ* is itself an experiential predicate: examples include *look tired* and *sound Bulgarian*.[1]

All the usual diagnostics mark these predicates as experiential (cf. the introductory portion of Chapter 2). They occur with overt experiencer phrases, and enforce direct experience presuppositions of a certain sort on their experiencer (240);[2] these presuppositions disappear in the usual intensional contexts (241); they give rise to exocentric readings in the right context, which contextually anchor their interpretation to a contextually relevant experiencer, especially in aspectual contexts that favor this interpretation, and allow for binding of experiencers (242); and in bare assertions, they give rise to acquaintance inferences anchored to the speaker (243).

(240) Alfonse looks tired to Bethany.

↪ Bethany has seen Alfonse.

(241) a. Alfonse must look tired.

b. If Alfonse looks tired, we shouldn't bother him.

c. Does Alfonse look tired?

1. These sensory verbs also compose to form experiential predicates in a variety of other syntactic configurations: they can appear in finite clauses beneath an expletive subject and a *like*-complementizer (e.g. in *It looks like Alfonse is tired*), or in an infinitival *to*-clause with proper subject (e.g. in *Alfonse looks to be tired* [apparently, this construciton is available only with *look*]), or in the so-called 'copy-raising' construction, which typically includes both a matrix subject and a coreferential pronominal in the lower clause, again with (what may be) a *like*-complementizer (e.g. in *Alfonse looks like he is tired*). Since the semantics of these don't differ interestingly from the construction to be addressed here (except where noted in §3.1.2), only the construction mentioned in the text will be directly addressed. Where needed, §3.1.2 draws some comparisons with the expletive and copy-raising constructions.

2. Unlike with predicates like *tasty*, and deverbal psych predicates, sensory experiential predicates like these cannot be embedded beneath *find* to demonstrate the same experiential effect, due to the clausal structure of the verb in English. However, this is likely a syntactic cross-linguistic accident that does not bear directly on the experiential semantics of the predicate: Coppock (2018: 126, ex. 1e) notes that Swedish *tycka*, which has certain experientially-oriented behaviors similar to English *find*, does felicitously embed clauses including sensory verbs like *se* 'look.'

d. I wonder whether Alfonse looks tired.

e. Alfonse will look tired.

f. Bethany believes that Alfonse looks tired.

(242) a. Bethany saw Alfonse emerge groggily from his office. He looked tired.

b. Every boy whose mother saw him looked tired.

(243) Alfonse looks tired.

⇝ The speaker has seen Alfonse.

In addition, these predicates clearly have an experiential denotation, in the sense exposited in Chapters 1 and 2: they are true of individuals who are disposed *simpliciter* to stimulate experiences of a certain sort. The difference is that the experiences in question, and so the kinds of experiential inferences noted above, clearly relate to a sensory modality – in the case of *look*, this is vision, but the same is of course true for auditory, haptic, olfactory, and gustatory experience with *sound*, *feel*, *smell*, and *taste*, respectively.

Predicates like *look tired* apparently do not relate to experiential kinds, unlike the deverbal psych predicates explored in Chapter 2. Instead, they pertain to the evidence that obtains for a proposition, in virtue of the experience that a stimulus produces through the relevant sensory modality. §3.1.1 notes some features of these predicates that need to be handled, and proposes some new machinery to do so; §3.1.2 then implements an analysis of them using these new tools.

3.1.1 Features of evidence-related predicates

Constructions involving evidence-related experiential predicates state something about the evidence that their subject is disposed to produce via a certain sensory modality. And so (244a) can be paraphrased: 'The visual experience that Alfonse is disposed to produce (*simpliciter*) evidences that Alfonse is tired,' and (244b) can likewise be paraphrased: 'The auditory evidence that the music is disposed to produce (*simpliciter*) evidences that the

music is Bulgarian.'

(244) a. Alfonse looks tired.

b. The music sounds Bulgarian.

In accounting for these constructions, a few notions will be helpful. First, reference needs to be made to the experience produced through the relevant sensory modality. The language has a way to do this, in virtue of certain relational nouns homophonous with the sensory verbs: these again include *look*, *sound*, *feel*, *smell*, and *taste*. And so expressions like the following refer in some way to the experience that a stimulus is disposed to produce.

(245) a. the look of Alfonse

b. the sound of the music

These nouns can be taken to relate stimuli to their looks, their sounds, etc.: and so for instance, '$look''(w)(x)(y)$' can be written to mean, 'y is the look of x at w,' and so on for each other relational noun and sensory modality. If these nouns simply denote that relation (246a), their internal arguments saturate the stimulus (246b), and the definite article has an ordinary Strawsonian denotation (246c), then the denotation of an expression like (245a) is as in (246d).

(246) a. $[\![\text{look}]\!]^w = \lambda x_e.\lambda y_e.look''(w)(x)(y)$

b. $[\![\text{look [of Alfonse]}]\!]^w = [\![\text{look}]\!]^w([\![\text{Alfonse}]\!]^w)$
$= \lambda y_e.look''(w)(a)(y)$

c. $[\![\text{the}]\!]^w = \lambda P_{et}.\iota x[P(x)]$

d. $[\![\text{the [look [of Alfonse]]}]\!]^w = [\![\text{the}]\!]^w([\![\text{look [of Alfonse]}]\!]^w)$
$= \iota x[look''(w)(a)(x)]$

And so *look of Alfonse* denotes the property of being a look of Alfonse, and the definite description denotes Alfonse's look – the analogous holds of the other four sensory modalities, such that e.g. *the sound of the music* denotes $\iota x[sound''(w)(\iota y[music'(w)(y)])(x)]$.

This notion of an individual's look, sound, etc. then needs to be characterized in a way appropriate to the stimulation of experience through the sensory modality. Roughly, an individual's look (at a world) is then its disposition to produce visual experience (at that world). Further, this disposition needs to be interpreted causally, via a specific sensory channel: thus Alfonse's look is the visual experience that Alfonse has a causal disposition to produce via contact with the visual receptors, and so on for each other modality. Therefore where an individual's look, sound, etc. is the stimulus of an experience, this is the same as to say that that individual is the causal source of the experience in question, through the relevant modality.

(247) $stim''(w)(\iota x[look''(w)(y)(z)])(z) \wedge exp''(w)(u)(z)$ only if:
y is the causal source of u's experiencing z at w,
via contact of y with the visual receptors of u
(and *mutatis mutandis* for the other sensory modalities).

The reason for this causal restriction is that evidence-related experiential predicates, unlike their deverbal psych counterparts that do not encode for a sensory modality, require a certain causal contact between the stimulus and experiencer to obtain: Alfonse cannot look tired to Bethany unless Bethany has seen Alfonse, and this requires attributing certain causal powers to Alfonse via his look – and not only this, but there is also a canonical means by which to do this, so that Bethany typically has to direct her eyes at Alfonse in order for this to be true.[3]

A construction like (245a) then says that Alfonse's look is disposed to be the stimulus of a certain sort of experience, which is the same as to say that Alfonse himself is causally disposed to produce a certain sort of visual experience. This experience in turn is taken to evidence a certain proposition. viz. that Alfonse is tired. Also required, then, is a notion of an experience providing evidence for a certain proposition. This can be encoded in the

3. This is not to deny the use of metaphorical extensions of verbs like *look* that do not literally make reference to the sense receptors, e.g. when speaking of seeing in the mind's eye, or being aware of an individual more generally. The strict sensory reading is what is at issue here.

relation $evid''$, which relative to a world holds between an experience and a proposition, and is true just in case the former evidences the latter.

(248) $evid''(w)(\phi)(x)$ is defined only if x is an experience, and ϕ a proposition;
where defined, it is true iff x evidences ϕ at w.

This evidencing can be thought of as a modal relation, which holds just in case, at all worlds at which the relevant experience is had and is ideally reliable, the proposition holds. And so, where ρ_{REL} is an accessibility relation related worlds by a relation of ideal reliability of experiences:

(249) $w\rho_{\text{REL}}w'$ iff:

a. w' preserves all the experiential states obtaining at w';

b. given that these states hold, the evidence that the relevant experiences provide for the truth of propositions at w' is ideally reliable.

(250) $evid''(w)(\phi)(x)$ iff $\forall w' : w\rho_{\text{REL}}w'[\phi(w')]$

This means that the $evid''$ relation does not encode which propositions are true, based on which experiential states obtain, but only which propositions, so to speak, 'look like' or 'sound like,' etc. that they are true: if the evidence provided by such experiences were ideally reliable, these propositions would be true, but that ideal reliability may or may not actually hold at w, the world at which the relevant experience evidences the proposition. As will be shown in §3.1.2 below, this means e.g. that *Alfonse looks tired* will be true just in case the visual experience that Alfonse is actually causally disposed to produce would indicate that he is tired, if these experiences were ideally reliable (as they very well may be, in which case that Alfonse looks tired is a reliable indication that he is tired).

Which experiences evidence which propositions is an enormously complicated question that exceeds not only the bounds of this work, but likely linguistics generally. It is doubtful that there is any linguistically encoded recipe for construing experiences as indicating the truth of propositions under ideally reliable conditions, and so speakers are apt to vary as

to which experiences they take to evidence which propositions, in a way that cannot be adjudicated by facts about the world, including facts about linguistic norms. And so the *evid″* relation is underdetermined in the same way that the preorder on qualities (cf. §2.1.1) and generic quantification (cf. §2.3.1 and §2.3.2) are (and so this relation is itself a source of subjectivity, unrelated in principle to the sort engendered by experiential semantics: cf. §5.3.2).

With these basic components in place, a denotation for sensory verbs can be constructed in terms of the notion of the look, sound, etc. of an individual being disposed to stimulate an experience that evidences a proposition. But before moving on, a feature of these sensory predicates are worth flagging.

The following is concerned solely with the interpretation of these verbs restricted to a specific sensory modality, such as vision or hearing. This is crucial, because as noted in Heycock (1994), these verbs have 'evidential extensions,' according to which the source of evidence that they pertain to is not sensory, but of a broader category that may roughly map onto types of evidence encoded by evidential markers cross-linguistically. Thus *look* in many constructions pertains to evidence generally, as opposed to visual experience, and *sound* often pertains to reportative evidence, rather than to auditory experience.[4]

With *look*, this extension is most easily seen where the verb takes a finite clause: in (251a), there may be any evidential indication whatsoever that Alfonse is going to be fired, not necessarily pertaining to visual experience. With *sound*, however, and as Lasersohn (1995) in particular notes, evidential extensions are rampant, even in the constructions covered here: and so (251b) easily takes the interpretation that according to available reportative evidence, Alfonse is nice, which has nothing to do in principle with actually hearing Alfonse.

(251) a. It looks like Alfonse is going to be fired.

4. The verbs *feel*, *smell*, and *taste* may also have evidential extensions, but they are somewhat harder to find. *Feel* in particular may relate to some sort of intuitive or 'endopathic' evidence, as in *It feels like something is wrong*, or *The house feels empty*. *Smell* and *taste* are much harder to push away from their strict sensory readings, although *smell* sometimes apparently relates to 'ambient' evidence, as in *This smells like a setup*.

b. Alfonse sounds nice.

These are mentioned only to be set aside as confounds to the genuinely sensory interpretations: their exact relation to these is unclear.[5] With that said, the following section will put the above notions to use in composing evidence-related experiential predicates.

3.1.2 Composing evidence-related predicates

With the notions discussed in §3.1.1 in place, a denotation for the sensory verb *look* can be given. *Look* composes with a property, denoted by the adjective, and an individual, denoted by the subject, to return true just in case the maximal experience stimulated by the look of that individual evidences the proposition that said individual has said property.

(252) $[\![\text{look}]\!]^w = \lambda P_{s,et}.\lambda x_e.evid''(w)(\lambda w'_s.P(w')(x))(\sigma y[stim''(w)(\iota z[look''(w)(x)(z)])(y)])$

This denotation also provides a schematic for sensory verbs generally: all that has to be done to characterize the denotation for *sound*, for example, is to swap '$look''$' with '$sound''$,' and so on. To compose (244a), first the verb takes the adjective, resulting in *look tired*.

(253) $[\![\text{look tired}]\!]^w = [\![\text{look}]\!]^w(\lambda w_s.[\![\text{tired}]\!]^w)$
$= \lambda x_e.evid''(w)(\lambda w'_s.tired'(w')(x))(\sigma y[stim''(w)(\iota z[look''(w)(x)(z)])(y)])$

This denotes the property true of an individual just in case the visual experience that it stimulates evidences that it is tired. There is no flavor of genericity to this predicate, and so this represents its exocentric reading, where the stimulus has actually produced a

5. It is tempting to think that the verbs involved are actually identical, and that there is just looseness in the relevant base of experiences that the verb can encode, with e.g. auditory evidence bleeding into reportative evidence due to some deeper cognitive association between the two. However, phenomena about the requirements on the stimulus make such an equation doubtful: as discussed in §3.1.2, sensory readings of these verbs typically require their subject to stimulate the relevant experience, while the evidential readings do not: thus, for *Alfonse sounds nice*, Alfonse himself does *not* need to be the stimulus of experience constituting reportative evidence (and indeed he generally is not), but where the sentence is read in an evidence-based way, e.g. if the way Alfonse sounds evidences that he is nice, then he must be, and the speaker of such a sentence is typically committed to actually hearing Alfonse. Cf. fn. 10.

visual experience in some (generally contextually relevant) experiencer, actually evidencing the proposition: cf. §2.3.2, and (238b).

The operator GEN then composes with the exocentric predicate, exactly as with deverbal psych predicates, to yield its 'genuinely' experiential, evaluative counterpart.[6] The denotation remains the same as in §2.3.2 (cf. (201)). As usual, the generic quantification targets worlds ideal to the actualization of the property over which it scopes. Where this pertained to a specific experiential kind for deverbal psych predicates, the relevant worlds for these sensory properties will be those ideal for the production of experience through the relevant sensory modality. The relevant accessibility relations can be called VIS, AUD, HAP, OLF, and GUS, to signal worlds optimal for the production of visual, auditory, haptic, olfactory, and gustatory experience respectively. For the intension P of *look tired*, $\rho_P = \rho_{\text{VIS}}$, and so on.

(254) a. $[\![\text{GEN}]\!]^w = \lambda P_{s,et}.\lambda x_e.\forall w' : w\rho_P w'[P(w')(x)]$

b. $[\![\text{GEN [look tired]}]\!]^w = [\![\text{GEN}]\!]^w(\lambda w_s.[\![\text{look tired}]\!]^w)$
$= \forall w' : w\rho_{\text{VIS}}w'$
$[evid''(w')(\lambda w''_s.tired'(w'')(x))(\sigma y[stim''(w')(\iota z[look''(w')(x)(z)])(y)])]$

This is the desired experiential predicate: it is true of an individual where, in all worlds ideal for the production of visual experience, that individual's causal disposition to produce visual experience stimulates an experience evidencing that the relevant individual is tired. In other words, it is true of an individual disposed to produce visual evidence that it is tired. Finally, *Alfonse looks tired* is derived by composing the experiential predicate with the subject, as follows.

6. It's taken for granted here that these evidence-based constructions are non-gradable, and hence there is no place for degree morphology here. Superificial evidence makes this plausible, like the absence of certain degree modifiers that can be used with psych verbs (cf. §2.2.2): compare, *The movie frightens Alfonse a lot* to ?*Alfonse looks tired a lot*, which can only have a frequency reading. This may ultimately be a simplification, and these constructions may ultimately be gradable in some way: the relevant gradable interpretation would be one that targets the degree to which the experience acts as evidence for the proposition, as is possibly in the case in *Alfonse looks tired more than Mary does*. This is to be distinguished from the reading on which degree morphology targets the adjective, and so this means not that there is visual evidence that Mary is more tired than Alfonse, but that of the visual evidence stimulated by Alfonse and Mary, it is that stimulated by Alfonse that more strongly evidences tiredness.

(255) $[\![\text{Alfonse [GEN [looks tired]]}]\!]^w = [\![\text{GEN [looks tired]}]\!]^w([\![\text{Alfonse}]\!]^w)$

$= \forall w' : w\rho_{\text{VIS}}w'$

$[evid''(w')(\lambda w''_s.tired'(w'')(a))(\sigma y[stim''(w')(\iota z[look''(w')(a)(z)])(y)])]$

And so, *Alfonse looks tired* is true on its experiential, or evaluative, reading just in case in all worlds ideally normal with respect to the production of visual experience, the maximal experience that Alfonse's look (or causal disposition to produce visual experience) stimulates evidences that Alfonse is tired. The reader can confirm that similar denotations will result for other adjectives and sensory verbs: and so, *sound Bulgarian* on its generic reading (cf. 244b) is true of an individual just in case the auditory experience that individual is causally disposed to stimulate evidences that said individual is Bulgarian, and so on.

There are two important things to note about this denotation for *look*. First, it does not have the argument structure of a classical raising verb, which only scopes over a proposition. Instead, it has two arguments, and composes with both a property and a stimulus. This mirrors the treatment of Germanic copy-raising in Asudeh & Toivonen (2012: 358, ex. 107): due to the clausal structure that the verb imposes, *look* composes first with its adjectival predicate, resulting in a stative verb *look tired*, which then composes with its subject in the ordinary way.

There are two reasons for this: the first is morphological, and the second is semantic. First, the morphology of English can apparently target the stative verb and its adjectival complement without its subject for derivation: this happens with the adjectival suffix *-ing*, which results in adjectives that incorporate the sensory verb and adjective with which it composes, but not a subject, like *tired-looking*. If this adjective is to be derived compositionally, then *look tired* must form a constituent capable of taking morphology, and this is not possible on a classical raising verb treatment, where *look* would scope over the small clause *Alfonse tired*: *-ing* would then only be able to suffix to a complex containing *look* that included the subject.[7]

7. This raises the question of how the morphology works to get the morpheme-ordering right. The result

The morphological relation between *look tired* and *tired-looking* is clear, and the two are at least close to synonymous, such that (256a) and (256b) are paraphrases of each other.

(256) a. Alfonse looks tired.

b. Alfonse is tired-looking.

On the present treatment, this is unproblematic, treating *-ing* in this construction as a semantically vacuous adjectival morpheme (257a).[8]

(257) a. $[\![\text{-ing}]\!]^w = \lambda P_{et}.P$

b. $[\![\textsc{gen}\ [\text{tired-looking}]]\!]^w = [\![\textsc{gen}]\!]^w(\lambda w_s.[[\![\text{-ing}]\!]^w([\![\text{look}]\!]^w(\lambda w_s.[\![\text{tired}]\!]^w))])$

$= \lambda x_e.\forall w' : w\rho_{\text{VIS}}w'$

$[evid''(w')(\lambda w_s''.tired'(w'')(x)(\sigma y[stim''(w')(\iota z[look''(w')(x)(z)])(y)])]$

$= 254b$

Tired-looking is then simply an adjectival variant of *look tired*, and its stimulus is not yet saturated.[9]

---

is correct if the sensory verb, just like *-ing*, behaves suffixally, or as a head-mover:

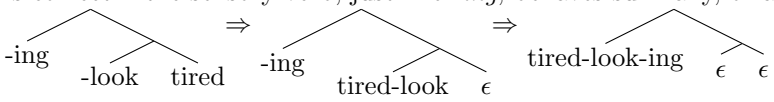

But this raises the question of why the verb cannot behave this way except in construction with *-ing*, and so **Alfonse tired-looks* is out. The text won't answer this question, although here it could be suggested vaguely that the morphological requirements on *-ing* prevent it from infixing, as it would if *look* were not treated as a head-mover, yielding **looking tired* as an adjective. The requirement that *-ing* suffix to a single morphonological unit then somehow drives the fact that *look* must be treated sufixally as well, as it can't in its normal function as a verb.

8. It might be puzzling to treat this *-ing* in this way: shouldn't it be assimilated to the *-ing* discussed in §2.3.1, which composes with object-experiencer psych verbs to form experiential predicates? If this were so, then complexes like *look tired* would have to have genuine semantic experiencer arguments, like psych verbs, to be desaturated by *-ing* for the suffix to function type-theoretically. But the experiencer phrases, like *to Alfonse*, capable of attaching to these verbs (see below) syntactically do not look like genuine verbal arguments, in their optionality or their ability to be dislocated, and for the reasons given in Chapter 1, there is independent semantic evidence is that these constructions do not have experiencer arguments anyway. The *-ing* here therefore is distinct, and in fact patterns with the homophonous suffix that attaches to transitive verbs that have incorporated nouns, as in *window-washing*. Note that this noun-incorporating *-ing* actually cannot be used with object-experiencer psych verbs to produce an adjective with the relevant experiential meaning: something like *tourist-frightening* can only have the causal sort of reading mentioned in §2.2.1. This can be seen from the fact that psych verbs without a causal reading available cannot take this *-ing* with an incorporated noun, as witnessed by **student-interesting*.

9. This account leaves a substantive problem dangling, viz. that these predicates combine less readily

Second, the subject of a verbal complex like *looks tired* plays two semantic roles: it functions both as the subject of predication in the visually evidenced proposition, and as the stimulus of the visual experience that evidences this proposition (what Asudeh & Toivonen *ibid.* term the 'psource' of this evidence). And so *Alfonse looks tired* is not true just in case there is simply disposed to be visual evidence that Alfonse is tired: this visual evidence must come from the causal dispositions of Alfonse himself. This means that in order for the experiential predicate to obtain of Alfonse, it must be that Alfonse is disposed to stimulate a certain kind of visual experience, and that the resulting experiential implications (cf. §3.1 above) involve the experiencer having seen Alfonse specifically.

And so for instance, it is insufficient to see that Alfonse's work desk, while he is absent from it, has a pillow on it, to utter *Alfonse looks tired* felicitously. In this scenario, there is visual evidence that Alfonse is tired, but it comes from seeing his work environment, and not him. This is dependent on the clausal structure of the sensory verb: when *look* is used with an expletive subject and a *like*-complementizer, this requirement does not hold. Compare the acquaintance inferences and direct experience presuppositions in the following examples: in each case, the (b)-sentences *are* felicitous when only seeing Alfonse's work station, or anything whatsoever that provides evidence that Alfonse is tired, whether it is Alfonse himself or not.[10]

(258) a. Alfonse looks tired.

with overt experiencer phrases than it would predict: ?*tired-looking to Alfonse* is distinctly odder than *looks tired to Alfonse*. Given that the two have the same denotation (ignoring the contribution of tense), there is no explanation for this difference at present. However, note that overt experiencers composing with e.g. *look tired* must be verbal modifiers, while those that would compose with *tired-looking* would be adjectival. It stands to reason that if there is a subtle difference between the semantics of these two sorts of overt experiencers, such that the adjectival counterpart cannot compose with sensory verbs of this sort that do not make reference to experiential kinds, then the contrast would be explained: the denotation is good with the verbal modifier, but not the adjectival. But what this difference might be will be passed over here, in an attempt to emphasize the common features of these experience phrases cross-categorially.

10. At least as far back as Rogers (1972), it's been noted that the requirement that the subject be the stimulus also looks to hold for copy-raising constructions: *Alfonse looks like he is tired* apparently signals that Alfonse himself is the stimulus of the relevant visual evidence. Landau (2011: §2.3) has some reservations here, though they look to stem from interpretations of these verbs that are evidentially extended, and so not strictly speaking sensory: cf. the caveat at the end of §3.1.1.

$\rightsquigarrow$ The speaker has seen Alfonse.

b. It looks like Alfonse is tired.

$\not\rightsquigarrow$ The speaker has seen Alfonse.

(259) a. Alfonse looks tired to me.

$\hookrightarrow$ The speaker has seen Alfonse.

b. It looks like Alfonse is tired to me.

$\not\hookrightarrow$ The speaker has seen Alfonse.

In the case where *look* is capable of taking an expletive subject, both the surface syntax and the semantics signal that the verb is scoping over an entire clause, viz. *Alfonse is tired*, and taking that as a single argument, consonant with a classical treatment of raising verbs. For such a use of *look*, a denotation with a simplified argument structure, that existentially quantifies over stimuli, is appropriate: the interpretation is only that something or other has provided visual evidence for the embedded proposition.

(260) $[\![\text{look}]\!]^w = \lambda\phi_{st}.\exists x[evid''(w)(\phi)(\sigma y[stim''(w)(\iota z[look''(w)(x)(z)])(y)])]$

Composing *It looks like Alfonse is tired*, assuming for simplicity a clausal counterpart of GEN, then yields the following result.

(261) $[\![\text{GEN [It looks [like [Alfonse [is tired]]]]}]\!]^w$

$= [\![\text{GEN}]\!]^w(\lambda w_s.[\![\text{look}]\!]^w(\lambda w'_s.[[\![\text{tired}]\!]^{w'}([\![\text{Alfonse}]\!]^{w'})]))$

$= \forall w' : w\rho_{\text{VIS}}w'$

$[\exists x[evid''(w')(\lambda w''_s.tired'(w')(a))(\sigma y[stim''(w')(\iota z[look''(w')(x)(z)])(y)])]]$

This is true just in case something or other is disposed to produce visual evidence that Alfonse is tired. From this denotation, no reference is made to Alfonse as a perceptual source. But for just this reason, the denotation in 260 is inappropriate for the construction under consideration here: if one is to make the subject necessarily the stimulus for the proposition in question, then that stimulus must be 'recovered' from the denotation of the clause over

which the verb scopes. But there is no clear way to recover Alfonse from a constituent that denotes the proposition that Alfonse is tired: this is just a function from worlds to truth values.[11] If, on the other hand, Alfonse is a genuine semantic argument of the verb, his needing to perform a double role, as both subject of the evidenced proposition and stimulus for the evidence, is easy to capture, as above.

The second thing to note about this denotation for *look* is that *look P* does not imply *P* of what it's predicated of. That Alfonse looks tired does not imply that he is tired, independent of a supplementary premise that the visual evidence in question is reliable: this follows from the modal semantics outlined in §3.1.1, that Alfonse's tiredness is relegated to those worlds ideally reliable with respect to this evidence, and which says nothing about whether the actual world is among these. And so any combination of evidence matching or mismatching with the evidenced property is possible: Alfonse can both look and be tired, look tired but not be tired, be tired without looking tired, and neither look nor be tired. And speakers can use verbs of perception to implicate either that the predicate taken by the sensory verb actually holds of the stimulus in the actual world or not, depending on whether these independent premises about the reliability of evidence are taken to be true or not. Where the evidence is taken to be positively misleading, *Alfonse looks tired* may in fact implicate that he is not tired – the reading is, that he 'merely' looks so.

Speakers can then commit to belief in these experiential matters, and dispute about them, at two different levels of description: first, it may be denied that Alfonse even looks tired, that is, the experiential predicate itself may fail to hold. This says something either about the experiences that Alfonse is in fact disposed to produce, or about what propositions those experiences in fact evidence. But the second level at which a dispute can arise is with respect to how reliable the evidence in question is, i.e. with respect to whether or not the

11. Of course, one could try to gerrymander a way to fish Alfonse out of this function: perhaps by identifying that individual such that, for some property *P*, in all worlds at which the proposition is true, that individual has *P*. This is an extremely risky move, however, since it will fail to uniquely recover Alfonse if his being tired entails that absolutely any other individual has any property at all (including trivial properties, like being self-identical, which somehow have to be ruled out).

experiential property holding says anything about the embedded predicate, e.g. whether Alfonse is actually tired.

With the above said, it can be shown how these experiential predicates also compose with overt experiencer PPs, in the same way that deverbal psych predicates do (cf. §2.3.3). But to do this, one change needs to be made to the notion of experiential alternatives, as introduced in that section. There, experiential alternatives were defined in terms of degrees of experiential kinds (cf. (221)). Since these sensory verbs do not encode amounts of experiential kinds on their evidence-based readings, the notion needs to be generalized: an individual's experiential alternatives are all those worlds in which stimuli are ideally disposed to stimulate precisely those experiences, of any sort, that they actually stimulate in the anchor-world, in the individual *qua* experiencer.

This can be expressed using the accessibility relation ρ_ϵ, which relates worlds according to the ideally normal conditions under which an individual produces experience of any kind: cf. (181), which invokes the same sort of accessibility relation, only anchored to a specific experiential kind.

(262) $w\rho_\epsilon w'$ iff:

a. w' preserves all intrinsic properties holding at w of potential stimuli relevant for their production of experience of any sort;

b. w''s conditions, given these properties, are ideally normal for the production of experience of any kind by potential stimuli.

Given this, an individual's experiential alternatives can be redefined as in (263a). This is a generalization of the original notion of experiential alternatives provided in §2.3.3. It entails the old definition as a condition of experiential alternatives (263b): if at all an individual's experiential alternatives, stimuli are disposed to produce *simpliciter* exactly those experiences that they actually produce in the anchor-individual at the anchor-world, then it follows that the degrees of experiential kinds that these stimuli are disposed to produce

will match the degrees actually produced in the anchor-individual as well. Experiential alternatives will therefore function as before, with respect to experiential predicates sensitive to these kinds.

(263) a. $Exp_{x,w}$
$:= \{w' : \forall y[\iota z[\forall w'' : w'\rho_\epsilon w''[z = \sigma u[stim(w'')(y)(u)]]] = \sigma u[stim''(w)(y)(u) \wedge exp''(w)(x)(u)]]\}$

b. $\forall w' \in Exp_{x,w}, \delta, k, y[\delta'_k(w')(y) = \delta''_k(w)(x)(y)]$

With this redefinition in mind, the old denotation for *to*, which relativizes the intension of a property to an individual's experiential alternatives (cf. (1.3.1)) can be preserved. It can then be taken to apply to the sensory predicates under consideration here, acting as a verbal as opposed to an adjectival modifier, but performing the same semantic function.

(264) $[\![\text{to}]\!]^w = \lambda x_e.\lambda P_{s,et}.\lambda y_e.\forall w' \in Exp_{x,w}[P(w')(y)]$

This yields the desired results, as can be shown for the case of *Alfonse looks tired to Bethany*, taking it that all those worlds ideally normal for the production of experience of any sort are also worlds ideally normal for the production of visual experience.

(265) a. $[\![\text{to Bethany}]\!]^w = [\![\text{to}]\!]^w([\![\text{Bethany}]\!]^w)$
$= \lambda P_{s,et}.\lambda y_e.\forall w' \in Exp_{b,w}[P(w')(y)]$

b. $[\![\textsc{gen}\ [\text{looks tired}]]\!]^w = [\![\textsc{gen}]\!]^w(\lambda w_s.[\![\text{look}]\!]^w(\lambda w'_s.[\![\text{tired}]\!]^{w'}))$
$= \lambda x_e.\forall w' : w\rho_{\text{VIS}}w'$
$[evid''(w')(\lambda w''_s.tired'(w'')(x))(\sigma u[stim''(w')(\iota z[look''(w')(x)(z)])(u)])]$

c. $[\![[\textsc{gen}\ [\text{looks tired}]]\ [\text{to Bethany}]]\!]^w$
$= \lambda y_e.\forall w' \in Exp_{b,w}[\forall w'' : w'\rho_{\text{VIS}}w''$
$[evid''(w'')(\lambda w'''_s.tired'(w''')(y))(\sigma u[stim''(w'')(\iota z[look''(w'')(y)(z)])(u)])]]$
$= \lambda y_e.evid''(w)(\lambda w'_s.tired'(w')(y))$
$(\sigma u[stim''(w)(\iota z[look''(w)(y)(z)])(u) \wedge exp''(w)(b)(u)])$

d. $[\![$Alfonse [GEN [[looks tired] [to Bethany]]]$]\!]^w$

$= [\![$[GEN [looks tired]] [to Bethany]$]\!]^w([\![$Alfonse$]\!]^w)$

$= evid''(w)(\lambda w'_s.tired'(w')(a))$

$(\sigma u[stim''(w)(\iota z[look''(w)(a)(z)])(u) \wedge exp''(w)(b)(u)])$

The result is true just in case the maximal experience that the look of Alfonse actually stimulates in Bethany evidences that Alfonse is tired, which is what's wanted. In order for this denotation to be defined, Alfonse's causal disposition to produce visual experience through the visual receptors must have stimulated an experience in Bethany – this means that Bethany must have seen Alfonse. Further, her visual experience of Alfonse must have a certain evidence-oriented character: it must be that, in those worlds in which such evidence is reliable, Alfonse is tired.

Thus, the experiencer phrase composes correctly, and preserves the relevant experiential presuppositions and truth conditions.The former, of course, will project in all of the expected environments, due to the definedness conditions of the σ-term, such that e.g. *Alfonse does not look tired to Bethany* has a denotation defined only if Bethany has seen Alfonse, and true just in case the visual experience thereby produced does *not* evidence that Alfonse is tired, as the interested reader can confirm.

3.2 Hybrid experiential predicates

Hybrid experiential predicates are those experiential predicates that simultaneously encode sensitivity to a sensory modality and an experiential kind. They can be composed phrasally, again by means of a sensory verb and an adjective, in which case the former encodes the sensory modality and the latter encodes the experiential kind: thus, *sound funny* is true of an individual whose causal disposition to produce auditory experience is itself disposed to stimulate humor. §3.2.1 treats hybrid predicates phrasally composed in this way. These predicates also occur more idiosyncratically, via certain morphological mechanisms, or just

simplex in the lexicon, but in these cases their meanings pattern like their more explicitly compositional counterparts. §3.2.2 treats these latter predicates, with particular attention to *tasty*.

3.2.1 Phrasally composed hybrid predicates

When sensory verbs compose with experiential adjectives, they sometimes yield special readings, resulting in experiential predicates that doubly specify the sort of experience a stimulus is disposed to stimulate, both in terms of sensory modality and experiential kind. So for example, *looks disgusting* may be true of an individual not if the visual experience it is disposed to produce evidences that it is disgusting, but rather if it has a disgusting look; and *sound funny* may be true of an individual, not when it is disposed to produce auditory evidence that it is sunny, but rather if it makes a funny sound.

These hybrid readings, therefore, predicate something of the causal disposition of the stimulus itself: that an individual that sounds funny is one that has a funny sound means that the adjective *funny* is effectively being predicated of the modality encoded by *sound* itself.

It is tempting to think that these hybrid readings are simply the result of the semantics of the adjective, which itself already lies at the core of an experiential predicate: one would then like an explanation for why these non-evidence-related readings arise specifically when the adjective is itself experiential in this way. However, this cannot be right, and the two constructions appear to be genuinely different, as with experiential predicates, both of them often survive. Thus (266) has two paraphrases, with (266a) being the evidence-related one, and (266b) being the hybrid one.

(266) Alfonse sounds funny.

a. ↔ The auditory experience that Alfonse is disposed to produce evidences that Alfonse is funny.

b. ↔ The auditory experience that Alfonse is disposed to produce is funny.

The first reading is applicable where the speaker is looking for a stand-up comedian, and listening to Alfonse, thinks he might be worth a shot. The second reading implies that the sounds that Alfonse makes themselves have a funny quality to them – perhaps because he has an odd voice, for example.

In the hybrid constructions, then, *sound* must have a slightly different denotation, and in particular one that composes with an adjective and simply applies the property that adjective denotes to the sensory modality in question, as stimulated by the subject. A denotation for this construal of *sound* would be as follows.

(267) $[\![\text{sound}]\!]^w = \lambda P_{et}.\lambda x_e.P(w)(\iota y[sound''(w)(x)(y)])$

Composition with the core of an experiential predicate that denotes P will then convert the result into a gradable predicate that encodes an experiential kind: for *sound funny*, this is humor, just as with *funny*, only now the stimulus is restricted to its causal disposition to produce auditory experience.

(268) a. $[\![\text{funny}]\!]^w = \lambda x_e.\delta_{\text{HUMOR}}(\sigma y[stim''(w)(x)(y)])$

b. $[\![\text{sound funny}]\!]^w = [\![\text{sound}]\!]^w([\![\text{funny}]\!]^w)$
$= \lambda x_e.\delta_{\text{HUMOR}}(\sigma y[stim''(w)(\iota z[stim''(w)(x)(z)])(y)])$

This is the core of the experiential predicate, which maps an individual to the degree of humor instantiated by the maximal experience that its sound actually stimulates. Other than the restriction to a specific sensory modality, this then has the semantics of the core of an experiential predicate as usual, and the 'genuinely' experiential predicate can then be derived as usual by adding degree morphology and genericity. So on composing with POS and GEN, this is the result.

(269) $[\![\text{GEN [POS [sound funny]]}]\!]^w = [\![\text{GEN}]\!]^w(\lambda w_s.[\![\text{POS}]\!]^w(\lambda w'_s.[\![\text{sound funny}]\!]^{w'}))$
$= \lambda x_e.\forall w' : w\rho_{\text{AUD}}w'[0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w')(\iota y[sound''(w')(x)(y)])]$

And so *sound funny* is true of an individual whose causal disposition to produce auditory experience is disposed to stimulate experience of humor. Finally composing with the subject gets the desired result for *Alfonse sounds funny*.

(270) $[\![\text{Alfonse [GEN [POS [sounds funny]]]}]\!]^w = [\![\text{GEN [POS [sounds funny]]}]\!]^w([\![\text{Alfonse}]\!]^w)$
$= \forall w' : w\rho_{\text{AUD}}w'[0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w')(\iota y[sound''(w')(a)(y)])]$

This is true where Alfonse is causally disposed to produce sound that itself stimulates humor, which is the correct interpretation. A similar result will obtain for any experiential predicate and sensory verb, so that e.g. *Alfonse looks disgusting* will be true just in case Alfonse is causally disposed to produce visual experience that itself stimulates disgust. The application of overt experiencers will further apply as normal, as in §3.1.2, yielding direct experience presuppositions that require both the relevant sensory contact with the stimulus, and the stimulation of the experiential kind in question, as the reader can confirm: thus, *Alfonse sounds funny to Bethany* presupposes that Bethany has heard Alfonse, and requires for its truth that the auditory experience she has of him stimulates humorous experience in her.

3.2.2 Idiosyncratic hybrid predicates: tasty

Certain experiential predicates combine a sensitivity to a sensory modality and sensitivity to an experiential kind more idiosyncratically, in virtue of restricted morphological composition, or in virtue of just being hard-coded that way in the lexicon. The result is that these predicates behave in roughly the same way as those composed in §3.2.1: they are true of individuals disposed to produce a certain experiential kind by the medium of a certain sensory modality. This section demonstrates this for the experiential predicate with which Chapter 1 began: *tasty*.

Several authors have either stated (e.g. MacFarlane 2014: §7.1) or implied through the metalanguage of their lexical entries (e.g. Stephenson 2007a) that *tasty* is synonymous with

taste good, on the experiential as opposed to evidence-based reading, and this equivalence in itself helps demonstrate that the adjective is decomposable into two components, much like its verbal counterpart: one that governs the sensory modality, and one that governs the quality of experience. The semantics of *good* is a thorny issue beyond the scope of this work, so a more harmless substitute can be used for illustration of this synonymy: *tasty* appears to be synonymous with *taste pleasant*, or equivalently using the verb-incorporating adjectives mentioned in §3.1.2, *pleasant-tasting*.[12]

Given the above, the composition of *taste pleasant* is now unproblematic: it is simply the result of taking the sensory verb *taste* on its experiential reading with the experiential predicate *pleasant*, which itself can be taken to be composed via the object-experiencer psych verb *please*, which is sensitive to the experiential kind PLEASURE, and the adjectival *-ant*, which performs the same semantic role as *-ing* (cf. §2.3.1).

(271) a. $[\![\sqrt{\text{PLEASURE}}]\!]^w = \lambda x_e.pleasure'(x)$

b. $[\![\text{-}\varnothing]\!]^w = \lambda P_{s,et}.\lambda x_e.\lambda y_e.\delta_{\kappa'(P)}(\sigma z[stim''(w)(y)(z) \wedge exp''(w)(x)(z)])$

c. $[\![\text{please}]\!]^w = [\![\text{-}\varnothing]\!]^w(\lambda w_s.[\![\sqrt{\text{PLEASURE}}]\!]^w)$
$= \lambda x_e.\lambda y_e.\delta_{\text{PLEASURE}}(\sigma z[stim''(w)(y)(z) \wedge exp''(w)(x)(z)])$

(272) a. $[\![\text{-ant}]\!]^w = \lambda R_{s,\langle e,ed\rangle}.\lambda x_e.\delta_{\kappa''(R)}(\sigma y[stim''(w)(x)(y)])$

b. $[\![\text{pleasasnt}]\!]^w = [\![\text{-ant}]\!]^w(\lambda w_s.[\![\text{please}]\!]^w)$
$= \lambda x_e.\delta_{\text{PLEASURE}}(\sigma y[stim''(w)(x)(y)])$

(273) $[\![\text{GEN [POS [taste pleasant]]}]\!]^w$
$= [\![\text{GEN}]\!]^w(\lambda w_s.[\![\text{POS}]\!]^w(\lambda w'_s.[\![\text{taste}]\!]^{w'}([\![\text{pleasant}]\!]^{w'})))$

12. It is a substantive issue if *tasty* really is synonymous with *taste good*, and if so, why *good*, which ordinarily is semantically not tied to any experiential kind specifically, takes on this pleasure-oriented reading in this construction. It may be that there is only one cognitively salient way for a taste to produce a good experience *simpliciter*, and *tasty* apparently cannot predicate goodness of a taste in any other sense, as when relativized to a goal in *This food tastes good for making people vomit* – cf. the impossible paraphrase ?*This food is tasty for making people vomit*. While in the latter case, the *for*-phrase may be ruled out for mundane syntactic reasons, it is simply not possible for *This food is tasty* alone ever to mean this either, or to mean anything other than that its taste is pleasant. With *This food tastes good*, perhaps things are less clear – and so *taste pleasant*, in addition to being theoretically safer, is probably empirically safer as a synonym anyway.

$= \lambda x_e.\forall w' : w\rho_{\text{GUS}}w'$

$[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta_{\text{PLEASURE}}(\sigma y[stim''(w')(\iota z[taste''(w')(x)(z)])(y)])]$

$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$

And so *taste pleasant* is true of an individual just in case under ideally normal conditions for the production of gustatory experience, that individual's taste, or causal disposition to produce gustatory experience, stimulates an experience that instantiates a non-zero degree of pleasure. Or, in short: it is true of individuals disposed to produce gustatory pleasure. Or, yet shorter: it is true of individuals that are tasty. This is what *tasty* means in its generic, positive form, and so here is the denotation for that adjective *qua* experiential predicate.

But *tasty*, while apparently denotationally equivalent to *taste pleasant*, is composed by different means: on the surface, it looks to be derived from *taste*, which encodes the sensory modality, and the adjectival suffix *-y*, elsewhere common in English, that forms the adjective and somehow encodes the experiential kind. It can be assumed that *taste* is just the relational noun noted in §3.1.1 above (274a).[13] The function of *-y*, applied most straightforwardly, would then be to encode the production of pleasure (274b).

(274) a. $[\![\text{taste}]\!]^w = \lambda x_e.\lambda y_e.taste''(w)(x)(y)$

b. $[\![\text{-y}]\!]^w = \lambda R_{e,et}.\lambda x_e.\delta_{\text{PLEASURE}}(\sigma z[stim''(w)(\iota y[R(z)(y)])(x)])$

This gets the right result, but there is something funny about it: it is implausible that English contains an adjectival morpheme with such a specific function and limited distribution (and that apparently only composes with *taste*). More plausibly, *-y* is related to the general adjectival morpheme that, like *-ful*, forms a predicate true of individuals that possess

13. While this assumption is easiest for present purposes, one might like to have cross-linguistic evidence that this is how the composition of such predicates works generally. In English, it is not clear from the surface morphology whether *taste* here is not instead the stative perceptual verb, or even an agentive verb, as in *Alfonse tasted the soup*. To do a broader morphological survey would take the present chapter too far afield, but even the most low-hanging data from familiar languages suggests that words synonymous with *tasty* do at least as a tendency get composed via an adjectival morpheme composing with a relational noun meaning *taste*. Cf. e.g. Korean *mas / masjoheun*, Spanish *sabor / sabroso*, or Tibetan *bro ba / bro ba chen po*. This is not to say that such adjectives are not derived by other means in other languages.

a certain amount of something. On such a view, *tasty* would be like *salty* or *watery*: the latter are true of individuals that contain a certain amount of salt or water, while the former is true of individuals that 'contain' a certain amount of taste (cf. Kennedy 2016: 122 ff). *Tasty*, in other words, is like *flavorful*, true of individuals that contain some amount of flavor.

At this point, this can only remain a vague allusion: a lot would need to be explained with respect to how these morphemes relate, and why it is that an adjective that ostensibly means 'has a taste' or 'has a certain amount of taste' yields a reading having to do with pleasure. There may ultimately be cognitive reasons for this: for some reason, having a taste relates to having a pleasant taste. This is seen in words like *flavorful*, which can mean either 'having a flavor,' or 'having a certain amount of flavor,' or 'having a pleasant flavor.' *Tasty* is odd, however, in requiring the pleasure-oriented interpretation absolutely: it is apparently impossible for it to mean anything but that the stimulus' taste is pleasant. That this is a peculiarity is shown by the fact that even what is apparently the same morpheme yields different results depending on the sensory modality: *smelly* can apparently mean 'having a smell,' or 'having a certain strength of smell,' as well as its hedonic interpretation, which is bewilderingly linked to unpleasant smells, rather than pleasant. Unless there is a huge profusion of hyper-specialized adjectival morphemes, this ought to receive explanation at some deeper level.[14]

But lacking that explanation, the denotation for *-y* in (274b) can be settled for as a descriptively adequate shortcut. The composition proceeds as follows, to yield the core of the predicate, which maps individuals to the degree of pleasure stimulated by their tastes.

(275) $[\![\text{tasty}]\!]^w = [\![\text{-y}]\!]^w([\![\text{taste}]\!]^w)$
$= \lambda x_e.\delta_{\text{PLEASURE}}(\sigma z[stim''(w)(\iota y[taste''(w)(x)(y)])(z)])$

14. Peter Lasersohn (p.c.) has suggested that the association between taste and positive experience, and between smell and negative experience, is cross-linguistically robust. Determining whether this is so, and what it might mean, would require more serious investigation into the cross-linguistic encoding of sensory modality, and probably straddles the border between linguistics and more general cognitive science: perhaps there is some cognitive reason that having a strength of smell is associated with having a bad smell (so that the sensory verb *smell* itself, without composing with an adjective, by default is quasi-synonymous with *smell bad*), and the reverse is true for taste.

The degree morphology and generic operator then work as usual, yielding the appropriate result for the experiential predicate.

(276) $[\![\text{GEN [POS tasty]}]\!]^w = [\![\text{GEN}]\!]^w(\lambda w_s.[[\![\text{POS}]\!]^w(\lambda w'_s.[\![\text{tasty}]\!]^{w'})])$
$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$
$= (273)$

Here then is *tasty*, as the present work has been implicitly talking about it since Chapter 1. It is true of individuals who are disposed, under ideally normal conditions of tasting, to causally produce pleasant gustatory experiences *simpliciter*, via contact with the taste receptors.[15]

Finally, there is a diagnostic that can help demonstrate that a certain predicate really does incorporate reference to a sensory modality analogously to a sensory verb. First, where such predicates are adjectival, like *tasty*, this means that they are themselves capable of composing with sensory verbs. If the sensory verb in question encodes the same sensory modality as the adjective, then the result is either oddity or redundancy. The former happens when there is morphological relatedness between the adjective and verb, as in ?*tastes tasty*, or equivalently, ?*tasty-tasting*: cf. e.g. ?*smells smelly*, or any composition involving a sensory verb with an adjective that has already incorporated that verb, like ?*looks tired-looking*.

Where the predicate encodes a sensory modality without this morphological relatedness, the result is instead that the adjectival predicate alone is synonymous with the result of composing it with the verb: thus, *taste delicious* is synonymous, on its experiential reading, with *delicious*. The same is not true for other predicates that one might expect encode a sensory modality: for instance, *look red* is not synonymous with *red*, which is evidence that color expressions do not encode visual modality in this way, and they do not license an experiential reading related to vision. This latter sort of redundancy follows if there is a

15. It's worth stressing again this canonical causal route through the taste receptors: if Alfonse is hooked up to a machine that directly stimulates his gustatory cortex to induce pleasant gustatory experiences in him, still he cannot rightly say on this basis that the machine is tasty. To find this out, he must lick the machine!

certain sort of collapse of causal dispositions to experience: for instance, the taste of a taste of an individual is just that individual's taste (277).

(277) $\iota y[taste''(w)(\iota z[taste''(w)(x)(z)])(y)] = \iota y[taste''(w)(x)(y)]$

Then assuming that *delicious* is synonymous with *tasty* (278) (ignoring its status as an extreme adjective), the denotation of *taste delicious* collapses into that of *delicious*.[16]

(278) $[\![\text{GEN [POS [delicious]]}]\!]^w = \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$

(279) $[\![\text{taste [GEN [POS delicious]]}]\!]^w = [\![\text{taste}]\!]^w([\![\text{GEN [POS delicious]}]\!]^w)$
$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(\iota z[taste''(w)(x)(z)])(y)])$
$\lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$
$=$ (278)

16. The stronger unacceptability of e.g. ?*taste tasty* is presumably due to morphological reasons, though these signal semantic redundancy as well, since if morphological material related to a lexical item like *taste* is included twice, then so is the corresponding semantic reference to taste as a sensory modality. And so to the extent that such predicates can be acceptable, they also receive interpretations synonymous with the adjective alone.

CHAPTER 4

DIRECT EVIDENCE AND COMMITMENT TO EXPERIENCE

This chapter explores how the semantics of experiential predicates, in concert with independently active evidential norms on speech acts and attitudes, causes speakers to commit to having had certain sorts of direct experience when using these predicates, and to commit to the truth of certain claims involving experiential predicates on the basis of their experiences. Previous chapters have shown a close link between experiential semantics and commitments to experience, especially through the use of overt experiencer phrases to encode presuppositions of direct experience (cf. §2.3.3). But the very same experiential semantics also gives rise to experiential commitments that cannot be assimilated to these, and has pragmatic consequences that cannot be reduced to factual reports of experiential states.

What follows demonstrates that experiential predicates are constitutively related to direct evidence, in the sense encoded by direct evidential markers cross-linguistically. In virtue of this relation between experiential semantics and direct evidence, mechanisms that independently enforce evidential commitments often require speakers using experiential predicates to commit to having direct evidence for the applicability of these predicates, which in turn requires these speakers to have had experiences of a sort lexically encoded by those predicates.

Speakers' commitment to the accuracy of their own direct evidence, in experiential matters as in others, then has consequences for how their experiences are taken to track experiential properties, and how speakers in a linguistic community construe experiences as involving accurate versus inaccurate perception. Because speakers can have conflicting experiences of the same stimuli, this in turn has consequences for how speakers evaluate the truth of experiential claims when the direct evidence for experiential properties is contradictory across speakers. Finally, these notions are shown to have a normative component, by which speakers can construe experiences as appropriate or inappropriate, praiseworthy or defective.

§4.1 deals with acquaintance inferences, the phenomenon introduced in §1.1.5, by which

assertions of experiential predications commit their speakers to having a certain sort of direct experience of the subject of predication. Previous approaches to these inferences are discussed and rejected, and it is shown that acquaintance inferences arise as a result of commitment to direct evidence as grounds for the assertion of an experiential predication. This is then shown to pattern with the cross-linguistic expression of direct evidentiality using direct evidential markers. It is then argued that acquaintance inferences pertain not only to assertion, but to belief in experiential propositions generally, which like assertions are governed by commitment based on an evidential source. Languages that do not encode direct evidentiality grammatically are shown to give rise to acquaintance inferences in virtue of the default evidential commitments required by belief, which require that speakers align their beliefs with their own experiences in a certain way. This means that acquaintance inferences arise as a result of direct evidential commitments, whether these are grammatically encoded or not, and that experiential predicates specifically enforce a direct evidentiality requirement in a way that non-experiential predicates do not.

§4.2 then addresses the pragmatics of using direct evidence to assess whether stimuli have experiential properties. It's argued that speakers are typically required to take their own experiences as accurately tracking experiential properties: what this amounts to is elucidated, and some cross-linguistic evidence is provided in support for it as an empirical generalization. Some pragmatic and normative consequences of this are then derived, especially in relation to how the experiential dispositions of speakers differ, and how speakers using the truth conditions of experiential predicates can cast certain experiential dispositions as normatively appropriate or defective.

4.1 Acquaintance inferences

As discussed in §1.1.5, acquaintance inferences are commitments, enforced by assertion of experiential predications, that the speaker makes to having had direct experience of a certain sort of the subject of predication. This phenomenon was introduced with the predicate *tasty*,

and it was noted how it extended to deverbal psych predicates in the introduction to Chapter 2, and to sensory verbs in §3.1. That a predicate gives rise to acquaintance inferences can be seen as of the core diagnostics signaling that it is in fact an experiential predicate.

The sort of acquaintance inference to which an assertion gives rise is encoded by the lexical or phrasal semantics of the experiential predicate, and each comes with a sensory and qualitative component: a *sensory acquaintance inference* pertains to the sensory channel through which a speaker is committed to having experienced the subject of predication, and a *qualitative acquaintance inference* pertains to what character that experience must have.

(280) The movie is frightening.

⇝ The speaker has watched the movie.

⇝ The speaker has experienced fear of the movie.

(281) Alfonse looks tired.

⇝ The speaker has seen Alfonse.

⇝ The speaker's visual experience of Alfonse evidences that Alfonse is tired.

(282) Alfonse's voice sounds funny.

⇝ The speaker has heard Alfonse's voice.

⇝ The speaker's auditory experience of Alfonse's voice is comedic.

It is entirely predictable from the semantics of each predicate what character these sensory and qualitative components have. Where the predicate encodes a sensory modality conventionally, as with complexes containing sensory verbs like *look* and *sound*, or adjectives like *tasty*, the speaker is committed to having had an experience through whatever sensory modality it encodes. Deverbal psych predicates like *frightening*, by contrast, only enforce a specific sensory acquaintance inference to the extent that it is supposed from world knowledge that experience of a certain stimulus occurs through a certain modality: thus, *frightening* itself carries no information about sensory modality, but (280) may imply that the speaker has watched the movie, since extra-linguistic knowledge suggests that this is the

canonical way of experiencing fear of a movie.

The qualitative component of the acquaintance inference is in turn controlled by what sort of experiential kind or evidence is is encoded by the predicate. For deverbal psych predicates like *frightening*, this follows straightforwardly from the kind associated with the corresponding psych verb (here, fear), and similarly for hybrid phrases like *sound funny*, where *sound* encodes the auditory sensory inference, and *funny* encodes that the relevant experiential kind is humor. Where there is a sensory verb and adjective with a purely epistemic interpretation, the quality of the experience pertains to what that experience provides evidence for, as with *looks tired*, which requires that the speaker have (visual) evidence that the subject is tired.

These inferences are in fact the same in their content as those produced by overt experiencer phrases, anchored to the speaker: cf. §1.1, §2.3.3, and §3.1.1 – the lexical and phrasal semantics already provided will thus take care of the content of the inferences themselves. What remains is to characterize in virtue of what mechanisms these inferences actually arise, and the goal of the present section is to show that this happens via speaker commitment to a direct evidential source for an assertion.

§4.1.1 is negative, arguing that acquaintance inferences are not speaker-oriented presuppositions on at-issue asserted content, nor are they constraints on speaker knowledge. §4.1.2 argues for the positive proposal that acquaintance inferences are to be assimilated to commitments to direct evidence, of the sort encoded cross-linguistically by direct evidential markers. §4.1.3 provides a compositional derivation of acquaintance inferences in languages that overtly mark for direct evidentiality, showing how acquaintance inferences arise as speaker-oriented not-at-issue contents in virtue of the speaker-oriented and not-at-issue nature of evidential commitments generally. §4.1.4 then shows how these inferences arise even in languages that do not mark for direct evidentiality, and which typically allow assertion to be evidentially neutral, and extends the notion of acquaintance inferences to attitude reports. Throughout, evidential effects are illustrated using Standard Tibetan, which

overtly marks for direct and indirect evidence in predication using adjectives, as an exemplar language.[1]

4.1.1 Against presuppositions on at-issue content and epistemic effects

Acquaintance inferences were discussed in §1.1.5, with respect to the problems they caused for a contextualist semantics. It was shown there, following Ninan (2014), that the experiential commitments that these inferences enforce are not reducible to speaker-oriented presuppositions on at-issue content, of the sort encoded by overt first-person experiencer phrases. The crucial data on which the distinction turned involved contexts in which acquaintance inferences, but not at-issue presuppositions, are voided, including beneath epistemic modals, beneath external negation, and in future contexts. All the experiential predicates examined in this work are entirely consistent on this score, and so for instance, on analogy with the examples in that section:

(283) a. The movie must be frightening to me.[2]

$\hookrightarrow_{proj}$ The speaker has watched the movie, and it must be that the speaker has experienced fear of it.

$\hookrightarrow_{acc}$ It must be that: the speaker has watched the movie, and experienced fear of it.

b. The movie must be frightening.

$\not\hookrightarrow$ The speaker must have watched the movie.

1. Thanks to Karma Ngodup for help with constructing and giving judgments on the Tibetan examples. Tibetan examples are given Roman transcription using the Wylie system. Key to the Tibetan glosses: ASST = assertive/neutral evidential, COP = copula, DIR = direct evidential, EGO = egophoric, ERG = ergative, FUT = future, IND = indirect evidential, LOC = locative, NEG = negation, PRES = present, Q = polar question. Glosses are approximate, and simplify subtleties in the tense/aspectual/evidential system where these are irrelevant.

2. Note the crucial difference between (283a) and (283b) with respect to the epistemic state of the speaker – the former, but not the latter, implies that the speaker is unfamiliar with her own experiences, perhaps because of forgetfulness. Anand & Korotkova (2018: 70, ex. 43) predict that sentences like (283a) have contradictory readings. This is a bad result – the epistemic state committed to in virtue of uttering this is slightly odd, but not impossible.

$\not\hookrightarrow$ The speaker must have experienced fear of the movie.

(284) a. Alfonse *doesn't* look tired to me, because I haven't seen him.

b. ?Alfonse *doesn't* look tired, because I haven't seen him.

(285) a. Alfonse's voice will sound funny to me.

$\hookrightarrow$ It will be that: the speaker hears Alfonse's voice.

$\hookrightarrow$ It will be that: the speaker's auditory experience of Alfonse's voice is humorous.

b. Alfonse's voice will sound funny.

$\not\hookrightarrow$ The speaker will hear Alfonse's voice.

$\not\hookrightarrow$ The speaker's auditory experience of Alfonse's voice will be humorous.

In fact the contexts in which acquaintance inferences are voided is far wider than this (cf. Anand & Korotkova 2018) – they disappear in any assertion made using a hedge or qualifier of some sort, which cannot be read as hedging or qualifying experiential presuppositions directed at the speaker.

(286) a. Apparently, licorice is tasty.

$\not\rightsquigarrow$ Apparently, the speaker has tasted licorice.

$\not\rightsquigarrow$ Apparently, the speaker's gustatory experience of licorice is pleasant.

b. Licorice is tasty, I hear.

$\not\rightsquigarrow$ The speaker hears that the speaker has tasted licorice.

$\not\rightsquigarrow$ The speaker hears that the speaker's gustatory experience of licorice is pleasant.

(287) a. Apparently, licorice is tasty to me.

$\rightsquigarrow$ Apparently, the speaker has tasted licorice.

$\rightsquigarrow$ Apparently, the speaker's gustatory experience of licorice is pleasant.

b. Licorice is tasty to me, I hear.

$\rightsquigarrow$ The speaker hears that the speaker has tasted licorice.

⇝ The speaker hears that the speaker's gustatory experience of licorice is pleasant.

These inferences are therefore tied to the commitments that a speaker makes or fails to make in virtue of assertion, but in such a way that they are not binding on the at-issue content of that assertion, or its presuppositions.

Ninan's (*ibid.*) suggestion is that acquaintance inferences arise not due to the semantic content of propositions composed using experiential predicates, but rather due to epistemic Moorean constraints on the speech act of asserting such propositions. As Hintikka (1962) notes, bare assertion of a proposition typically commits a speaker to knowing that proposition.[3] This can be seen from the fact that denying knowledge of a proposition just asserted is infelicitous, creating an epistemic counterpart to the classical Moore's Paradox.

(288) ?It's raining, but I don't know whether it's raining.[4]

The idea is that in asserting that it's raining, the speaker commits to knowing that it's raining, and so to deny this commitment via the second clause is bizarre. The commitment to knowledge is a general condition on assertion, and has nothing to do with the at-issue content of the first clause itself, which expresses the proposition that it is raining *simpliciter*, not that the speaker or anyone else knows this: it is the act of asserting this proposition bare that gives rise to the speaker-oriented knowledge effect.

In Gricean pragmatics, this fact is often taken to justify a normative general requirement on assertion, to the following effect (cf. Gazdar 1979: 46, ex. 24; Garrett 2001: 61, ex. 4; Ninan 2014: 302).

3. As should be clear from the comments on indirect evidentiality voiding commitment to the truth of an asserted proposition in §4.1.2 below, this observation does not hold universally across languages, but is dependent on the evidential system a language employs. However, it is a robust generalization where assertions are made in languages lacking evidential markers, without the use of hedges, modal expressions, etc.

4. Authors often use a propositional clause, rather than an embedded question, in demonstrating this effect, e.g. *?It's raining, but I don't know that it's raining.* A polar question is used here instead, since asserting *I don't know that it's raining* is already bizarre due to Moorean effects combined with the factivity of *know* (and for this reason, such an assertion tends to coerce a non-factive reading of the verb). Using a question allows the effect to be shown more clearly across two clauses without this complication.

(289) If x felicitously asserts p, then x knows p.

Ninan (2014: 303 ff.) notes that inferences arising from this epistemic constraint on assertion behave similarly to acquaintance inferences, and so proposes assimilating the two. For instance, just as seen for acquaintance inferences in §1.1.5, epistemic inferences can't be felicitously denied 288, they project out of negation (290a) (as shown by the infelicitous cancellation), and refuse cancellation even beneath external negation (290b).

(290) a. ?It isn't raining, but I don't know whether it's raining.

b. ?It *isn't* raining, because I don't know whether it is.

Possibly, they also are cancelled beneath those presuppositional holes and filters that allow speaker-oriented presuppositions of experience to project or be locally accommodated (shown by the possible felicity of the denial in (291)).[5]

(291) It must be raining, but I don't know whether it's raining.

For acquaintance inferences, the idea is then that a speaker must have the relevant direct experience to know the truth or falsity of propositions expressing experiential predications. If assertion commits a speaker to knowledge, and knowledge commits a speaker to experience of a certain sort, it follows that assertion commits a speaker to experience of a certain sort, and the acquaintance inferences are derived. So a principle like the following is additionally posited (cf. Ninan 2014: 302), where 'x' and 'y' are variables over individuals, and 'ϵ' is a variable over properties denoted by experiential predicates.

(292) If x knows whether y has ϵ, then x has the experience of y encoded by ϵ.

5. It's not entirely obvious that (291) is felicitous – it improves with stress on *know*, which suggests a strengthening of the standards for knowledge. And so the voiding of the epistemic inference here may be more contentious than the voiding of the acquaintance inference in (283b). Goodhue (2017) has explicitly argued that *must* is only felicitous in an assertion when its prejacent isn't known by the speaker, but this is at least debatable. Other classical presuppositional holes and filters fare better: compare *Is it raining? I don't know whether it's raining*, or *If it's raining, we shouldn't go outside, but I don't know whether it's raining*, which impeccably void the knowledge requirement. The predictive *will* neither voids the inference nor projects it into the future, making the parallels imperfect, but it's not clear to what extent this might be for independent reasons, and so it's not obvious to what extent this should matter to Ninan.

The derivation of the acquaintance inference is then as follows.

1. x felicitously asserts that y has ϵ.

2. By (289) and [1], x knows that y has ϵ.

3. By (292) and [2], x has the experience of y encoded by ϵ.

There is no recipe provided for what the 'experience of y encoded by ϵ' is, but one can substitute on a case-by-case basis, and so this schema instantiated for the assertion of *Licorice is tasty* runs as follows.[6]

1. The speaker felicitously asserts that licorice has the property of being tasty.

2. By (289) and [1], the speaker knows that licorice has the property of being tasty.

3. By (292) and [2], the speaker has tasted licorice, and the speaker's gustatory experience of licorice is pleasant.

Ninan's proposal makes a crucial prediction. Indirective evidential markers that can be independently shown to void acquaintance inferences for experiential predications ought to be infelicitous alongside claims to knowledge of the propositions expressed by those predications. This is because, if to know whether an individual has an experiential property requires the relevant direct experience of that individual, then a grammatical marker that commits a speaker to *not* having such experience must also commit the speaker to *not* knowing whether the individual has that property.

To the extent that speakers are sensitive to these restrictions, they ought to read claims to knowing whether an individual has an experiential property via indirect evidence as

6. Note that the derivation of the acquaintance inferences hinges on the speaker making the assertion felicitously. If one doesn't assume the assertion is felicitous, and adds as an independent premise that the speaker *hasn't* tasted licorice, etc., then the correct conclusion is that the speaker's assertion is *not* felicitous by the same principles.

infelicitous *tout court*, since according to the epistemic proposal, one cannot know such things by indirect evidence. But this prediction is not borne out.

The sorts of markers implying indirect evidence that the literature has primarily so far been concerned with, as in (283b), (286), and (287), make this prediction plausible, as each of these markers at least arguably does void the knowledge requirement. This was already seen for *must* generally in (291), but an argument could be made that it holds for experiential predications as well, along with *apparently* and slifted *I hear*, as shown by the following plausibly felicitous denials of knowledge.[7]

(293) a. Licorice must be tasty, but I don't know whether it is.

b. Apparently, licorice is tasty, but I don't know whether it is.

c. Licorice is tasty, I hear, but I don't know whether it is.

But there are periphrastic constructions that are troubling for the epistemic account: explicit claims to knowledge, alongside explicit claims to having (only) indirect evidence as the source of that knowledge don't look to be infelicitous in general.

(294) I know that the licorice is tasty...

a. because Alfonse made it.

b. because it's Finnish.

Both continuations in (294) provide an indirect justification for the speaker's knowledge (based not on having tasted the licorice, but on who made it, or its origin), and the sentence as a whole is felicitous even where it's clear to everyone that the speaker hasn't tasted the licorice in question. The epistemic account fails to predict this felicity: hearers should take the sentences to be bizarre, since the speaker is predicted not to be able to know that the licorice is tasty on these grounds.

7. As with (292), it's not entirely clear these denials are felicitous (without stress on *know*), at least not always. But the point can be granted, to proceed from the stronger argument.

But such constructions are slippery, and it would be better to have a more firmly grammaticized juxtaposition of knowledge with indirect evidence. This is provided by languages that have grammaticized indirect evidential markers. Since such markers often must accompany otherwise bare assertions, which typically require knowledge, this is a promising domain for testing the epistemic proposal. But a caveat needs to be made regarding the cross-linguistic relation between indirect evidential markers and knowledge, and commitment to truth more broadly.

In many languages, certain indirect evidentials loosen speaker commitment in assertion, causing the speaker not to commit to knowledge of an asserted proposition generally, and not just with experiential predicates. This is especially prevalent with reportative evidentials: a survey of the 'epistemic extensions' of these evidentials can be found in Aikhenvald (2004: §5.4.3). This lack of speaker commitment can be extreme: assertions using reportative evidentials in Cheyenne do not even commit the speaker to the truth of the asserted proposition (Murray 2017: §3.2.2), which can be seen from the fact that assertions made using a reportative evidential can be followed by an overt denial of the very same proposition using a direct evidential.[8]

(295) *É-hó'tȧheva-sėste Annie naa oha é-sáa-hó'tȧhévá-he-∅.*
3-win-RPT.3SG Annie and CNTR 3-not-win-NEG$_{AN}$-WTN
'Annie won, they say, but I witnessed that she didn't.'

[Cheyenne. Murray 2017: 76, ex. 3.8]

But this is not universally true of indirect evidentials cross-linguistically, and in many languages the use of inferential and even reportative evidentials does not lessen speaker commitment to the asserted proposition (cf. Faller 2002: §3.2.2 for brief discussion and cited examples). As Aikhenvald (2004: 3-8) has been at pains to explicate, evidentials encode the source of evidence for information, which in principle differs from epistemic matters having

8. Cf. Faller (2002: 194, ex. 158) for an illustration that the reportative *-si* in Cuzco Quechua vitiates commitment to the speaker's belief in the asserted proposition, so that Moore's Paradox doesn't arise when using it.

to do with the reliability of information, knowledge, and certainty. Not only are these notions logically distinct, but in many languages the encoding of evidential source, and the encoding of knowledge, come apart, and the use of an indirect evidential implies knowledge of the asserted claim just like the use of a direct evidential does.[9] It is these languages that are of interest for evaluating Ninan's (2014) epistemic position on acquaintance inferences.

Standard Tibetan is one such language.[10] It has an indirective copular form, *yod red*, which is used in predication with adjectives, when the speaker lacks direct evidence that the subject has the predicated property.[11] Where the adjective in question is an experiential predicate, like *bro ba chen po* 'tasty,'[12] use of *yod red* marks that the speaker has not had the experience of the subject relevant to an acquaintance inference.

(296) *kha lag 'di bro ba chen po yod red.*
food this tasty COP.IND
'This food is tasty.'

[Standard Tibetan]

⇝ The speaker has not tasted the food.

9. This is not to take a stand on the difficult question of whether, and to what extent, evidentiality and epistemic modality can be assimilated, nor is it to reject Matthewson et al.'s (2007) claim against Aikhenvald that the two are not fully distinct grammatical categories. It is an observable fact that in many languages, markers indicating indirectness of evidence still require commitment to knowledge, and this is the only sense in which it's claimed that epistemic matters and matters of evidential source vary independently.

10. The use of the term 'Standard Tibetan' is a bit of a cheat, and implies a host of problems. The term 'Lhasa Tibetan' is sometimes used to mean the standard variety of Tibetan spoken across the diaspora, but as Hill (2012: 391) notes, this term doesn't clearly distinguish between dialects of 'Central Tibetan' and the Tibetan spoken in the city of Lhasa itself, neither of which are really the same as Standard Tibetan (cf. Hill & Gawne 2017: 31). This inadequacy won't be addressed here.

11. *yod red* has a number of other functions: it is an existential-locative marker, and indicates alienable possession. But its evidential meaning is consistently indirective. For an introduction to Tibetan copular forms, including their evidential functions and how they behave with respect to predication using adjectives, see Denwood (1999: §9.2.3). The othrographic form of the indirective copula is not entirely consistent: one sometimes finds *yog red* or *yod pa red* attested in the literature. Some authors use the term 'assertive' (e.g. Tournadre & Dorje 2003) or 'factual' (e.g. DeLancey 2017) instead of 'indirect' to describe *yod red*, and DeLancey (*ibid.* §27.4.1) claims that 'factual' forms in general imply not indirect evidence, but unmarked evidential source. But *yod red* is unquestionably indirective in the sense relevant here: it typically requires lack of direct evidence for predications using adjectives.

12. *Bro ba chen po* is derived from the noun *bro ba* 'taste,' with the adjectival *chen po* 'big [positive],' similarly to English *tasty* (cf. §3.2.2), in what appears to be a cross-linguistically common pattern, and so it literally means 'having a big taste' (cf. Ch. 3, fn. 13).

(296) has a 'counter-acquaintance inference:' its assertion is felicitous where the speaker has not tasted the food, and so her gustatory experience of it has no particular quality.[13] But *yod red* in general does not void the speaker's commitment to the truth of the asserted proposition, and a speaker making an assertion using *yod red* is still committed to knowing what is asserted. The speaker simply must know the asserted proposition by indirect, rather than direct, means.[14]

This can be seen from the fact that explicit denials of knowledge following the use of *yod red* are in general infelicitous, as expected from the Hintikka-paradox illustrated in (288).

(297) *?bod la g.yag yod red, yin n'i bod la g.yag yod med shes gi*
Tibet.LOC yak COP.IND but Tibet.LOC yak be-not-be know PRES
med.
COP.NEG.EGO
'There are yaks in Tibet, but I don't know whether there are yaks in Tibet.'

The first clause of (297) asserts that there are yaks in Tibet, and commits the speaker to knowing this only by indirect evidence – she cannot have been to Tibet and seen yaks there, but must have read about them, heard about them from someone, etc. The second clause then denies that she knows what she just asserted with an indirect justification, and

13. Cf. Garrett (2001: 39, ex. 16), cited from Chang & Chang (1984), also using (an alternate form of) *bro ba*. Karma Ngodup (p.c.) informs me that (296) can also be used if the speaker has tasted the food some time in the past, but has forgotten this, or forgotten what it tastes like, or it has been so long that he is not sure whether his tastes or the flavor of the food have changed, etc. Presumably this is because tasting the food counts as direct evidence only insofar as the speaker can recall that experience for evidential purposes.

14. Garrett (2001: 37) makes this point forcefully about indirectivity in Tibetan: "Although indirect is associated with indirect forms of evidence, the knowledge it represents is still presented as certain knowledge. In other words, the speaker must be committed to what she is saying. Therefore, violations of Moore's paradox – which accounts for the pragmatic infelicity of %*John left, but I don't believe he did* – are disallowed here as with other evidentials; one cannot assert p and profess disbelief at the same time." And he continues (*ibid.* 38): "Similarly, one can't assert a proposition on one evidential grounds and deny it on another [contrast with the Cheyenne in (295) above] ... although indirect invokes a less direct form of evidential grounding than other evidentials, it does not imply that there is any lower degree of commitment to the proposition expressed." The requirement of knowledge by *yod red* is strong enough that it is mentioned spontaneously in pedagogical contexts: thus, the textbook Tournadre & Dorje (2003: 110) explains: "The assertive auxiliary verbs *re'* and *yo:re'* [i.e., *yod red*] specify that the information ... is considered by the speaker to be **certain** [emphasis in the original]. In general, the assertive implies that the speaker is letting it be known that while the assertion is certain, he or she has not personally witnessed it."

infelicity results.[15]

If using *yod red* for adjectival predication both commits speakers to knowledge of what they assert (297), and voids acquaintance inferences (296), then the fact that this copula is felicitous at all when used in experiential predications is surprising on the epistemic account: it ought to be odd, since the commitment to knowledge and the commitment to lack of direct experience (e.g. having tasted the food) ought to conflict. But they do not, and the assertion of (296) is felicitous.

This can be seen more directly by placing knowledge attributions, or disavowals of knowledge, directly alongside the experiential predication itself. A knowledge attribution alongside such a predication with *yod red* is felicitous, while a denial of knowledge is infelicitous (but on this latter point, see the caveat in fn. 15). Hence speakers are taken to know that an individual has an experiential property, even where the acquaintance inference is voided.

(298) *kha lag ʻdi bro ba chen po yod red. ngas shes gi yod.*
food this tasty COP.IND I.ERG know PRES COP.EGO
'This food is tasty. I know it.'

(299) ?*kha lag ʻdi bro ba chen po yod red, yin nʻi kha lag ʻdi bro ba chen po yod med*
food this tasty COP.IND but food this tasty be-not-be
shes gi med.
know PRES COP.NEG.EGO
'This food is tasty, but I don't know whether this food is tasty.'

The epistemic account predicts the reverse: (298), and not (299), ought to be infelicitous. In Standard Tibetan, the encoding of acquaintance inferences and the encoding of knowledge of asserted propositions therefore do not coincide. It is possible, as far as the language is

15. There may be some complications with this data. Karma Ngodup (p. c.) informs me that the denial of knowledge, while strange, may be passable to emphasize that the speaker has no direct evidence for the claim made in the first clause. He claims that the felicity drastically improves if one says instead that one does not know the fact in question *very well*, while asserting that one does not know it *simpliciter* intuitively conflicts with a claim to knowledge in the first clause. This suggests a gradability in the security of knowledge, which *yod red* might indeed weaken, but not to the extent that it voids commitment to knowledge altogether. This contrasts with the behavior of the direct evidential copula *ʻdug*, use of which accompanied by denial of knowledge is unequivocally bad. If this intuition is replicated across speakers, then one must allow that indirectivity can weaken claims to knowledge without voiding them. Nonetheless, even this weaker commitment to knowledge serves the purposes here.

concerned, to know that an individual has an experiential property by indirect evidence, i.e. by evidence other than having direct experiential contact with the individual. Insofar as this pattern replicates cross-linguistically – insofar as indirect markers that do not void commitment to knowledge generally are usable felicitously with experiential predications – this is evidence against the epistemic account as it stands. What follows proposes that acquaintance inferences are evidential effects, and not epistemic effects, and so tentatively rejects the epistemic proposal.

4.1.2 Acquaintance inferences as direct evidential effects

The hypothesis to be entertained here is that acquaintance inferences are direct evidential effects – that is, they result from speakers committing to a direct evidential source, in the sense of traditional evidential taxonomies such as Willett (1988) and Aikhenvald (2004), as the justification for the at-issue content of an assertion. The link between direct evidence and experiential predicates is suggested in passing by MacFarlane (2014: 3), and Anand & Korotkova (2018) suggest that commitment or lack of commitment to direct evidence is tied to the presence or absence of acquaintance inferences in assertions. There are three reasons for believing this claim, which pertain to the similarity between the distribution of acquaintance inferences and the distribution of overt markers of direct evidence.

First, in languages that have overt markers of direct evidence, those constructions that allow direct evidential marking give rise to acquaintance inferences just where this direct evidential marking appears. Anand & Korotkova (*ibid.*) note that the Turkish indirective marker *miş* voids acquaintance inferences. This can be seen from the fact that this marker makes denial of direct experience felicitous as a continuation from an experiential predication, while it otherwise isn't.

(300) a. #Durian güzel, ama hiç dene-me-di-m.
durian good but ever try-NEG-PST-1SG
'Durian is good, but I've never tried it.'

b. Durian güzel-miş, ama hiç dene-me-di-m.
durian good-IND but ever try-NEG-PST-1SG
'Durian is good, *I hear/infer*, but I've never tried it.'

[Turkish. Anand & Korotkova 2018: 58, ex. 8]

If acquaintance inferences arise as a result of commitment to direct evidence in assertion, then an explicit marker of lack of commitment to direct evidence unsurprisingly voids those inferences – and in fact, this was already observed to be the case in Standard Tibetan as well in §4.1.1 (cf. (296)).

The association runs in the other direction as well: not only does overtly marking for indirect evidence void acquaintance inferences, but overtly marking for direct evidence makes these inferences obligatory. To take the case of Standard Tibetan again, the copula *'dug* is a marker of direct evidence, and when used in predication with adjectives, it requires the speaker to have witnessed 'firsthand' that the subject has the property in question.[16] When *'dug* accompanies an experiential predication, the acquaintance inference indeed arises in the expected way (cf. Hill & Gawne 2017: 16, ex. 10c).

(301) *kha lag 'di bro ba chen po 'dug.*
food this tasty COP.DIR
'This food is tasty.'

⇝ The speaker has tasted the food.

⇝ The speaker's gustatory experience of the food is pleasant.

These inferences display the sorts of superficial behaviors expected from acquaintance inferences, in that they can't be overtly canceled (302), and they project out of negation (303).

(302) ?*kha lag 'di bro ba chen po 'dug, yin n'i ngas bro ba bltas med.*
food this tasty COP.DIR but I.ERG taste COP.NEG.EGO

16. *'dug* is sometimes referred to as 'testimonial' (e.g. Tournadre & Dorje 2003) or 'attestative' (e.g. Agha 1993: 112) in the Tibetan literature, but these terms refer to the same category as direct or perceptual evidentials in the literature more broadly. Hill (2012: §§2.1-2.2) briefly outlines the copular system, and the evidential function of the direct evidential. *'dug* isn't restricted to visual evidence (cf. Tournadre 2017: §4.4.1).

'This food is tasty, but I haven't tasted it.'

(303) *kha lag 'di bro ba chen po mi 'dug*
food this tasty NEG COP.DIR
'This food isn't tasty.'

⇝ The speaker has tasted the food.

⇝ The speaker's gustatory experience of the food is not pleasant.

The pattern is therefore that where a construction allows a semantically enforced distinction between commitment to direct evidence and commitment to indirect evidence as the basis of an assertion, acquaintance inferences arise with experiential predication just in case that construction is marked for direct evidential source.

Second, those environments where acquaintance inferences are voided, even in languages that lack evidential markers, tend to be exactly those where, cross-linguistically, overt direct evidential marking is impossible, even in languages that have such markers. Direct evidentials are restricted in their distribution, and cannot occur in a variety of constructions, especially modal constructions (for complications, examples, and cross-linguistic variation, cf. Aikhenvald ch. 8).[17] Their required absence is typical beneath epistemic modals, in the antecedents of conditionals, and in the future, exactly those sorts of environments where acquaintance inferences are also absent. (304c) and (305c) are to be imagined as uttered while the food in question is still being prepared.

(304) a. The licorice might be tasty.

b. If the licorice is tasty, I'll buy some.

c. The licorice will be tasty.

⇝̸ The speaker has tasted the licorice.

17. In fact Klecha (2014: §1.3) uses the absence of acquaintance inferences as a diagnostic test for modal constructions – this makes sense if the presence of acquaintance inferences is due to commitment to direct evidentiality, and commitment to modal claims on the basis of direct evidentiality is typically disallowed. Why it should be disallowed so robustly is a deeper question that won't be addressed in this work, but the fact that languages eschew encoding for direct perception of claims that depend on non-actual circumstances is hardly shocking.

$\not\rightsquigarrow$ The speaker's gustatory experience of the licorice is pleasant.

In Standard Tibetan, the direct evidential *'dug* cannot appear in these constructions, signaling that it is not possible for a speaker to commit to the claims made using them based on a direct evidential source.

(305) a. *kha lag 'di bro ba chen po* {*yod / *'dug*} *srid pa* {*red / *'dug*}.
good this tasty be may COP.ASST
'This food might be tasty.'

b. *kha lag 'di bro ba chen po* {*yod / *'dug*} *na ngas nyo gi yin.*
food this tasty be if I.ERG buy FUT COP.EGO.
'If this food is tasty, I'll buy it.'

c. *kha lag 'di bro ba chen po yong gi* {*red / *'dug*}.
food this tasty come FUT COP.ASST
'This food will be tasty.'

The generalization is that those constructions in which acquaintance inferences are systematically voided, even in languages without evidential markers, are constructions in which direct evidential markers are typically disallowed in languages that have them. If acquaintance inferences arise as the result of commitment to direct evidence, then it's expected that they would vanish in those constructions that cross-linguistically disallow commitment to direct evidence.

Third, acquaintance inferences, like commitments made with direct evidential markers generally (and with other sorts of evidentials as well), undergo *interrogative flip*. This means that the commitment to direct experience, and to direct evidence, are relocated from the speaker to the addressee in questions: cf. Aikhenvald (2004: §8.1.1.ii), Murray (2017: §2.3), and Sun (2017: §2.5.1) for descriptive expositions of this phenomenon for evidentials generally.[18] And so where an experiential predication is questioned, there is no acquaintance

18. This is not to say that all evidentials anchor to the addressee in all interrogative constructions in all languages: Maslova (2003), for example, reports that indirect evidentials in Yukhagir reflect the speaker's commitment to lack of direct evidence for the query in questions; cf. Sun (*ibid.*) for some similar examples. The author is unaware of any speaker-anchored direct evidential markers in questions, and their meaning

inference enforced on the speaker, but the addressee is expected to have the relevant direct experience: such a question is infelicitous otherwise (cf. fn. Ch. 1, 15).[19]

(306) Is licorice tasty?

⇝ The addressee has tasted licorice.

This is exactly what happens with overt direct evidential markers. To take the case of Standard Tibetan again, it is normal with respect to interrogative flip with direct evidentials: *'dug* marks a speaker presumption that the addressee has direct evidence verifying an answer to the question asked. This holds in general, but when experiential predications are questioned, the result is that the required direct evidence takes on the form of the by-now familiar acquaintance inference.[20]

(307) *kha lag 'di bro ba chen po 'dug gas?*
food this tasty COP.DIR Q

---

would perhaps be odd, given that the speaker would be asking about that for which he has direct evidence, clashing with the combination of commitment to accuracy of one's own direct evidence (cf. §4.2.1), and commitment to ignorance in sincere questions. Nevertheless, there is no *a priori* reason to rule out that they may exist. The tentative typological consensus at present is that interrogative flip is the 'normal' case for evidentials, and so what is primarily to be accounted for.

19. This is yet another problem for the sort of contextualism addressed in Chapter 1: there would need to be some mechanism shifting the experiencer parameter away from the speaker in questions, or else the resulting speaker-oriented meaning, required to capture acquaintance inferences in assertions, yields the utterly wrong meaning e.g. of *Is licorice tasty to me?* This is not to say that the contextualist could not come up with some technical device to accomplish this – perhaps questions are quasi-monstrous contexts for experiencer variables, shifting their value to the addressee – but there are two problems with such an approach, even if it could be made to work. First, questions do *not* shift the values of contextually fixed implicit arguments generally, and in particular do not display addressee-anchoring for non-experiencer parameters, meaning that the mechanism would have to be a special one proposed for experiencers, and is at risk of being *ad hoc*. Secondly, even if it were made to work, the resulting meaning would be wrong: *Is licorice tasty?* simply does not mean *Is licorice tasty to you?* The former, but not the latter, asks as to whether licorice is tasty *simpliciter*, based on the addressee's direct evidence. Thus if the addressee says *yes*, and the questioner tastes licorice and doesn't like it, she is within her rights to say the addressee was wrong in the first case, but not in the second (so long as the *addressee* in fact likes the taste of licorice).

20. Note that the acquaintance inference does not extend to the addressee's having liked the taste of the food, unlike with its speaker-oriented counterpart. This is to be expected if what is required is that the addressee have direct evidence for some answer to the question, not any particular one: both a positive negative answer to the question of whether the food is tasty require the sensory acquaintance inference, but the addressee's evidence may be constituted by either liking or not liking that taste, depending on the answer to be given.

'Is this food tasty?'

⇝ The addressee has tasted the food.

This is to be expected if acquaintance inferences arise as the result of commitment to direct evidence: those environments where the expectation of direct evidence anchors to the addressee are those in which the acquaintance inference anchors to the addressee as well. The following sections will therefore proceed with the hypothesis that acquaintance inferences are direct evidential effects in this sense.

The evidential account going forward preserves the original insight of the epistemic account laid aside in §4.1.1: it will be shown how acquaintance inferences arise due to commitments made in virtue of assertion that are not part of an asserted proposition's at-issue content, first by using direct evidential markers (in §4.1.3), and second through commitment to belief via assertion with a neutral evidential source (in §4.1.4). Just as Ninan (2014) proposes, acquaintance inferences in assertion therefore arise through Moorean commitments on the speaker – but these are commitments to direct evidence, rather than to knowledge.

4.1.3 Inferences with direct evidential markers

The account of experiential semantics presented so far, when combined with a simple and independently plausible semantics of direct evidentiality, and along with one further auxiliary assumption, allows acquaintance inferences to be compositionally derived when direct evidential markers are present.

Any treatment of direct evidential markers must capture two of their basic features: first, they conventionally commit speakers of assertions to having had a certain sort of perceptual evidence for the asserted proposition; and second, this commitment to 'direct evidence' does not form part of the at-issue content of what is asserted, and so has the sort of 'not-at-issue,' yet semantically-encoded and non-cancellable, force typically associated with a Gricean conventional implicature.[21]

21. Cf. Murray (2017: §§2.1-2.2) for diagnostics demonstrating this: in addition to being non-cancellable,

To introduce the notion of a speaker committing to having direct perceptual evidence for an asserted proposition, the treatment of indirect evidentials given in Izvorski (1997) can be adapted to their direct counterparts: evidentials are quasi-modal expressions, which require that a proposition be compatible with a base of evidence anchored to an individual, determined by the conventional content of the evidential marker. For direct evidentials, the 'modal base' in question has to do with the sorts of perceptions that an individual has had at the context of utterance. This can be represented by the *perceptual alternatives* of that individual, viz. the set of worlds compatible with what that individual has 'directly perceived.'[22]

(308) $Per_{x,w} := \{w' : w' \text{ is compatible with } x\text{'s direct perceptions at } w\}$

Some formal features of perceptual alternatives will need to be addressed, to make clear their relation to experiential semantics and the derivation of acquaintance inferences. But what it is *in general* for a world to be compatible with an individual's direct perceptions is a deeper question that will be left more or less open, according to the following intuitive gloss: all x's perceptual alternatives w' at w are such that $\phi(w')$, just in case x has directly perceived that ϕ at w. This is then constitutive of one's direct evidence *verifying* ϕ.

(309) x's direct evidence at w verifies ϕ iff:

$\forall w' \in Per_{x,w}[\phi(w')]$

And so a direct evidential conventionally commits the speaker of an assertion of ϕ to having direct evidence verifying ϕ (read non-factively, so that one can have direct evidence verifying what isn't actually so, e.g. in hallucinations), and this is the same as committing the speaker to having perceived that ϕ in the right 'direct' way.

---

commitments to direct evidentiality, like evidential commitments generally, tend to outscope all at-issue operators (and so resist embedding interpretations), and are not directly challengeable in conversation.

22. Izvorski's (*ibid.* 230) treatment of evidentials makes explicit use of the two-tiered modal semantics of Kratzer (1986): a modal base confines the relevant worlds to those in which the speaker has the relevant evidence, while an ordering source then ranks those worlds according to the speaker's presumed reliability of that evidence. This is what the text intuitively glosses, and it will reappear in slightly more explicit form, in the treatment of belief based on indirect evidence in §4.1.4.

$Per_{x,w}$ can be thought of as that set of worlds, all of which preserve the experiences that x has had at w (such that for all $w' \in Per_{x,w}$, if at w, x has experience z stimulated by stimulus y, then this is so at w' as well), and all of which further are compatible with these experiences being ideally reliable as to what they imply. The idea is that to directly perceive that ϕ is to have an experience stimulated by some stimulus, such that where this experience reliably tracks some feature of the world, that feature of the world obtaining entails that ϕ is true. Which experiences ideally justify which propositions being true is an enormously complicated matter that won't be addressed here – most likely, as with the preorder on qualities (cf. §2.1.1), the felicity conditions on generics (cf. §2.3.1), and especially the nature of $evid''$ (cf. §3.1.1), this relation between having an experience and having direct evidence for a proposition is underdetermined by linguistic convention and facts about the world, and so speakers will differ as to how to construe it (it therefore constitutes one of the many sites of 'weak subjectivity,' cf. §5.3.2).

To implement the second criterion, that the direct evidential enforce this commitment on the speaker at a not-at-issue level of content, a technical stopgap will be introduced, which distinguishes an expression's at-issue extension from its not-at-issue extension. Where an expression has an ordinary denotation, as with the denotations so far presented in this work, that denotation is assumed to be an expression's at-issue extension. If by contrast an expression has a denotation consisting of an ordered pair, the first member of the pair is taken to be its at-issue extension, while the second member is taken to be its not-at-issue extension. And so a direct evidential marker DIR composes with a proposition, and adds to it the not-at-issue content that the speaker's direct evidence verifies this proposition, as follows, where s_c is the speaker in the context of utterance c, and w_c the world of that context.[23]

23. The reason for introducing this device is that there is no standard method in the literature for compositionally representing the not-at-issue contribution of an evidential. Where a compositional treatment of evidentials is provided, as in Izvorski (1997: 226, esp. ex. 8) and Matthewson et al. (2007: 245, ex. 91), the evidential contribution is placed in some combination of the at-issue asserted and presupposed content, while accounts that place the evidential contribution as not-at-issue, as in Faller (2002) and Murray (2017), are non-compositional (*pace* Murray's insistence to the contrary regarding her own account). The treatment in

(310) $[\![\text{DIR}]\!]^{c,w} = \lambda\phi_{st}.\langle\phi(w), \forall w' \in Per_{s_c,w_c}[\phi(w')]\rangle$

Not-at-issue content represented by the second member of the ordered pair can be formally integrated in the pragmatics however one pleases, e.g. as updating the common ground directly, as not licensing propositional anaphora, etc. This is not intended as a serious general mechanism for dealing with not-at-issue content: for this, along with the standard two-dimensional account in Potts (2005), see Murray (2017) for a dynamic implementation of not-at-issue evidential content, using the framework of Bittner (2011). The denotation in (310) can be taken as that of a direct evidential such as the Tibetan *'dug*, illustrated in §4.1.2 above.[24] The derivation of *kha lag 'di bro ba chen po 'dug*, 'This food is tasty' (cf. (301)), where the experiential predicate is in the positive, generic form, is then as follows, taking *bro ba chen po* to be synonymous with English *tasty* (311d), and where d_c is the object demonstrated in the context of utterance c.[25]

(311) a. $[\![\text{'di}]\!]^{c,w} = \lambda P_{et}.\iota x[x = d_c \wedge P(x)]$

b. $[\![\text{kha lag}]\!]^{c,w} = \lambda x_e.food'(w)(x)$

c. $[\![[\text{kha lag}]\ \text{'di}]\!]^{c,w} = [\![\text{'di}]\!]^{c,w}([\![\text{kha lag}]\!]^{c,w})$
$= \iota x[x = d_c \wedge food'(w)(x)]$

d. $[\![\text{GEN}\ [\text{bro ba chen po}]]\!]^{c,w}$

---

the text is a compromise, to introduce the not-at-issue content compositionally in the most uncontroversial, theoretically-neutral way possible. The denotation in (310) hard-anchors the evidential base to the speaker, in the manner of first-person marking: as should be clear from the comments on interrogative flip in §4.1.2, this is not right, since evidentiality is typically not speaker-oriented in questions. This is left unresolved, since there is no extant compositional account of origo flip, and providing one would require getting into the semantics of illocutionary force, which is beyond the scope of this work. Murray (2017: ch. 5) implements origo flip, but non-compositionally, in a way not derivable from the interacting semantics of the evidential and interrogative force. Cf. also McCready (2007), for treatment of a related phenomenon involving monstrous operators.

24. This cannot be the final treatment of *'dug*: it, like most evidentials, fuses multiple semantic contributions, and acts as an existential-locative marker, which is also used for predication with adjectives and alienable possession. The Tibetan auxiliary paradigm makes clear, however, that this latter function is isolable from direct evidentiality (though cf. Garrett 2001 for a more complex story, in which Tibetan direct evidentiality with *'dug* is actually derived in part from its locative component).

25. The silent morpheme POS is not included in this derivation, since Tibetan *-po* is apparently the positive morpheme, and so positive degree morphology is already included in the citation form of the adjective.

$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$

e. $[\![$[[kha lag] 'di] [GEN [bro ba chen po]]$]\!]^{c,w}$

$= [\![$GEN [bro ba chen po]$]\!]^{c,w}([\![$[kha lag] 'di$]\!]^{c,w})$

$= 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)])$

f. $[\![$'dug$]\!]^{c,w} = \lambda\phi_{st}.\langle\phi(w), \lambda w' \in Per_{s_c,w_c}[\phi(w')]\rangle$

g. $[\![$[[[kha lag] 'di] [GEN [bro ba chen po]]] 'dug$]\!]^{c,w}$

$= [\![$'dug$]\!]^{c,w}(\lambda w_s.[\![$[[kha lag] 'di] [GEN [bro ba chen po]]$]\!]^{c,w})$

$= \langle 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)]),$

$\forall w' \in Per_{s_c,w_c}$

$[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota y[taste''(w')(\iota x[x = d_c \wedge food'(w')(x)])(y)])]\rangle$

And so the sentence's at-issue content is true just in case the food demonstrated in the context of utterance has a taste that is disposed to produce a non-zero degree of pleasure. Its not-at-issue content, to which the speaker commits in virtue of making an assertion using the sentence with a direct evidential, is true just in case the speaker has directly perceived this.

If, as suggested in §4.1.2, the acquaintance inference is a direct evidential effect, then it is in virtue of this contributed not-at-issue content that the acquaintance inference arises: and so, for one to have direct evidence that a stimulus produces gustatory pleasure (*simpliciter*) entails that one is the experiencer of gustatory pleasure stimulated by that stimulus. This result follows, if one's perceptual alternatives and experiential alternatives, in the sense of §2.3.3, are related in a certain way, such that all of one's perceptual alternatives must be compatible with one's experiential alternatives, i.e. if for each individual, the former are a subset of the latter.

(312) $Per_{x,w} \subseteq Exp_{x,w}$

What this means is that if a world is compatible with an individual's direct perceptions, then it is also compatible with that individual's experiences, in the special sense outlined in

§2.3.3 – in other words, to directly perceive that a stimulus is disposed to produce a certain experience is simply to have the relevant experience, as stimulated by that stimulus. The definition of experiential alternatives is in (313a), repeated from (263). Recall from there that this definition entails the special condition in (313b): in all an individual's experiential alternatives w', each stimulus is ideally disposed to produce that degree of the experiential kind that it actually produces in x at the anchor-world w.

(313) a. $Exp_{x,w}$
$:= \{w' : \forall y[\iota z[\forall w'' : w' \rho_\epsilon w''[z = \sigma u[stim''(w'')(y)(u)]]] = \sigma u[stim''(w)(y)(u) \wedge exp''(w)(x)(u)]]\}$

b. $\forall w' \in Exp_{x,w}, \delta, k, y[\delta'_k(w')(y) = \delta''_k(w)(x)(y)]$

If as concluded in (311), the extension of the not-at-issue content contributed by the direct evidential *'dug* is as in (314), then by (312), it follows from the not-at-issue content that the demonstrated food produces a non-zero degree of gustatory pleasure at some of the speaker's experiential alternatives; but since by (313b), these alternatives are uniform as to the degree of gustatory pleasure that the food produces, it follows as in (315) that it produces a non-zero degree of gustatory pleasure in all of the speaker's experiential alternatives.

(314) $\forall w' \in Per_{s_c,w_c}$
$[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota y[taste''(w')(\iota x[x = d_c \wedge food'(w')(x)])(y)])]$

(315) $\forall w' \in Exp_{s_c,w_c}$
$[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota y[taste''(w')(\iota x[x = d_c \wedge food'(w')(x)])(y)])]$
$= 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w_c)(s_c)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)])$

But this, again by (313b), is just to say that the degree of pleasure that the taste of the demonstrated food has actually stimulated in the speaker is non-zero, which entails the relevant acquaintance inference: for this to be so, the speaker must have tasted the food and liked its taste.

And so the acquaintance inference is derived via the direct evidential, as a not-at-issue proposition to which the speaker commits in virtue of making an assertion using an experiential predicate with a direct evidential source. This treatment carries over to other experiential predicates in the expected way, enforcing the relevant sensory and qualitative acquaintance inferences mentioned in §4.1.

It follows from this as well that acquaintance inferences project out of a number of environments, as typical of not-at-issue content, preserving the inference in the desired way. For instance, where the above sentence is negated, to form *kha lag 'di bro ba chen po mi 'dug*, 'This food is not tasty' (cf. (303)), it yields the following denotation, taking *mi* for simplicity to denote extensional negation (316a).

(316) a. $[\![\text{mi}]\!]^{c,w} = \lambda p_t.\neg p$

b. $[\![[[[\text{kha lag}]\ \text{'di}]\ [\textsc{gen}\ [\text{bro ba chen po}]]]\ \text{mi}]\!]^{c,w}$
$= [\![\text{mi}]\!]^{c,w}([\![[[[\text{kha lag}]\ \text{'di}]\ [\textsc{gen}\ [\text{bro ba chen po}]]]\!]^{c,w})$
$= 0_{\textsc{pleasure}} =_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)])$

c. $[\![[[[[\text{kha lag}]\ \text{'di}]\ [\textsc{gen}\ [\text{bro ba chen po}]]]\ \text{mi}]\ \text{'dug}]\!]^{c,w}$
$= [\![\text{'dug}]\!]^{c,w}([\![[[[\text{kha lag}]\ \text{'di}]\ [\textsc{gen}\ [\text{bro ba chen po}]]]\ \text{mi}]\!]^{c,w})$
$= \langle 0_{\textsc{pleasure}} =_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)]),$
$\forall w' \in Per_{s_c,w_c}$
$[0_{\textsc{pleasure}} =_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w')(\iota y[taste''(w')(\iota x[x = d_c \wedge food'(w')(x)])(y)])]\rangle$

And so the at-issue content is true just in case the taste of the demonstrated food is disposed to produce a zero degree of pleasure, and the not-at-issue content is true just in case the speaker has directly perceived this. By parity of reasoning with the above, this not-at-issue commitment is equivalent to (317).

(317) $0_{\textsc{pleasure}} =_{\textsc{pleasure}} \delta''_{\textsc{pleasure}}(w_c)(s_c)(\iota y[taste''(w)(\iota x[x = d_c \wedge food'(w)(x)])(y)])$

That is, the degree of pleasure that the taste of the demonstrated food has actually produced in the speaker is the zero-degree. And so what the speaker commits to is both

having tasted the food, and not having liked its taste: the sensory acquaintance inference is preserved, while the qualitative inference is negated, exactly as desired. Note also that these inferences escape even external negation, as noted in §4.1.1: this is because at-issue negation, whether external or not, composes with at-issue content, and cannot target not-at-issue content.

This derivation hinges on accepting (312), which amounts to the supposition that to directly perceive that a stimulus is disposed to stimulate a certain degree of an experiential kind is just to have the relevant experience, stimulated by that stimulus. While this may be an independently plausible principle, it would be nice to derive it somehow – here is a sketch of how this might be done.

Taking for granted that to directly perceive that an individual has a certain experiential disposition is just to witness that disposition being actualized, it follows from the above that this direct perception entails directly perceiving that an experiencer is in the relevant experiential state, as stimulated by that individual. However, there is independent reason to believe that, as direct evidentiality and experiential states are encoded linguistically, it is not possible to directly perceive that an experiencer other than oneself is in an experiential state. Data suggesting this have long been observed in the literature on Japanese predicates that require a subject to be in an experiential state, like *sabisii* 'lonely' (tracing back to before Kuno 1973), and similar observations have been made in Korean as well, for predicates like *ecirep* 'dizzy' (cf. Lee 2013).

In these languages, it is typically infelicitous to predicate an experiential state using a non-first person subject (pending some complications with tense: cf. Lee *ibid.*), so long as the speaker commits to a direct evidential source (318a). But an indirective evidential marker, like the Japanese *noda*, rescues the felicity (318b).[26]

26. The first datum traces to Kuno (1973). Bylinina's (2017) citation of it is given here, since she relates these observations to experiential predicates in the way relevant for this work, and her glosses are preserved. It's taken for granted that an unmarked evidential source in Japanese requires commitment to direct evidence, where the indirective marker is available for use. As expected, not only *noda*, but just about any device for signaling lack of direct evidence, including verbs of seeming and epistemic modals, also vitiate the first-

(318) a. {watasi-wa / *anata-wa / *kare-wa} sabisii desu
I-TOP you-TOP he-TOP lonely COP
'{I'm / *You're / *He's} lonely.'

b. Mary-wa sabisii noda
Mary-TOP lonely EVID
'Mary must be lonely.'

[Japanese: Bylinina 2017: 301, exx. 30a-b]

If these predicates require their subjects to be in experiential states, and if as proposed above this is the same role that experiential predicates play when overtly relativized to an experiencer, then this evidential restriction ought to carry over to experiential predicates overtly relativized to an experiencer other than the speaker. This is borne out in Japanese, where overt experiencers can be encoded using dative marking, with *oisii* 'tasty.' Here again, a direct evidential source with an overt experiencer other than the speaker is infelicitous, but an indirective marker removes the infelicity.

(319) a. {watasi-ni-(totte)-wa / *John-ni-(totte)-wa} kono keeki-wa oisii
I-DAT-TOTTE-TOP John-DAT-TOTTE-TOP this cake-TOP tasty
'This cake is tasty {to me / *to John}.'

b. John-ni-(totte)-wa kono keeki-wa oisii noda
John-DAT-TOTTE-TOP this cake-TOP tasty EVID
'This cake must be tasty to John.'

[Bylinina *ibid.* exx. 32a-b]

If this is a cross-linguistically robust pattern,[27] then it may be that in general the expe-

person requirement. The ban on direct evidentials can be seen more clearly in Korean, which has an overt direct evidential marker *-te* that shows similar restrictions. Again, cf. Lee (2013), and for a more general formal treatment of the evidential behavior of *-te*, cf. Lim & Lee (2012).

27. There are some potential counterexamples to this generalization: McLendon (2003: 104) reports that in Eastern Pomo, one can report that another individual is afraid, using a visual evidential, if one has visual evidence of their fear (say, from their bodily reactions). This is complicated by the fact that visual evidentials often pertain not only to the direct visual perception of states themselves, but also to the direct visual perception of that from which a state can be inferred. This corresponds to the difference flagged in Matthewson (m.s.) between two types of direct evidence – evidence can be direct either in virtue of being located with respect to an event in a certain way, or in virtue of being received through a certain sensory channel. If one can report another's state of fear using a visual evidential, this may suggest that

riential states of others cannot be witnessed in the way required for direct perception. This would be so, and the restrictions in Japanese and Korean would be derived, if to have direct evidence that an individual is in an experiential state is simply to be that individual in that experiential state. Direct evidence would then be *intrinsically autocentric*, and there would be a collapse between direct evidence and experience when it comes to experiential states.

(320) $\forall w' \in Per_{x,w}[stim''(w')(y)(z) \wedge exp''(w')(u)(z)]$ iff:
$x = u \wedge stim''(w)(y)(z) \wedge exp''(w)(u)(z)$

This principle states that where an individual's perceptual alternatives uniformly decide that an individual in an experiential state, this is just to say that the perceiver and the experiencer are identical, and that the perceiver is actually in said experiential state. From this the first-person requirements above follow: speakers can only commit to directly perceiving their own experiential states, and to do so is just to commit to being in that state.[28]

Putting all the above together, then, the derivation of the principle in (312) is as follows: (i) to directly perceive that a stimulus has an experiential disposition is to witness that disposition being actualized; (ii) to do this is just to directly perceive that an experiencer is in a certain experiential state, stimulated by the stimulus; (iii) to directly perceive this is just to be that experiencer, in that experiential state; so (iv) to directly perceive that a stimulus has an experiential disposition just is to be the experiencer of the relevant experiential state

this evidential is direct in the latter sense, but indirect in the former crucial sense, meaning that it signals direct perception of the 'results' of an event. One would need to distinguish these two types of evidentials by a prior diagnostic before testing the present generalization.

28. Why this principle in turn holds is a deeper question that won't be addressed here. But a good place to search for an explanation might be the way that experiential states are more fundamentally conceived in terms of experiential kinds being located in experiencers (cf. §2.1.1). Matthewson (m.s.) has suggested that direct evidentiality is to be broken up into three dimensions, one of which (and which likely corresponds to the notion of directness being assumed here) crucially depends on the location of the evidential anchor in relation to the witnessed event. If directly perceived events must be in some special locative relation to the perceiver, then it may be that, due to the relation between having an experience and that experience being located 'in' an experiencer that experiential states intrinsically encode, that it is just not possible for an experiential state to be located relevantly with respect to the evidential anchor unless it is 'inside' that anchor, which in turn requires just having the relevant experience. The inability to do this for anyone but oneself is then a (conceivably contingent) physiological-psychological barrier: so long as one cannot literally have another's experiences, the identity condition follows.

associated with that disposition, as stimulated by that stimulus.

A final note on the relation between experiential predicates and direct evidence bears mentioning. It follows from the above that experiential predicates have the curious property of being partially definable in terms of direct evidence: for any experiential predicate that denotes a property ϵ, for any x and w, x has ϵ at w just in case x is ideally disposed to stimulate direct evidence that it has ϵ at w. In other words, for a predicate to be experiential requires that it denote a property that an individual has just in case it is directly perceived to have it (in ideal conditions). This follows from the definition of the disposition to stimulate an experience *simpliciter* in §2.3.1: to be disposed in ideal conditions to stimulate such-and-such an experience is to actually stimulate that experience when the felicity conditions on the generic are met. This means that some individual is the experiencer of that stimulated experience, and that individual's having such an experience is identical to its having direct evidence that the stimulus has the corresponding experiential property.

4.1.4 The default to direct evidence

If the treatment in §4.1.3 is on the right track, then the question of how acquaintance inferences arise in languages like Standard Tibetan, in constructions that overtly encode direct evidentiality, is not mysterious. These inferences are enforced semantically, by the conventional content of evidentials, and the not-at-issue commitments that these force the speaker to make.[29] Similarly, where this sort of direct evidential marking is absent, the acquaintance inference vanishes, as when indirect evidentiality is overtly marked, in which case there is conventional commitment to only an indirect source of evidence, or beneath epistemic modals, futures, and in the antecedents of conditionals, in which case having direct evidence for the claims made by these constructions is impossible.[30]

29. Among these languages and constructions are assumed to be those that have an evidential paradigm featuring overt morphemes marking only indirect evidentiality of some sort, and which therefore require commitment to a direct evidential source in the unmarked case.

30. In fact where indirect evidentiality is overtly marked, a stronger conclusion might be drawn, that the speaker has *not* had the relevant direct experience verifying the asserted proposition, as in the 'counter-

But acquaintance inferences are also enforced in languages, like English, that do not have markers semantically encoding evidential source, and so typically do not conventionally commit speakers to direct sources of evidence in assertion. This is reflected in the fact that with non-experiential predicates, assertive predications in English can typically be made indifferently on the basis of either a direct or an indirect source of evidence. Yet acquaintance inferences in these languages behave just like the sorts of direct evidential effects derived in §4.1.3, as shown in §4.1.1 and §4.1.2: the inferences are identical to those that arise with direct evidential markers in languages that contain them, and they show the same sort of not-at-issue behavior.

It is as if, even in these languages, assertions using experiential predicates made use of a direct evidential marker. This is so for predicates like *tasty*, which give rise to acquaintance inferences exactly paralleling Tibetan *bro ba chen po*, and so which could be derived in exactly the same way, by composing English sentences like *Licorice is tasty* with a silent direct evidential DIR. The same can also be said of the other sorts of experiential predicates explored in this work, including deverbal psych predicates like *frightening* and sensory verbs like *look tired*. Their derivations, with such a silent direct evidential, would be as in 321a and 322a, and by parity of reasoning with §4.1.3, the resulting not-at-issue content would reduce to that in (321b) and (322b).

(321) a. $[\![\text{DIR [[the movie] [is [GEN [POS frightening]]]]}]\!]^{c,w}$
$= \langle 0_{\text{FEAR}} <_{\text{FEAR}} \delta'_{\text{FEAR}}(w)(\iota x[movie'(w)(x)]),$
$\forall w' \in Per_{s_c,w}[0_{\text{FEAR}} <_{\text{FEAR}} \delta'_{\text{FEAR}}(w)(\iota x[movie'(w)(x)])]\rangle$

acquaintance inferences' mentioned in §4.1.1. Since evidentials are typically arranged in a hierarchy (cf. e.g. Barnes 1984: §3, and §4.2.1 of the present work), where commitment to an evidential source ranked lower on the hierarchy implies a lack of evidence of a sort higher on the hierarchy, commitment to the truth of an asserted proposition based on indirect evidence will typically imply a lack of direct evidence on the basis of which to assess the truth of that proposition (since direct evidence is, apparently universally, ranked higher than indirect evidence on said hierarchies). Given the above, this will entail a lack of experience of the relevant sort. This counter-acquaintance inference will not hold when commitment to indirect evidence for a proposition is not typically paired with commitment to its truth. Placing a proposition beneath an epistemic modal or in the antecedent of a conditional also often leads to counter-acquaintance inferences for the embedded proposition, though the reasons for this are disputed: cf. von Fintel & Gilles (2010) and Giannakidou & Mari (2016) for treatments of this issue with epistemic modals.

b. $0_{\text{FEAR}} <_{\text{FEAR}} \delta''_{\text{FEAR}}(w_c)(s_c)(\iota x[movie'(w)(x)])$

(322) a. $[\![\text{DIR}\ [\text{Alfonse}\ [\text{GEN}\ [\text{looks tired}]]]]\!]^{c,w}$
$= \langle \forall w' : w \rho_{\text{VIS}} w'$
$[evid''(w')(\lambda w''_s.tired'(w'')(a))(\sigma x[stim''(w')(\iota y[look''(w')(a)(y)])(x)])]$,
$\forall w' \in Per_{s_c,w}[\forall w'' : w' \rho_{\text{VIS}} w''$
$[evid''(w'')(\lambda w'''_s.tired'(w''')(a))(\sigma x[stim''(w'')(\iota y[look''(w'')(a)(y)])(x)])]]\rangle$

b. $evid''(w)(\lambda w'_s.tired'(w')(a))$
$(\sigma x[stim''(w)(\iota y[look''(w)(a)(y)])(x) \wedge exp''(w_c)(s_c)(x)])$

And so the not-at-issue content to which the speaker committed by making an assertion using *The movie is frightening* on the basis of a direct evidential source would be that the movie has actually stimulated a non-zero degree of fear in the speaker. This in turn means that the speaker must have experienced the movie in whatever way is required for that stimulation (paradigmatically, watching it). The not-at-issue content of *Alfonse looks tired*, meanwhile, would be true just in case the maximal visual experience that Alfonse stimulates in the speaker evidences that Alfonse is tired, which in turn means that the speaker of this sentence commits to Alfonse being the causal source of a visual experience in the speaker, and so that the speaker has seen Alfonse.

Since these languages apparently do not have any direct evidential markers to play the role of DIR, the question remains why it is that there is a conventional commitment to a direct evidential source in assertions of experiential predications, but not with non-experiential predicates. To answer this, the nature of acquaintance inferences must first be correctly situated. While the previous literature, and the present work, have so far made acquaintance inferences primarily a matter of assertion using experiential predicates, these inferences apparently are not tied to assertion *per se*. Rather, acquaintance inferences are conventionally associated with *belief* in a proposition expressible using an experiential predication. Thus, when speakers are taken to believe such propositions, as an extremely robust default they are taken to do so on the basis of direct evidence. This can be seen from the fact that doxastic

reports typically preserve the acquaintance inference for the subject of the attitude.

(323) Alfonse thinks that licorice is tasty.

→ Alfonse has tasted licorice.

→ Alfonse's gustatory experience of licorice is pleasant.

If this is so, then following Ninan (2014), those acquaintance inferences that arise from assertion are plausibly due to Moorean effects, but resulting from the speaker's commitment to belief, rather than to knowledge. That is, if belief in such propositions conventionally implies a direct evidential source, then so does assertion of these propositions, which conventionally commits the speaker to belief in them. But as seen in §4.1.1, knowledge of these experiential propositions doesn't require a direct source of evidence in all cases – and still less does belief in them, since it is possible to believe the relevant propositions on the basis of indirect evidence.

The conventional link between belief and direct evidence is vitiated wherever there is a semantically-encoded requirement of an indirect evidential source. This is true not only where various semantic markers (such as conditionals, futures, and epistemic modals) make commitment to a direct source of evidence impossible, but also in doxastic reports, where there is lexical encoding of indirect evidence. The latter case occurs with the verb *believe*, which as Stephenson (2007b: §5.2) points out, implies a lack of direct evidence for the believed proposition on the part of the believer (cf. Lasersohn 2009: 366 for a similar comment, comparing *consider* to x*believe*).

(324) Alfonse believes that licorice is tasty.

↪ Alfonse has not tasted licorice.

(324) requires that Alfonse believe that licorice is tasty on the basis of some indirect evidential source, and so it is strange to report Alfonse's opinion about licorice's taste using *believe* if he has tasted it before. The plausible reason for this is that *believe* lexically encodes an argument slot for this indirect evidential source: it can appear overtly as an argument to

the verb (325a). The propositional and evidential source arguments can even appear overtly together (325b). Neither is possible with *think* (326).

(325) a. Alfonse believes me.

b. Alfonse believes me that licorice is tasty.

(326) a. *Alfonse thinks me.

b. *Alfonse thinks me that licorice is tasty.

And so the behavior of acquaintance inferences in attitude reports is the same as in assertions: overt markers of indirectness void them, showing that direct evidence is not required for belief in experiential propositions in all cases, but nevertheless in the absence of such a marker, the inference that the belief is based on direct evidence is a robust default.

Now this reframes the puzzle: why is there such a default commitment to direct evidence in believing these propositions? Answering this question requires examining the relation between experiential semantics, evidence, and belief, first to show how default or evidentially neutral belief requires direct evidence, and only for experiential properties, and second to show how belief rooted in an indirect evidential source voids this requirement. To begin, recall from §4.1.3 that an agent's perceptual alternatives are a subset of its experiential alternatives – that is, for any x, w: $Per_{x,w} \subseteq Exp_{x,w}$ (cf. (312)). In addition, one more stipulation can be made: an agent's doxastic alternatives, the worlds compatible with its beliefs, are a subset of its perceptual alternatives.

(327) $Dox_{x,w} \subseteq Per_{x,w}$

This amounts to the assumption that agents successfully take themselves to be *accurate perceivers* of a certain sort: they do not hold beliefs that do not align with their direct perceptions (their direct evidence), or in other words, they believe what they perceive. Justification for this assumption, as well as discussion of what it amounts to and when it is violated, is in §4.2.1. From these two principles, it follows that an agent's doxastic alternatives are a subset of its experiential alternatives – for any x, w: $Dox_{x,w} \subseteq Exp_{x,w}$.

And from this it follows that to commit to belief in a proposition expressed by an experiential predication, one must commit to having direct evidence verifying this proposition. To see this, suppose that Alfonse thinks that licorice is tasty at w, that is, that his doxastic alternatives at w uniformly verify this.

(328) $\forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

From the above, it follows that $Dox_{a,w} \subseteq Exp_{a,w}$, so as long as Alfonse's set of doxastic alternatives is non-empty, it follows that licorice is tasty in some of Alfonse's experiential alternatives (329a). But since by definition Alfonse's experiential alternatives are uniform with respect to the degree of pleasure that the taste of licorice stimulates, it follows that licorice is tasty in all of his experiential alternatives (329b).

(329) a. $\exists w' \in Exp_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

b. $\forall w' \in Exp_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

Since Alfonse's experiential alternatives are a superset of his perceptual alternatives, it then follows that licorice is tasty in all of his perceptual alternatives (330a), which is equivalent to Alfonse having direct evidence verifying that licorice is tasty (cf. (309)). Further, by the definition of experiential alternatives (cf. (313)), this means that the taste of licorice actually stimulates (at w the anchor-world) a non-zero degree of pleasure in Alfonse. And so Alfonse must have tasted licorice, and liked its taste, in virtue of thinking that licorice is tasty in this sense (330b).

(330) a. $\forall w' \in Per_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

b. $\delta''_{\text{PLEASURE}}(w)(a)(\iota x[taste''(w)(l)(x)])$

Then so long as *think* merely enforces a condition on its agent's doxastic alternatives in the usual way (331), a report such as *Alfonse thinks that licorice is tasty* entails that Alfonse has tasted licorice, and likes its taste (332).

(331) $[\![\text{think}]\!]^{c,w} = \lambda x_e.\lambda \phi_{st}.\forall w' \in Dox_{x,w}[\phi(w')]$

(332) $[\![\text{Alfonse [thinks [that [licorice [is [GEN [POS tasty]]]]]]]}]\!]^{c,w} =$

$[\![\text{think}]\!]^{c,w}(\lambda w_s.[\![\text{that [licorice [is [GEN [POS tasty]]]]]}]\!]^{c,w})([\![\text{Alfonse}]\!]^{c,w})$

$= \forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

$=$ (328) $\rightarrow$ (330b)

As the reader can confirm, the same will hold for believing the negation of such an experiential predication: *Alfonse thinks that licorice isn't tasty* entails that Alfonse has direct evidence that licorice is not tasty, i.e. that he has tasted licorice and not liked its taste. Further, where Alfonse has not tasted licorice at w, it follows by definition that for all $w' \in Exp_{a,w}$, $\delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w')(l)(x)])$ is undefined, since $\delta''_{\text{PLEASURE}}(w)(a)(\iota x[taste''(w)(l)(x)])$ is undefined. The degree of pleasure that the taste of licorice produces *simpliciter* is therefore undefined in all of Alfonse's doxastic alternatives, so long as he has never tasted licorice – and so insofar as he commits to being an accurate perceiver in the sense of (327), he cannot think that licorice is or is not tasty, in the sense that a *think*-report encodes, until he tastes it.

This direct evidence requirement does not hold for non-experiential predicates – this is because an agent's experiential alternatives by definition are only ever settled with respect to experiential matters, viz. which stimuli are disposed to produce which experiences. Non-experiential predicates do not encode matters of disposition to stimulate experience, and so an agent's experiential alternatives always leave the application of non-experiential properties undecided.

To take the example of *smart*, which encodes a non-experiential property: regardless of what experiences Alfonse has, of Bethany or otherwise, some of his experiential alternatives will verify that she is smart, and some will not. These experiences may in turn perceptually justify, or not, the truth of the proposition that Bethany is smart, pending on whether Alfonse has an experience that constitutes witnessing that she is or is not smart. But there is nothing stopping his perceptual alternatives from being heterogenous on the matter, since no uniformity is forced on them by his experiential alternatives. And so it's possible in turn

that his doxastic alternatives, still assumed to be a subset of his perceptual alternatives, uniformly verify or falsify that Bethany is smart, even where he has not witnessed either of these things – a *think*-report like *Alfonse thinks that Bethany is smart* is therefore compatible with Alfonse thinking that Bethany is smart due to direct evidence or not, as desired, and the same holds for every non-experiential predicate.

(333) $[\![\text{Alfonse [thinks [that [Bethany [is smart]]]]}]\!]^{c,w}$
$= [\![\text{think}]\!]^{c,w}(\lambda w_s.[\![\text{that Bethany is smart}]\!]^{c,w})([\![\text{Alfonse}]\!]^{c,w})$
$= \forall w' \in Dox_{a,w}[smart'(w')(b)]$

So it is that the commitment to direct evidence based on evidentially neutral belief commits an agent to having direct evidence for the belief only where predication of an experiential property is concerned – this follows from agents' commitment to their perceptions being accurate, coupled with the nature of experiential semantics.

But it is possible to believe that an individual has, or lacks, an experiential property on the basis of only indirect evidence. Because evidentially neutral belief, as encoded by *think*, requires a direct evidential source, this sort of indirect belief needs to be reported using a device that limits the justification for the belief to indirect evidence, like the verb *believe*. *Believe*-reports are not evidentially neutral, and this can be represented by having the verb compose with an indirect source argument (334), which then factors into the way the doxastic alternatives representing the belief are arranged (335).

(334) $[\![\text{believe}]\!]^{c,w} = \lambda x_e.\lambda\phi_{st}.\lambda y_e.\forall w' \in Dox_{y,x,w}[\phi(w')]$

(335) $Dox_{y,x,w} := \{w' : w'$ is compatible with x's beliefs at w justified by y as an indirect evidential source$\}$

In the spirit of Izvorski's (1997) treatment of indirect evidentials, this set of doxastic alternatives $Dox_{y,x,w}$ can again be spelt out by casting it in terms of a modal base and an ordering source: unlike with perceptual alternatives, the set is sensitive not to the experiences

that the believer has, but to the indirect evidence that the believer has received from y. The set is defined as follows.

(336) $w' \in Dox_{y,x,w}$ iff:

a. w' preserves all indirect evidence that x receives from y at w;

b. w' is maximally compatible with x's beliefs at w about the reliability of indirect evidence x receives from y at w.

(336a) reflects a modal base, viz. the set of all worlds preserving the indirect evidence that x receives from y at w. This means, for instance, that if at w Bethany has claimed that licorice is tasty, and Alfonse heard her, then the set of worlds characterized by (336a), where $y = b$, includes only those worlds in which Bethany claims that licorice is tasty and Alfonse heard her, and so on for all pieces of indirect evidence that she has produced.

$Dox_{y,x,w}$ is then a subset of the worlds in this modal base, characterized by the restriction in (336b), which itself corresponds to an ordering source. This source restricts the worlds in the set of doxastic alternatives to those in which all pieces of indirect evidence that y has produced for some ψ at w, and which x takes to be reliable indications of the truth of ψ at w, are actually reliable in this way (and so *a fortiori*, ψ is true in all these worlds).[31] In other words, $Dox_{y,x,w}$ is that set of worlds compatible with the indirect evidence that y has produced for x, and that x takes to be reliable, at w. Given this, the report *Alfonse believes Bethany that licorice is tasty* is as follows.

(337) $[\![$Alfonse [[believes Bethany] [that [[licorice [is [GEN [POS tasty]]]]]]]$]\!]^{c,w}$

$= [\![$believe$]\!]^{c,w}([\![$Bethany$]\!]^{c,w})([\![$that [licorice [is [GEN [POS tasty]]]]$]\!]^{c,w})([\![$Alfonse$]\!]^{c,w})$

31. And so more explicitly, where the relevant modal base = $\{w' : \phi(w')$, for all ϕ such that ϕ is a proposition true iff y has produced indirect evidence for x for some proposition $\psi\}$, and the ordering source = $\{\phi$: x takes ϕ to be true at w, and ϕ is true iff some indirect evidence that y has produced for x for some ψ reliably indicates that $\psi\}$, then the set of 'best' worlds derived from this modal base and ordering source is identical to $Dox_{y,x,w}$. Where x's commitments to the reliability of y's indirect evidence is consistent (which the text assumes), this set will just be that in which all the propositions in the ordering source are true (given the modal base), i.e. that set of worlds in which all of y's produced indirect evidence that x takes to be reliable at w, actually is reliable. Cf. the comments in Ch. 2, fn. 32 on implementing the Kratzerian modal semantics.

$= \forall w' \in Dox_{b,a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLESURE}}(w')(\iota x[taste''(w')(l)(x)])]$

$= 1$ iff according to the indirect evidence produced by Bethany for Alfonse and that Alfonse takes to be reliable at w, licorice is tasty.

What this report says is that that Bethany is actually the source of some indirect evidence that licorice is tasty, and that Alfonse believes this (e.g., reportative evidence: perhaps she said so),[32] and that Alfonse takes this evidence to reliably indicate the truth of the proposition (and hence, he takes licorice to be tasty). Alfonse's experiential and perceptual alternatives at w are irrelevant to $Dox_{a,b,w}$: his beliefs based on this indirect source are unconstrained by them, and so even if he hasn't tasted licorice, and so has no direct evidence that it is or is not tasty, the report is appropriate, given that he trusts whatever indirect evidence that he's received from Bethany.

So it is that even where Alfonse is in the same doxastic state, so long as he has not tasted licorice, it is misleading to report his beliefs evidentially neutrally (*Alfonse thinks that licorice is tasty*), since this commits him to having had an experience that he hasn't, but not to report his beliefs on the basis of an indirect evidential source (*Alfonse believes Bethany that licorice is tasty*).[33] Belief reports in general thus have evidential consequences, and there is no notion of an individual's doxastic alternatives *simpliciter*, only alternatives as they pertain to an evidential modal base.[34]

32. It may be that *believe*, with an overt indirect evidential source argument, requires specifically reportative evidence. However, the notion of indirect evidence is left more open than this, because the same seems not to be true when the argument is left implicit: *Alfonse believes that licorice is tasty* is true even if Alfonse believes this only inferentially, based on its ingredients. Where the indirect evidential source is left implicit, it may be existentially quantified: and so this report effectively states that Alfonse believes that licorice is tasty, on the basis of some indirect evidence or other.

33. It may of course be possible to introduce overt material that composes with a *think*-report to form non-neutral belief reports, like justification clauses, as in *Alfonse (only) thinks that licorice is tasty because Bethany says so*, which may have a reading with no inclination to an acquaintance inference anchored to Alfonse. Presumably, the *because*-clause here acts as some sort of restrictor of the evidential base.

34. It is for this reason that there is no contradiction in the fact that Alfonse holds licorice to be tasty with respect to an evidential base that ignores his perceptual evidence, and neither holds it to be tasty nor not tasty with respect to an evidential base that includes that evidence. This oddity simply reflects the oddity of the facts: that where Alfonse has not tasted licorice, one can felicitously report Alfonse as *believing* that licorice is tasty, but cannot felicitously report him as *thinking* this – and, given that assertion commits one to a direct evidential source in the same way (see below), that Alfonse can believe that licorice is tasty, but

That beliefs in general are evaluated with respect to evidential bases is suggested by the fact that in some languages, belief reports can include evidentials in the embedded clause, which do not outscope the attitude environment, and which specify on what evidential basis the subject holds the belief. In Standard Tibetan, the expected difference emerges based on whether the embedded evidential is direct (*'dug*) or indirect (*yod red*). In the examples below, only (338b) requires Trashi to have the belief about the food on the basis of having tasted it (or having just tasted it), while (338a) suggests that he has not tasted it, or at least not recently, making *bsams* 'think' behave more like English *believe* when paired with an indirect evidential.[35]

(338) a. *bra shis kyis kha lag 'di bro ba chen po yod red bsams kyi yod red.*
Trashi.ERG food this tasty COP.IND think PRES COP.IND
'Trashi thinks (believes) that this food is tasty.'

b. *bra shis kyis kha lag 'di bro ba chen po 'dug bsams kyi yod red.*
Trashi.ERG food this tasty COP.DIR think PRES COP.IND
'Trashi thinks that this food is tasty.'
→ Trashi has (just) tasted the food.

'Ordinary' doxastic alternatives as encoded by *think*, and notated without a reference to an evidential source, are then evidentially neutral precisely in that their evidential modal base is unrestricted, and has access to every sort of evidence that the believer has access to, perceptual and otherwise.

In the case of reporting Alfonse's attitude towards licorice, this means that a *think*-report, insofar as Alfonse successfully takes himself to be an accurate perceiver, cannot avoid the restrictions put on the evidential base by Alfonse's experiential and perceptual alternatives: they, and so the doxastic alternatives, must be uniform based on the experiences that Alfonse has actually had, and so relative to this evidential base Alfonse is forced to believe neither

cannot felicitously say what he believes.

35. Note that there is a further evidential *yod red* in each of these examples in the matrix clause. These pertain to the evidence that the *speaker* has for the fact that Trashi has this belief – since *yod red* is used, this means the speaker has not witnessed that Trashi thinks this.

that licorice is tasty, nor that it is not tasty, so long as he has never tasted licorice. The restriction on doxastic alternatives in (327) would then pertain to those alternatives defined in terms of evidential bases that allow access to direct evidence: where direct evidence is invoked, an agent presuming itself to be an accurate perceiver must believe in line with that evidence.

And so the characterization of the evidentially neutral $Dox_{x,w}$, for any x, w, can be recast to mirror the characterization of doxastic alternatives anchored to an indirect evidential source, as in (336), the only difference being that the modal base and ordering source are unconstrained with respect to what the evidential source relevant to the belief is, or what sort of evidence it provides.

(339) $w' \in Dox_{x,w}$ iff:

a. w' preserves all evidence that x receives at w;

b. w' is maximally compatible with x's beliefs at w about the reliability of evidence that x receives at w.

Since the modal base now includes propositions about what the believer has perceived, the assumption of accurate perception will force beliefs anchored to this more open modal base to be in line with the believer's perceptions, in the way outlined above, and so belief in a proposition expressed by an experiential predication will require a direct source of evidence, unless the norm on treating one's own direct perceptions as accurate is violated (on which, more in §4.2.1).

To illustrate the violation in these new terms: suppose that Alfonse has not tasted licorice at w. Then $\delta''_{\text{PLEASURE}}(w)(a)(\iota x[taste''(w)(l)(x)])$ is undefined. So by definition, for all $w' \in Exp_{a,w}$, $\delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w')(l)(x)])$ is undefined, and since $Per_{a,w} \subseteq Exp_{a,w}$, then for all $w' \in Per_{a,w}$, $\delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w')(l)(x)])$ is undefined. But this is just to say that Alfonse's direct evidence at w verifies that licorice is neither tasty nor not tasty. Now suppose that Alfonse thinks that licorice is tasty, viz. that $\forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w')(l)(x)])]$.

By (339), since the modal base preserves *all* evidence available to Alfonse, it follows that it will include the proposition that Alfonse's direct evidence verifies that licorice is neither tasty not not tasty. But then, $Per_{a,w}$ and $Dox_{a,w}$ are disjoint: and since $Dox_{a,w}$ tracks which evidence Alfonse has that he takes to be reliable, it follows that thinking that licorice is tasty commits Alfonse to treating his own direct evidence as unreliable: for if he did treat it as reliable, then by definition what his direct evidence verifies would be true in all of $Dox_{a,w}$, which it is not. Therefore, if Alfonse has never tasted licorice, but *Alfonse thinks that licorice is tasty* is true, it follows that Alfonse is committed to his own direct evidence not being reliable. No such commitment follows from the truth of the report *Alfonse believes Bethany that licorice is tasty.*

$Dox_{x,w}$ preserves the original notion of doxastic alternatives, as that set of worlds compatible with whatever x believes on the basis of all x's evidence taken together, and in what follows in this work doxastic alternatives *simpliciter* will continue to be used in this original signification – but the consequence of keeping the modal base open to direct evidence for reports using experiential predicates should be kept in mind.

This treatment of acquaintance inferences in attitude reports can then be carried over to the original examples involving assertion, as suggested above: where evidentially neutral belief in an experiential predication requires a direct source of evidence, and where assertion commits a speaker to evidentially neutral belief in what is asserted, the relevant restriction to a direct source of evidence straightforwardly follows.

The commitment to belief via assertion must be implemented in some way. For present purposes, it doesn't matter how this is done, whether it is taken to arise as a result of a silent assertion operator in the compositional semantics, or whether it results purely conventionally via an independent pragmatic norm on speech acts. The commitment can be represented by the composition of an operator ASST, which composes with a proposition and adds to it the not-at-issue content that its speaker believes that proposition. Since as noted above, in languages like English that do not have evidential markers, assertion is evidentially neutral,

it follows that the speaker's committed belief ought to be evidentially neutral, and so it is constrained by Dox_{s_c,w_c}, i.e. the set of worlds compatible with that the speaker believes, on the basis of all available evidence, at the world of assertion.

(340) $[\![\textsc{asst}]\!]^{c,w} = \lambda\phi_{st}.\langle\phi(w), \forall w' \in Dox_{s_c,w_c}[\phi(w')]\rangle$

Assertion in general then commits the speaker to not-at-issue belief in the asserted proposition, but per the above this in turn commits the speaker to having a direct source of evidence for this belief when experiential predications are asserted. And so the relevant not-at-issue content in (341), which reflects the content of the speaker's assertion that licorice is tasty, entails (342).

(341) $[\![\textsc{asst}\ [\text{licorice}\ [\text{is}\ [\textsc{gen}\ [\textsc{pos}\ \text{tasty}]]]]]\!]^{c,w}$
$= \langle 0_{\textsc{pleasure}} <_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w)(\iota x[taste''(w)(l)(x)]),$
$\forall w' \in Dox_{s_c,w_c}[0_{\textsc{pleasure}} <_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w')(\iota x[taste''(w')(l)(x)]]\rangle$

(342) $\forall w' \in Per_{s_c,w_c}[0_{\textsc{pleasure}} <_{\textsc{pleasure}} \delta'_{\textsc{pleasure}}(w')(\iota x[taste''(w')(l)(x)]$
$\rightarrow 0_{\textsc{pleasure}} <_{\textsc{pleasure}} \delta''_{\textsc{pleasure}}(w_c)(s_c)(\iota x[taste''(w_c)(l)(x)])$

And so in virtue of asserting using *Licorice is tasty*, the speaker commits to the not-at-issue content that the taste of licorice has stimulated a non-zero degree of pleasure in the speaker, which requires that the speaker has tasted licorice and liked its taste. This is just the acquaintance inference derived from the overt direct evidential in §4.1.3, and it can be applied to the cases in (321) to yield analogous results, getting the right sensory and qualitative inferences in all the right places.

So it is that the acquaintance inferences triggered by experiential predicates in assertions are the result of speakers committing to direct sources of evidence for their assertions via the assertion's not-at-issue content: and this is so regardless of whether this is accomplished via an overtly marked direct evidential, or by the ordinary force of commitment to belief in evidentially neutral assertion.[36] The relation between these two sorts of commitments may

36. There is therefore a sense in which use of an overt direct evidential is 'overkill' as far as the acquaintance

ultimately be more deeply related: perhaps the conventional commitment to belief in assertion results from commitment to a neutral evidential source for an asserted proposition, and so languages without evidential markers in effect have one-evidential systems. But whatever the case, all languages are apparently conventionally sensitive to the encoding of direct evidentiality, though in these languages it may require examining experiential predicates to see this.

4.2 Pragmatics of experiential commitments

This section addresses some pragmatic consequences of the above treatment of commitment to direct experience in terms of direct evidentiality, and in particular with respect to the default commitment to the accuracy of own's own direct evidence, and what results when speakers differ in the direct evidence that they have for experiential properties. §4.2.1 addresses the default commitment to one's own direct evidence being accurate with respect to experiential properties, and the asymmetry between attributing inaccurate perception to oneself versus to others. §4.2.2 shows how the present treatment of experiential semantics, in virtue of not relativizing experiential predicates' extensions to any experiencer, allows for normative effects stemming from the accuracy or inaccuracy of perceptions, which in turn normatively compel speakers to experience in a certain way.

4.2.1 Accurate and inaccurate perceivers

In §4.1.4, the commitment to direct experience in virtue of evidentially neutral beliefs about experiential properties was derived due to a default commitment by speakers that they perceive accurately in a certain way. This was encoded by having speakers' doxastic alternatives

content is concerned: a neutral (neither solely direct nor solely indirect) evidential base derives this content just as well where experiential predicates specifically are concerned. But languages with a direct/indirect contrast often require this yet stronger commitment to direct evidence where commitment to indirect evidence is lacking, and so marking with a direct evidential may be the only way for the speaker to accomplish an analogous speech act.

be a subset of their perceptual alternatives: and so by default, speakers are committed to believing what they perceive. There are two issues to address with respect to this proposal – first, what reason there is to believe in this default commitment, and second, how this commitment manifests independently of acquaintance inferences with the use of experiential predicates. As to the first, there are a number of features of direct evidentials cross-linguistically that suggest their association with truth, speaker commitment, and certainty.

First, unlike indirect evidentials in a number of languages (cf. §4.1.1), direct evidentials apparently never allow the speaker to eschew commitment to the truth of an asserted proposition when used. Murray (2017: §3.2.1) implements this distinction using Cheyenne as an exemplar language, by casting direct evidentiality as requiring a speaker to propose an update to the common ground with the asserted proposition, while reportative evidentials propose no update at all, and inferential evidentials merely propose to update with the possibility that the asserted proposition is true.

Second, when a speaker has multiple sources of evidence for the truth of a proposition and attests to its truth, those sources of evidence are hierarchically arranged, such that evidentials marking sources of evidence higher on the hierarchy are systematically preferred to those lower on that hierarchy. Direct evidence is cross-linguistically preferred in this sense to indirect evidence: use of a direct evidential marker is expected of a sincere speaker who has the relevant direct evidence, even where that speaker also has indirect evidence, and use of an indirect evidential strongly implies that the speaker does *not* have the relevant direct evidence. While there is some cross-linguistic variation, this preference for direct over indirect evidence is plausibly universal. For expositions of such hierarchies in specific languages, cf. e.g. Barnes (1984) for Tuyaca and Oswalt (1986) for Kashaya: these are summarized in Aikhenvald (2004: §10.1), along with some further examples, and cf. Willett (1988), de Haan (1999), and Faller (2002: §2.4) for general expositions of the phenomenon.[37]

37. There is further complexity in these evidential hierarchies cross-linguistically: direct evidence itself is

Third, direct evidentials, and especially visual evidentials, acquire secondary meanings in a a number of languages, and are used to convey heightened certainty in, or responsibility for, a claim made. That this use of a direct evidential is semantically distinct from its canonical use is shown by the fact that when used to invoke certainty in this way, the relevant sensory experience (i.e. visual experience) is not necessarily required. Aikhenvald (2004: 170) provides an example from Tariana, where the speaker expresses conviction and certainty in what is stated using a visual evidential, despite not having directly perceived that this is so (and cf. *ibid.* for further similar examples).

(343) waha-wya-ka enu i-daki-ne-naka
we-LIM-DECL thunder INDEF-grandchild-PL-PRES.VIS
'We **are** grandchildren of Thunder.'

[Tariana. Aikhenvald 2004: 170, ex. 5.45]

The purpose of listing these cross-linguistic regularities is to demonstrate that direct evidentiality as a linguistically encoded cognitive category is highly 'adjacent' to the notions of commitment to truth and certainty, not that direct evidentiality *per se* encodes this in virtue of its logical structure: it may do this as well, but the above considerations do not show this. This means: (i) that where direct evidentials are used with their canonical meanings, they impose requirements of truth and certainty; (ii) that where evidence pertaining to truth and certainty is required, direct evidence is robustly preferred over indirect evidence; and (iii) where direct evidentials are extended to non-canonical epistemic meanings, these meanings themselves pertain to matters of truth and certainty. Why this association between direct evidence and truth or certainty holds, and why it is reflected in the speaker presumption of

typically hierarchically arranged, with the visual evidential preferred to the sensory non-visual evidential, where both exist (this is also reported in Barnes *ibid.* for Tuyaca). A number of Tibeto-Burman languages also contain 'egophoric' marking that may be evidential in nature (though this is disputed), which signals personal involvement and is preferred by default even over the direct evidential: cf. Gawne (2017) for an introduction. Faller (2002: ch. 4) takes this notion of preferred evidence and its connection to direct evidentiality a step further, suggesting that the Cuzco Quechua evidential *-mi* actually encodes the 'best' possible evidence that one can have for an asserted proposition, and that this is a type of direct evidence where applicable.

accurate direct evidence is a deeper issue, which won't be addressed here – it is sufficient to note that it is a robust cross-linguistic default.[38]

The default commitment to the accuracy of one's own perceptions shows up in the way speakers treat bare experiential predications in relation to statements about their own experiential states. This has been noted e.g. in Pearson (2013: 11ff.) – assertions that something has an experiential property are ordinarily infelicitous alongside assertions that one has an experiential state that conflicts with this (cf. also MacFarlane 2014: 4, exx. 1-3).[39]

(344) Licorice is tasty...

a. ?but it's not tasty to me.

b. ?but I don't find it tasty.

c. ?but I don't like its taste.

For speakers to commit to these sorts of statements is for them to hold that a stimulus has a certain experiential property *simpliciter*, but also to hold that they themselves are in an experiential state that directly evidences the opposite – in short, that they are inaccurate perceivers of a certain sort, and have perceived something that isn't so. On the present semantic treatment, such statements are not contradictory: the continuations in (344), for example, receives the following satisfiable interpretation.

(345) $0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota x[taste''(w)(l)(x)])$
$\wedge\, 0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(s_c)(\iota x[taste''(w)(l)(x)])$

38. A tempting explanation for this fact is to suppose that direct evidence is factive: one can't directly perceive what isn't so. This would be represented by making perceptual alternatives realistic with respect to the anchor-world: thus, for all x, w: $w \in Per_{x,w}$. But this has no interesting explanatory power, even if true: the question would still remain why speakers by default are committed to taking their own experiences to reflect direct evidence in this factive sense. It makes little difference whether one takes direct evidence to be intrinsically non-factive, and to make the question why speakers take their own direct evidence to be factive, or takes direct evidence to be factive, and to make the question why speakers take their own experiences to provide them with direct evidence. The text, as already noted, assumes the first approach: it is possible for one to perceive something that isn't so, and also for distinct individuals to have contradictory direct evidence.

39. MacFarlane makes use of operators that void the relevant acquaintance inference, which demonstrates that the oddity of these sentences is not merely due to the acquaintance inference conflicting with the at-issue asserted content (though the two do conflict, where the inference appears).

There is no reason that both these conjuncts cannot be true: the speaker can be in an experiential state of pleasure stimulated by the taste of licorice, while that same taste produces no pleasure in ideally general conditions. If accurate and inaccurate perception of a proposition is defined as follows, it can be shown that such a statement commits the speaker to being an inaccurate perceiver with respect to the asserted proposition.

(346) a. x is an accurate perceiver with respect to ϕ at w iff:

$\phi(w) \wedge \forall w' \in Per_{x,w}[\phi(w')]$

b. x is an inaccurate perceiver with respect to ϕ at w iff:

$\phi(w) \wedge \forall w' \in Per_{x,w}[\neg\phi(w')]$

Thus for any true proposition ϕ, an individual is an accurate perceiver of it just in case its direct evidence verifies it (i.e. if it directly perceives that it is so), and is an inaccurate perceiver of it if its direct evidence verifies the opposite. This is applicable to any proposition verifiable by direct evidence, whether it pertains to experiential properties or not – but where experiential properties are concerned, and direct evidence entails being in a certain experiential state, being an accurate or inaccurate perceiver will entail having a certain sort of experience, which aligns or misaligns with the experiential properties of some stimulus.

Suppose then that (345) is true: it follows that the speaker s_c is an inaccurate perceiver with respect to the proposition *that licorice is tasty* at w, since the second conjunct in (347), which instantiates (346b), reduces to the second conjunct in (345), making the two identical.

(347) $0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota x[taste''(w)(l)(x)])$

$\wedge \forall w' \in Per_{s_c,w}[0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w')(\iota x[taste''(w)(l)(x)])]$

The speaker is therefore committed to perceiving inaccurately, and so to thinking something that conflicts with her direct evidence, and infelicity results.

That there is nothing wrong with the truth conditions in (345) is shown by the fact that they are reportable of someone other than the speaker with no such infelicity. So where s_c

above is Alfonse, the following report spoken by someone other than Alfonse gets the same interpretation, and is not infelicitous.

(348) Licorice is tasty...

a. but it's not tasty to Alfonse.

b. but Alfonse doesn't find it tasty.

c. but Alfonse doesn't like its taste.

(349) $0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota x[taste''(w)(l)(x)])$
$\wedge\, 0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(a)(\iota x[taste''(w)(l)(x)])$

But this entails that Alfonse is an inaccurate perceiver. In the spirit of the Moorean analysis of §4.1, then, it is strange for speakers to cast themselves as being inaccurate perceivers with respect to what they assert, but there is typically no problem with casting others as being so. Direct evidence is then intrinsically autocentric in this way as well: the commitment is to the accuracy of one's own direct evidence.

The commitment to being an accurate perceiver is a powerful default, but it is defeasible: speakers might, for a number of reasons, take themselves to be inaccurate perceivers, and so cease to take their own direct evidence as constitutively tied to the truth of experiential propositions. This might be because of a sensory defect for predicates sensitive to sensory modality, such as lacking a sense of taste, but it might also be because speakers take themselves to be abnormal and so not in the 'target audience' for evaluating experiential properties, making their experiences irrelevant (cf. Pearson 2013: 144), or simply to be self-deprecating. And so the assumption that the speakers' direct evidence tracks the relevant experiential properties can be voided, where the speaker takes there to be some powerful reason that the two disconnect. Anthony (2016: §2.2) offers examples of this sort, which are made more felicitous when accompanied by reasons for the disavowal of one's own direct evidence.

(350) a. The show is funny, but it isn't funny to me
(anymore, since I've seen it so many times).

b. (I'm sure) math is fascinating, but it doesn't fascinate me
(, since I can't understand it).

And so where the speaker disavows the ability to perceive the experiential properties of a stimulus properly for whatever reason, the norm on accurate perception disappears. Better said: one is always free to evaluate the experiential properties of an individual in a way conflicting with one's own experiences, but the necessary consequence of this is that one commits to being an inaccurate perceiver. This is not an all-or-nothing commitment, but can be compartmentalized, so that one perceives some propositions accurately, and others inaccurately.

According to §4.1.4, the acquaintance inference in languages without direct evidential markers was taken to depend on the default commitment to accurate perception (of the proposition asserted). It follows that where a speaker no longer makes this commitment, the acquaintance inference ought to disappear. This is apparently right: suppose that Alfonse has tasted licorice and likes its taste, but then burns his tongue and loses his sense of taste. He can then still assert *Licorice is tasty*, despite the fact that from this there is no inference that he likes its taste (he might follow with, *it's too bad it isn't tasty to me (anymore).*)[40]

40. The inference may actually still be preserved, with a time lag – it is not obvious that Alfonse can utter *Licorice is tasty* felicitously if he had never tasted licorice before losing his sense of taste, or if he didn't like it beforehand. Something more complicated may therefore be going on here – perhaps Alfonse still has direct evidence that licorice is tasty, from his past experience, and takes his present experience not to yield trustworthy direct evidence that might contravene it. If this holds in general, then the requirement that speakers take themselves to be accurate perceivers with respect to experiential properties may be even more stringent than the text implies: (350a) works so well precisely because the speaker used to find the show funny, and when there is no implication of past direct evidence verifying the proposition, as in (350b), it seems that an epistemic qualifier (like *I'm sure*) may be needed for the assertion to palatable. However, the important thing to note is that, as predicted, the stringency of the accurate perception requirement patterns with the stringency of the acquaintance inference.

4.2.2 Experiential normativity

Experiential predicates can be used to construe stimuli as ideally disposed to produce certain experiences *simpliciter*, and not merely relative to some experiencer or other. In tandem with the fact addressed in §4.2.1, that speakers can be construed as perceiving accurately or inaccurately with respect to these experiential properties, this yields a normative effect expressed using the truth conditions of experiential predicates: some experiencers' dispositions to experience are in line with what stimuli are ideally disposed to stimulate, and some are not. In differently construing the felicity conditions on the generic component of experiential predicates, speakers can take these conditions to align or misalign with actual dispositions to experience, and this in turn is to cast experiential dispositions as aligning or misaligning with ideal normality. This notion of ideal normality then has the ability to take on a normative component, in construing which sorts of experiential dispositions are appropriate or defective.

This ability to compare experiencers' dispositions to an ideal standard is a special expressive power of the semantic position adopted in §1.3, that experiential predicates are not intrinsically relativized in their denotation to any experiencer whatsoever. If there is no reference in the semantics to a disposition to produce an experience *simpliciter*, but rather only in some experiencer or other, then this comparison cannot be made, and the normative effect goes missing – one can only use the truth conditions of these predicates to describe which experiencers are disposed to experience what from which stimuli.

In this sense, a bare semantics for experiential predicates is strictly more expressive than its relativist counterpart described in §1.2: the former, but not the latter, is capable of encoding ideal normality independent of an experiencer. While in §1.3, the difference between a relativist semantics and a bare semantics was shown to be trivial on compositional semantic matters, and pragmatic matters of truth assessment, here at a higher level of description, involving the normative stakes of truth conditions, a final difference between the two appears.

To take an example, suppose that a comedian has told an utterly tasteless, offensive

joke, and that Alfonse laughs uproariously at it.[41] Those around him take issue with his reaction – it is really wrong to laugh at such a heinous joke, and it is just *not funny*. They take Alfonse to have made an error of some sort, even a normatively quite serious one – and issues of the way the world is aside, the disagreement between Alfonse and his detractors is not at all 'faultless' *simpliciter* by the speakers' lights, but turns on a substantive issue of what it is appropriate to laugh at.[42] Recall now (cf. §1.2) what the relativist treatment of the semantics of the claim that the joke is funny is, and how speakers are pragmatically compelled to assess it for truth.

(351) a. $[\![\text{funny}]\!]^{w,x} = \lambda y_e.funny''(w)(x)(y)$

b. $[\![[\text{the joke}]\ [\text{is funny}]]\!]^{w,x} = funny''(w)(x)(\iota y[joke'(w)(y)])$

(352) a. $\lambda w_s.\lambda x_e.funny''(w)(\iota y[joke'(w)(y)])$ is correctly assessed as true in c' iff it is true *simpliciter* at $w_{c'}$.

b. $\lambda w_s.\lambda x_e.funny''(w)(\iota y[joke'(w)(y)])$ is incorrectly assessed as true in c' iff it is false *simpliciter* at $w_{c'}$.

(353) a. $\lambda w_s.\lambda x_e.funny''(w)(\iota y[joke'(w)(y)])$ is properly assessed as true in c' iff: $funny''(w_{c'})(a_{c'})$.

b. $\lambda w_s.\lambda x_e.funny''(w)(\iota y[joke'(w)(y)])$ is improperly assessed as true in c' iff: $\neg funny''(w_{c'})(a_{c'})$.

The relativist-proposition that *The joke is funny* epresses is true relative to an assessor at w just in case the joke in question at w is funny, or disposed to be funny, to that assessor. It

41. The text need not give the details of the joke. The illustration works best if the reader imagines something that actually repulses them in this way – pick the most vile joke that you can imagine, and suppose that Alfonse finds it hilarious.

42. Anthony (2016: 714, exx. 66-69) offers several politically charged examples of this sort, which make claims using experiential (deverbal psych) predicates designed to make the reader understand immediately the normative stakes involved in virtue of one's experiential reactions. He further claims (cf. *ibid.* §3.3), quite plausibly, that researchers have suffered from theoretical blindness to this non-faultless, normative dimension of differing experiential dispositions due to being themselves normatively disinterested in the constructed examples proffered when conducting research.

is never assessed correctly or incorrectly as true, so long as there are experiencers that non-trivially differ with respect to their humorous reaction to the joke at w (cf. (85) and (86b)), since then it is neither true nor false *simpliciter* at w. And one is by default pragmatically compelled to assess its truth autocentrically, i.e. with respect to oneself as assessor.

Supposing now that Alfonse at the example world w honestly assents to the claim that the joke is funny (on a non-exocentric reading), based on his own humorous reaction – what error can his detractors coherently take him to having made, according to the relativist semantics and pragmatics? He has not assessed the truth of the claim incorrectly, as this is not even possible, in the relevant factual sense (nor is it possible to assess it incorrectly). He has not assessed its truth improperly: in fact, it would be improper of him to assess otherwise, given his own reaction.[43] Further, the factual commitments he makes in virtue of thinking that the claim is true commit him only to a truth, viz. that the joke is funny *to him* (cf. §1.2.1, esp. (90b)).

(354) a. $[\![\text{think}]\!]^{w,x} = \lambda\phi_{s,et}.\lambda y_e.\forall\langle w', z\rangle \in Dox_{y,w}[\phi(w')(z)]$

b. $[\![\text{Alfonse [thinks [that [[the joke] [is funny]]]]]}]\!]^{w,x}$
$= \forall\langle w', z\rangle \in Dox_{a,w}[funny''(w')(z)(\iota u[joke'(w)(u)])]$

It is true that in virtue of Alfonse's detractors also assessing autocentrically and forming their beliefs on this basis, that they can be correctly reported as disagreeing with him (cf. (91b)-(91c)), can felicitously deny his claims, or claim that he has said something false, and so on. The problem is that the relativist machinery implies that to disagree with Alfonse in either of these senses *just is* to take oneself to have a distinct experiential reaction from him, and there are no normative stakes intrinsic expressing these differing experiential reactions

43. It does no good to say that where one's experiential dispositions are of a certain defective sort, it is inappropriate to assess autocentrically, because this is empirically just not true – indeed if Alfonse finds the joke funny, but sternly claims that it is not, he is in a real sense being *disingenuous* with his interlocutors, and therefore violating a pragmatic norm to save social face, so long as he does not commit to being an inaccurate perceiver with respect to humor (cf. §4.2.1). In other words, even though Alfonse's detractors think that something is wrong with his assessment, this is still the assessment that, pragmatically, he *ought to have*, in order to ingenuously reflect his own experiences as expected by use of the predicate.

at all.

This misconstrues the character of the disagreement – the idea is that Alfonse has a distinct experiential reaction from his detractors, and that this reaction, according to his detractors, makes his own opinion wrong in virtue of his experiences not being just *different*, but *defective* in virtue of this difference. And so the relativist semantics and pragmatics predicts no error on the part of Alfonse in assessing the way he does – the fact that his detractors take him to be making an error shows that they are either behaving incompetently, or the relativist picture is missing something crucial.

The relativist literature has expressed anxiety over this issue of what the 'substance' of disagreement is. Lasersohn (2005: 683-684) says: "What, then, are speakers disagreeing about? I think that the only answer one can give to this is that the two sentences cannot both be accommodated into a single coherent perspective. [...] Such disagreements are in some sense 'without substance.'" And MacFarlane (2007: 30): "But why should controversy feel uncomfortable even when the disagreement is entirely due to to differences in the interlocutors' respective contexts of assessment? One possible answer is: it just is. That's a brute psychological fact about us." These sorts of responses do not adequately address how it is possible for speakers to take one another to be in error in virtue of these truth assessments. Of course relativists can say whatever they want about the coherency of belief, the social role of conflict and disagreement, and which sorts of experiential reactions speakers approve and disapprove of. The problem is that none of this addresses how a speaker can be competently taken to make an error in virtue of an appropriate autocentric assessment, or how speakers are capable of using the truth-conditional contents of experiential predicates to express this fact: the problem is that while the joke is funny *to Alfonse*, it is according to his detractors *not funny simpliciter* (and not merely not funny to them).

Lasersohn has made two moves that might offer to address this shortcoming. The first is to introduce the notion of an *error of taste* (cf. Lasersohn 2017: 210). The idea is that belief in a relativist-proposition can be construed as an error relative to an assessor, where

that relativist-proposition's truth is non-trivially dependent on the value of the experiencer parameter, is false relative to the assessor in question. Adapting from *ibid.*, ex. 249:

(355) x commits an error of taste in believing ϕ at w according to y iff:
x believes ϕ at w, $\neg\phi(w)(y)$, and there is some z such that $\phi(w)(z)$.

Suppose that y is one of Alfonse's detractors at w, and that ϕ is the relativist-proposition *that the joke is funny*. Then since Alfonse thinks this is true, and the joke is not funny as properly assessed by y, but there is some z such that the joke is funny as properly assessed by z (at the very least, Alfonse), it follows from this definition that Alfonse commits an error of taste in believing that the joke is funny according to y at w.

But this is of no help, because for Alfonse to commit an error of taste in this way according to y is no more than for Alfonse to assess properly relative to his own experiential dispositions, and for these to differ from y's dispositions as to the truth of the relativist-proposition that the joke is funny. In other words, this notion of an error of taste adds no expressive power to the relativist semantics or pragmatics, and so the problem remains the same as above: to say that an error of taste was made only amounts on this account to saying that Alfonse's experiences differ from his detractors', which again has no normative interest.

To see this more clearly, note that if Alfonse's detractors are competent in recognizing errors of taste so construed, they will understand that they have made a symmetrical error of taste, according to Alfonse, in thinking that the joke is not funny (as the reader can confirm): but the problem is that the detractors take their own experiential dispositions to track the truth of the matter in a way that Alfonse's don't, and so they take Alfonse to have made a sort of error that they *have not*. What is required, therefore, is some means of privileging one experiencer over another, with respect to their role as experiencer in truth assessment.

Lasersohn (2017: §10.1) also entertains the idea that experiencers (corresponding to his 'perspectives') can be ranked, according to their 'objective superiority in assessing the truth of particular contents (*ibid.* 215):' thus, one might write '$\langle x, w\rangle <_\phi \langle y, w\rangle$' to express that

at w, y as an experiencer is superior to x at assessing the truth or falsity of the relativist-proposition ϕ. Where x is Alfonse, y one of his detractors, w the example world, and ϕ the content of *that the joke is funny*, this formulation would express that the detractor is an experiencer objectively superior to Alfonse in the assessment of the truth of whether the joke is funny. This in turn might be what the detractor is somehow expressing in insisting that Alfonse is wrong: and thus while the joke is funny from Alfonse's perspective, it is not according to the superior detractor's perspective, and it might be not funny *simpliciter* if the ranking on perspectives converges on a truth value of false for the relativist-proposition with respect to all best-ranked experiencers (cf. *ibid.* 217, ex. 252).[44]

But Lasersohn does not address in virtue of what one experiencer might be ranked above another in assessing the truth of a relativist-proposition. The relativist machinery itself provides no guide as to how this might be done – it only details what it is for a relativist-proposition to be true or false with respect to an experiencer (at a world). The idea is that one perspective is better than another *in assessing the truth of a content*, but it is not clear what this could possibly mean, since experiencer-sensitive relativist contents are only true with respect to some experiencer or other, and are not true or false *simpliciter* – in virtue of what, then, can one assessment from one experiencer be 'better' than another, if it tracks no truth independent from truth relative to that experiencer?

To allow that there is an independently defined notion of truth *simpliciter* at a world for these contents, and to rank perspectives on this basis, is just to collapse relativism into the bare treatment of experiential predicates, by assuming that the world determines whether an experiential predicate applies, and then various experiential dispositions might or might not adhere to it. If, by contrast, the ranking on perspectives as to truth is not

44. It is unclear if Lasersohn himself would countenance this machinery for ranking perspectives to treat predicates like *funny* – his own examples involve slightly different sorts of aesthetic predicates less directly tied to brute experiential reactions, like *high quality* as predicated of wine. This division between 'purely' subjective predicates and ones that can be objectively ranked according to perspective is not promising for present concerns: it seems that *any* experiential predicate can in principle be employed normatively in this way, and so would require appeal to ranked perspectives if this approach is going to work.

itself determined by the world, but is experiencer-dependent, then the issue recapitulates at this level: Alfonse's detractors cannot merely be taking their own perspectives to be superior to his, from their own perspectives: a suitable way of cashing this out in terms of differing experiential dispositions will make it equivalent to just another admission that their experiential dispositions differ from Alfonse's.

The relativist semantics and pragmatics as it stands thus has no way to make sense of these sorts of normative disagreements on the basis of differing experiential reactions. By contrast, the bare treatment makes sense of the situation in terms of plain old truth: Alfonse's detractors take the joke to by funny to Alfonse, but not funny. And they take their own reactions to be superior to his, since they are capable of recognizing the truth that the joke is not funny due to these reactions, while he is not.

Recall the bare treatment of the generic, positive form of *funny*, which is true of a stimulus just in case it is ideally disposed to produce a non-zero degree of pleasure. *The joke is funny* is thus true at w just in case the joke actually has this disposition at w.

(356) a. $[\![\text{GEN } [\text{POS funny}]]\!]^w = \lambda x_e.0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w)(x)$

b. $[\![[\text{the joke}] \text{ [is [GEN [POS funny]]]}]\!]^w = 0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w)(\iota y[joke'(w)(y)])$

Alfonse's detractors then take him to be an inaccurate perceiver at w with respect to the proposition *that the joke is not funny* (cf. §4.2.1, esp. (346b)). They think that it is true, while taking Alfonse to have direct evidence that it is false (and so he directly perceives that the joke *is* funny). And so where x is one of these detractors:

(357) $\forall w' \in Dox_{x,w}[0_{\text{HUMOR}} =_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w')(\iota y[joke'(w)(y)])$
$\wedge \forall w'' \in Per_{a,w'}[0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\text{HUMOR}}(w'')(\iota y[joke'(w)(y)]]$

It follows from this that the detractor takes the ideally normal conditions under which the joke is disposed to produce experiences are misaligned with its disposition to produce experiences in Alfonse in a certain way, and so Alfonse's experience of the joke is not ideally

normal, and therefore deviant. Recall from §2.3.3 (cf. (220b)) that the extension of *the joke is not funny* can be expanded as (358).

(358) $\forall w' : w\rho_{\text{HUMOR}}w'[0_{\text{HUMOR}} =_{\text{HUMOR}} \delta_{\text{HUMOR}}(\sigma x[stim''(w)(\iota y[joke'(w)(y)])(x)])]$

That is, the joke is not funny *simpliciter* at w just in case at all worlds ideally normal with respect to the intrinsic qualities of the joke at w for the production of humor, the maximal experience that the joke stimulates instantiates a zero-degree of humor *simpliciter*. Since at w this is not so, given that Alfonse has a humorous experience of the joke, it follows that if the joke really isn't funny at w, then w is itself not among these ideal worlds, nor is any world in which the intrinsic properties of the joke stay the same while someone finds it funny. Alfonse's experiential dispositions are thus sufficient to make the world non-ideal with respect to the production of humor, since in an ideally normal world no one would find the joke funny at all, since it isn't funny. Alfonse's actual experiential reactions thus necessarily constitute a departure from the ideally normal production of experiences – whether the joke is disposed to produce humor *simpliciter*, and whether it is disposed to produce pleasure in Alfonse, do not match.

This deficiency in turn takes on a normative aspect, in virtue of the fact that commitment to truth is normative: if Alfonse is normatively required to believe true things, and if the joke really is not funny, then in virtue of thinking that it is funny, he has committed a normative error. And what's more, he does so on the basis of commitment to the accuracy of his own perceptions. What this means is that insofar as one is obliged both to believe true things, and to take one's own direct perceptions to be accurate, it is not possible to refrain from error with respect to truth and uphold all of the relevant pragmatic norms on the accuracy of one's own experience, if one experiences in a way that is deviant from the ideally normal production of experiences by stimuli. It follows effectively that in virtue of norms on truth, speakers actually take each other to be obliged to experience things in a certain way, and take certain dispositions to experience to be unacceptable, in virtue of the truth conditions

of experiential predicates.[45]

This conclusion has two puzzling features. First, it follows from the above that Alfonse's detractors take Alfonse not only to have made a normative error, but a factual one as well – this is because on a bare treatment, the world determines whether a stimulus is ideally disposed to produce a certain experience or not. There would therefore seem to be a 'fact of the matter' about whose experiential dispositions are deviant, and perhaps this is problematic. Second, this treatment offers no explanation for why conversely many cases of disagreement stemming from differing experiential dispositions are not treated as normatively serious, and do have a genuinely 'faultless' air. Both these issues will be addressed in Chapter 5, where the possibility of construing normative error in virtue of bare experiential truth conditions is argued for, even in the absence of commitment to factual error, and where construal of this sort of normative error is taken to result from insisting on metasemantic behaviors pertaining to the application of predicates.

45. Note that the relativist cannot make use of a simple norm of truth to derive this behavior, since there cannot on a relativist semantics be an unqualified norm of truth on experiential propositions: some provision must be made for requiring belief in true propositions relative to some experiencer or other – and this then collapses into simply believing that certain stimuli are disposed to produce experiences in certain experiencers, and the problem begins anew.

CHAPTER 5

EXPERIENTIAL METASEMANTICS

Metasemantics is the study of the conditions under which linguistic expressions receive their semantic values.[1] In a compositional, truth-theoretic framework, where one goal of the semanticist is to assign truth-conditional denotations to the lexical items of a language, a major goal of metasemantics is to characterize what it is for an expression to have such a denotation.

This question is especially pressing for experiential predicates, which have been thought to exhibit behaviors problematic for a traditional truth-conditional semantics. Because these predicates' truth conditions relate in some way to experiential dispositions, and because speakers tend to assess them autocentrically, with respect to their own dispositions, the construal of their truth conditions varies across speakers as these dispositions vary. This raises the question of what if anything the truth-conditional content of an experiential predicate like *tasty* might be in the language *simpliciter*, since there do not seem to be any features of the world in virtue of which such predicates truly or falsely apply to individuals inviariantly across the linguistic community. This has given rise to various claims to the effect that experiential predicates are somehow 'subjective,' that they pertain to 'matters of opinion' as opposed to 'matters of fact,' and so on – and since Kölbel (2002), authors have attempted to expand truth-theoretic semantics to accommodate them, alongside other expressions thought to have similar behaviors.

If, as argued in §1.3, experiential predicates are not relativized to an experiential standard in their semantic content at a context of utterance, the question remains how to characterize their denoting an experiential property *simpliciter* relative to such a context. Two broad approaches to this question have been explored. The first is to claim that discourse involv-

1. Cf. Burgess & Sherman (2014) for an overview of this term, and some construals of what the goals and boundaries of metasemantics might be. The present chapter does not fully adhere to anything said there, or to any other extant characterization of metasemantics, which is still a somewhat nebulous region of inquiry, but one that can't be ignored in giving an account of the semantics of experiential predicates.

ing experiential predicates is radically defective in some way. Hirvonen (2016) epitomizes this approach in claiming that if experiential predicates really are treated by speakers as denoting unrelativized properties in the language, then either these expressions actually are relativized to a standard of experience as contextualists claim, and so speakers are suffering from widespread semantic blindness, or speaker treatment of these predicates ought to be taken at face value, in which case these predicates denote bare, unrelativized experiential properties – and since these don't exist (or so one might think), an error theory of experiential discourse results, and all experiential predications are strictly speaking false.[2]

The second approach is to accept that experiential predicates are not relativized to an experiential standard at a context of utterance, but deny that any systematic error results, by enriching the notion of a point of evaluation in the intensional semantics, and taking it to include more information than a possible world parameter alone supplies. Together with some supplementary pragmatic theses, this results in a relativist semantics, as covered in §1.2: if the extensions of expressions depend not just on the way the world is, but either on an additional parameter, such as a judge (cf. Lasersohn 2005), a center (cf. Egan 2010), or a gustatory standard (cf. MacFarlane 2014: ch. 7), or on a parameter that acts as a 'refinement' of a world and settles 'matters of opinion' as well as 'matters of fact,' like a perspective (cf. Kölbel 2002) or an outlook (cf. Coppock 2018), then speakers can assess for truth relative to these enriched points of evaluation, leading to competent differences in truth assessment that do not need to track differences in the state of the world.[3]

2. Some works that take the semantic blindness approach seriously are Stojanovic (2007: §2), Cappelen & Hawthorne (2009: 118), and Hirvonen (2014: ch. 14). It is not obvious how native speakers of a language could be in systematic behavioral error as to the meanings of a huge class of expressions in that language, since plausibly speaker behavior is what in large part constitutes the meanings of these expressions to begin with. The blindness-theorist therefore faces not an impossible task, but one beset from the start with extreme methodological difficulties: cf. Baker (2012: §4) and Zeman (2016: §6) for comments. The text will not entertain the blindness approach, since the treatment given in this chapter obviates it. The error-theoretic conclusion will be addressed briefly in §5.2.3.

3. So long as these theories' supplementary pragmatic theses bid speakers to set the values for these intensional parameters using features of the context in which the extensions of expressions are assessed (cf. §1.2), the relativist frameworks are all more or less identical, *modulo* some trivial technical details. Cf. Lasersohn (2013) for such an ecumenical account of the position.

The present work has taken a different approach, and claimed that the values of experiential predicates are not relativized to an experiencer at any level of semantic description. While §1.3 showed that such a bare account of experiential predicates is practically indistinguishable from its relativist counterpart on compositional semantic matters, and pragmatic matters of truth assessment, a hanging worry was left at the end of Chapter 1 – what is the status of the bare experiential properties that these predicates purport to denote? Whether there are bare experiential properties in the world, and what this would amount to, is a metaphysical matter, not a semantic one – but what is a semantic matter is whether speakers commit to their existence in using these predicates, and what it would mean to do so.

This chapter addresses the metasemantic import of assigning experiential predicates such experiencer-invariant denotations. It argues that the relativist semantics is right so far as it goes, and that experiential predicates do not simply denote solely world-sensitive properties in virtue of their linguistic meaning. But it goes on to demonstrate that this is a ubiquitous feature of natural language predicates. A metasemantic account must be given of how predicates get their truth-conditional contents, and once this is given, it is apparent that experiential predicates and non-experiential predicates are fundamentally on a par, in that speakers construe their truth-conditional contents in conflicting ways. It is a mistake, therefore, to think that extra parameters of evaluation pertaining to matters of experience or opinion must be invoked for experiential predicates' assessment: rather, predicates generally converge to a greater or lesser extent on truth-conditional contents in virtue of speaker behavior, and this metasemantic fact, combined with the nature of experiential semantics and direct evidentiality, converge to make speakers behave autocentrically in their behavioral construal of the truth-conditional content of experiential predicates.

§5.1 begins with a detour into the metasemantics of predicates generally, arguing that they are almost always semantically underdetermined, and receive a truth-conditional content only with respect to speaker behavior, which seldom if ever perfectly converges in a linguistic community. This is represented by updating the grammar to include a hyper-

intensional component, which is sensitive to the truth-conditional denotations assigned to expressions via models, which are construed as patterns of semantically relevant behavior. The result is that truth conditions are assigned in virtue of speaker behavior with respect to linguistic expressions, depending on how those speakers take the world to be. This means that truth assessment is sensitive in general to both facts about the world and behavioral dispositions towards the meanings of expressions, which has consequences for how truth-sensitive attitudes are modeled, and what they entail about speaker commitments to 'the way the world is.'

§5.2 then applies this metasemantic picture to experiential predicates, showing how speakers constrain their behavior in construing the truth-conditional content of these predicates based on their own experiential dispositions, and demonstrating how this gives rise to uses of these and other predicates that reflect speaker opinion. It is first demonstrated how differing speaker behavior with respect to experiential predicates and autocentric evaluation can be modeled given the above-introduced machinery. This behavior is then traced to its source, both in the compositional semantics and the pragmatics: the conclusion is that autocentricity is the result of the felicity conditions on generic quantification over experiential states being construed in line with norms on the accuracy of one's own direct evidence. It's then shown that the resulting picture allows speakers to express robust descriptive and normative attitudes using these predicates, in a way that does not require them to commit to a metaphysical picture on which experiential properties exist in the world, despite the fact that their semantics contains no relativization to an experiencer.

§5.3 closes this work by showing how this metasemantic machinery is reflected substantively in the compositional semantics itself, and how the special role of experience due to direct evidence is semantically encoded. Subjective attitudes and certain sorts of dative marking are discussed, and cast as hyperintensional operators that are specially sensitive to the metasemantic machinery outlined here. Finally, some some methodological comments are offered on the nature of linguistic subjectivity, and how it ought to be construed in line

with the semantic and metasemantic machinery outlined in this chapter.

5.1 The metasemantics of predicates

This section offers a sketch of the metasemantics of predicates generally, arguing that they do not denote single properties *simpliciter* in virtue of their linguistic meaning, but must be construed truth-conditionally more or less precisely by the behaviors of speakers in that community. This is shown using the exemplar of an uncontroversially non-experiential predicate, *athlete*, which is not typically considered evaluative in any way that would require the relativization of its extension to non-world parameters. The lessons drawn for predicates generally are then applied to the case of experiential predicates in §5.2.

§5.1.1 characterizes the semantic underdetermination of predicates generally, and hints at what sort of move would have to be made in order to accommodate it. §5.1.2 offers some new theoretical machinery to handle the metasemantics of predicates, that takes this semantic underdetemrination into account – this requires a grammar with a metasemantic component, that deals with meaning at a finer level of grain than ordinary truth conditions, and addresses how these truth conditions arise in virtue of speaker behavior. §5.1.3 then demonstrates that from this machinery it follows that speaker attitudes towards the meanings of expressions are systematically sensitive both to facts about the world and behaviors involved in construing those meanings truth-conditionally. This allows the descriptive and behavioral components of these attitudes, and how they interact, to be characterized.

5.1.1 Semantic underdetermination: athlete

When authors address the idea of not relativizing the extension of experiential predicates to an experiential standard, they tend to reject it, or draw unpalatable consequences from it, on grounds that it would mean that these predicates track some 'fact of the matter' as to whether something has an experiential property, presumably in the way that non-

experiential, non-evaluative predicates in general do with the properties they denote. This is then taken to be difficult to square with speakers' autocentric evaluation of these predicates.

There is typically a slide from a semantic thesis about what parameters a predicate is or is not relativized to, into a claim about some such 'fact of the matter.' Thus Lasersohn (2005: 654-655) begins by entertaining that such predicates are 'basic, with no indexing, implicit arguments, or relativization,' but is soon comparing this thesis to the idea that 'there is a definite fact of the matter' with respect to the application of the predicates, and claiming that the difficulty that would attend knowing this fact is in tension with speakers being able to 'speak with authority' as to e.g. whether something is tasty after tasting it. And MacFarlane (2014: 2-3) starts by entertaining the idea that '(a) "tasty" is true of some things, false of others, and (b) whether "tasty" is true or false of a thing, on a particular occasion of use, does not depend on the idiosyncratic tastes of the speaker, assessor, or anyone else,' but this soon morphs into the claim that 'there is a "fact of the matter" about whether a thing is tasty in the way that there is a fact of the matter about whether it is red or deciduous or acidic,' which is then taken to have various bad epistemic consequences (cf. §5.2.3).

Behind these moves is a certain metasemantic picture about the way that predicates in general have their truth-conditional content. For there to be a fact of the matter as to whether an individual is tasty is, presumably, for the state of the world to be a necessary and sufficient determinant of whether that individual is tasty. But then, for an individual to be tasty is the same thing as for the predicate *tasty*, as actually used, to truly apply to it: and so the state of the world, and it alone, determines what *tasty* as actually used is truly or falsely predicated of (or perhaps neither, e.g. for individuals that lack a taste), for an arbitrary individual. But then, it must be that the denotation of *tasty* determines a mapping from the state of the world to individuals to truth values, for any individual – and this is just what it is for *tasty* to denote a single property of individuals, viz. some single function of type $\langle s, \langle e, t \rangle \rangle$.

The metasemantic picture is therefore that *tasty*, and presumably other experiential predicates, denote a single property in the language unless they are relativized to some parameter other than the world – and since this is an unacceptable result, there must be some other such parameter. Presumably, non-experiential, non-evaluative predicates do actually denote a single property, and so the 'fact of the matter' result there is welcome. Experiential predicates are therefore special in their not denoting a single property, which is why they call for special relativization by the introduction of a new parameter.

This is not a plausible metasemantics of predicates. It is true that experiential predicates do not denote single properties *simpliciter* in virtue of linguistic practice, but neither does almost any predicate in any natural language. The truth-conditional content of predicates generally is *semantically underdetermined* – given a predicate, it is in general not possible to single out the one property that this predicate denotes, in a linguistic community of any size, or even with respect to a single speaker (with perhaps a couple of exceptions: cf. fn. 11). And the semanticist generally doesn't attempt this.

Rather, predicates have vague conditions of application, with individuals acting as more or less canonical instantiators of whatever property or cluster of properties the predicate is construed as denoting. Can *dog* apply to a coyote – what about to a coydog? Does *wooden* apply to things made of particle board? These are not supposed to be evaluative or experiential predicates – yet the question of *which* property they denote is no more settled than for the case of experiential predicates, and speakers can differ, and even contradict one another, in their construal of how these predicates apply. The ability to competently construe a single expression as denoting distinct properties is therefore not a reason to introduce relativization to a parameter specifically for experiential predicates, or other predicates that purportedly pertain to things about which there is no unrelativized 'fact of the matter.'

In order say what the meaning of a predicate actually is, some answer must be given to what Pérez Carballo (2014) calls the 'hermeneutic question:' what is it for a certain lexical item, and in particular a lexical predicate, to have a certain semantic denotation?

The implicit answer given to this question by the treatments cited above is that to assign a property denotation to a predicate is to take that predicate to denote some single property in the language. But the truth-conditional content of a lexical item is plausibly constituted by speaker behavior: for a predicate to denote a certain property in a linguistic community just is for speakers in that community to treat it as denoting that property, as construed by their relevant truth-conditional semantic behavior (cf. Caie 2014 on the 'fit' of a semantic theory to a linguistic community). This means, roughly, that speakers take a predicate to apply to an individual just in case they take that individual to have that very property.

But if denotations of predicates are constituted by speaker behavior in this way, it follows that if speaker behavior is not fully determinate as to what property is denoted by a predicate, then there simply is no single property that predicate denotes in the language, and so this metasemantic picture cannot be accurate. Rather, speakers behaviorally construe predicates as having truth-conditional content, and their behavior is not fully determinate as to which properties they construe each predicate as having, resulting in semantic underdetermination.[4]

A vivid example of semantic underdetermination is provided by Ludlow (2008: 118), who reports a debate on sports talk radio as to whether the racehorse Secretariat can legitimately be considered an athlete.[5] Some callers were upset about the inclusion of Secretariat on a

4. Two other possibilities one might consider, to save the idea that predicates have fixed truth-conditional contents assigned to them in the language, are (i) to supervaluate with respect to speaker behavior, so that predicates are true or false of individuals according to their properties *simpliciter* just where a linguistic community's behaviors perfectly converge, and neither true nor false otherwise; or (ii) to hold a quasi-epistemicist position, that there really is just one truth-conditional denotation for each predicate, though there is an in-principle inability to know which denotation exactly that is. There is no point in discounting these positions – they are just not relevant to the issues at hand. A supervaluationist theory will simply fail to yield semantic values when speaker behaviors diverge, and so is too coarse-grained to have anything to say about precisely the cases under consideration; and a quasi-epistemicist theory will make the true semantics of expressions in principle not discernible from any empirically observable behavior of speakers (cf. Williamson 1994: §7.5 for complications on the connection between meaning and use on an epistemicist theory), making it irrelevant to empirical semantics even if true. A third path is proposed in Caie (2014), according to which it is metaphysically indeterminate what the exact semantic features of a natural language are – moving to a deeper metasemantic level of explanation in terms of speaker behavior makes this move implausible.

5. Although the show's exact date and audio aren't publicly available, this debate took place on 'Mike and the Mad Dog,' on WFAN Sports Radio, 101.9 FM, in response to the December 27, 1999 issue of *Sports Illustrated*, which included the article, 'The Master List: The 50 greatest sports figures from each of the 50 states.' In this article, Secretariat was listed in third place for Virginia.

list of esteemed sports figures, claiming that only humans were properly athletes, and so only they deserved to be on such a list – the host defended the opposing position, that a racehorse could be an athlete, and so the inclusion was justified.

If *athlete* really does denote some single property in the language, and the speakers in question were sincere and fully-informed as to the relevant facts about the world, including who Secretariat was and what he did, on its surface such a debate becomes unintelligible – and in fact it has just the sort of character that debates concerning the application of experiential or evaluative predicates purportedly give rise to in a special way. What exactly is the denotation of *athlete*? A first gloss, past which the compositional semanticist never truly goes (cf. Glanzberg 2014 for comments on this fact), is as follows.

(359) $[\![\text{athlete}]\!]^w = \lambda x_e.athlete'(w)(x)$

This denotation isolates a certain formal, type-theoretic feature of *athlete*, that it denotes a property of individuals. But which property – what does '$athlete'$' actually denote in the metalanguage? It does no good to say 'the property of being an athlete,' since this description allows the semanticist to isolate a single property no better than the metalanguage predicate, or the natural language predicate itself. A property, in the relevant sense for this compositional semantics, is a function from worlds to individuals to truth values – which one is denoted here, and does it map Secretariat to true at the actual world or not?

The semanticist could answer this question by in effect voicing an opinion as to whether Secretariat is an athlete or not – a similarly opinionated answer could be provided as to whether e.g. licorice is tasty, to settle what property *tasty* denotes. But the semanticist's opinion isn't what is at issue, but rather a description of some fact about English, and in English usage as a whole, there simply is no one property that speakers treat *athlete* as denoting. The debate over Secretariat is just one scenario in which this fact becomes apparent: it becomes equally apparent in considering what sort of athletic events one must participate in to qualify as an athlete (are intramurals enough? or a single minor-league game?), whether some level of physical skill is required, etc.

One might respond to this by bringing predicates like *athlete* in line with the relativist treatment of experiential predicates, and taking them to be evaluated with respect to some intensional parameter that determines the precise truth-conditional content of the predicate.[6] Call it p, for 'precisification' – and so where $athlete'$ is a function from precisifications to worlds to individuals to truth values, and $athlete'(p)$ is the intension of a property, viz. a function of type $\langle s, \langle e, t\rangle\rangle$:

(360) $[\![\text{athlete}]\!]^{w,p} = \lambda x_e.athlete'(p)(w)(x)$

But a relativist treatment cannot end here: it requires also some principle that guides speakers in assessing for truth correctly and properly, by providing a value for the intensional parameter p (cf. §1.2). And so on par with e.g. (85) and (86) for experiential predicates, something like the following might be proposed (and *mutatis mutandis* for falsity), where 'ϕ' is a variable over functions from precisifications to traditional propositions – cf. the notion of a 'relativist-proposition' in §1.2.

(361) a. ϕ is true *simpliciter* at w iff $\forall p[\phi(p)(w)]$.

b. ϕ is false *simpliciter* at w iff $\forall p[\neg\phi(p)(w)]$.

(362) a. ϕ is correctly assessed as true in c' iff ϕ is true *simpliciter* at $w_{c'}$.

b. ϕ is incorrectly assessed as true in c' iff ϕ is false *simpliciter* at $w_{c'}$.

Then where the set of precisifications p over which (361) universally quantifies are those licensed by the relevant linguistic community as a whole, a relativist-proposition is true or false *simpliciter* just where all these precisifications converge on a truth-value. Where the

6. The text will not consider the possibility that lexical predicates like *athlete* are to be given a contextualist treatment, with some parameter at the context of utterance acting as a precisification of truth-conditional content, perhaps in the manner of Lewis' (1980: §3) 'standards of precision.' Such a view seems to be what Sundell (2011) and Barker (2013) are getting at for the case of experiential predicates, but the interested reader can confirm that such a treatment just recapitulates all the contextualist's problems with cross-contextual truth assessment and attitude reports covered in Chapter 1 with respect to experiential predicates. What is needed, as before, is a notion of e.g. believing that an individual is an athlete *simpliciter*, not relative to some parameter set at a context of utterance, if speakers' cross-contextual behavior is to be captured.

precisifications diverge, there is no 'fact of the matter' – and perhaps Secretariat falls within this extension gap, so there is no fact of the matter as to whether he is an athlete.

But still, even where linguistic behavior in the community is not uniform, speakers can and do adopt conflicting precisifications, and assess the truth of the relevant propositions accordingly – this is why the debate over Secretariat is possible. In these cases, opinionated speakers must choose which precisifications to adopt for purposes of truth assessment, since the extensions of truth-apt expressions, such as *Secretariat is an athlete*, will non-trivially depend on this. In §1.2, the corresponding fact with experiential predicates was cast in terms of proper truth assessment, i.e. a pragmatic norm as to how to supply a value for the experiencer parameter.

The problem is that unlike with the case of experiential predicates, for which proper truth assessment can be defined autocentrically in terms of the assessor's own experiential dispositions (cf. (86b)), there is apparently no pragmatic principle compelling speakers to assess *athlete* in any particular way, at least within the vague boundaries that linguistic convention countenances. There is, in other words, no way to offer a clause on proper truth-assessment with respect to precisifications, no way to fill in the '?' in the following.

(363) ϕ is properly assessed as true in c' iff $\phi(?)(w_{c'})$
(and is improperly assessed as true otherwise).

It does no good to say that an assessment is proper where the precisification in question is in line with the assessor's relevant opinions, or usage of words, dispositions to judge that a predicate applies, etc., since these things are constituted by which precisifications the assessor employs in the first place, making any such explanation viciously circular. The entire reason the autocentric treatment of experiential predicates is plausible for the relativist is that this kind of circular treatment of perspectivality in terms of opinion can be avoided, by appealing to facts about the world independent of the assessor's opinion, viz. experiential dispositions of the assessor: cf. Lasersohn (2017: 93-94) for comments.[7]

7. Though Kölbel (2002: 103-104) makes no such requirement, and apparently allows that norms on belief

What one wants to say is that in for precisifications, there is a highly restricted space, within which no independent pragmatic norm governs how they are adopted: it is simply a fact that speakers must behaviorally construe truth-conditional contents in some way or other, and the way that they do this constitutes which precisifications they are disposed to employ. This is close to what will be implemented in the following sections, but the drifting away from a relativist account in spirit is clear: what is at stake is general underdetermination of meaning, and the fact that speakers are ubiquitously non-uniform in their behavioral construal of truth conditions, not special parameters required for experiential predicates sensitive to the properties of assessors.

The lesson to be drawn from this is that the construal by different speakers of the same predicate as having distinct truth-conditional contents is the norm, not an exception for experiential predicates, and so there is no special relativization to an experiencer to be expected there. It is true that experiential predicates display some special behavior reflecting the experiential dispositions of assessors, in virtue of their experiential semantics (cf. §5.2), but this behavior makes use of the same core features of semantic underdetermination that the vast majority of predicates share. In resolving truth-conditional content, the appropriate target is therefore not primarily experience, evaluation, or opinion, but the plain metasemantic fact that language use is semantically underdetermined, because speaker behavior is not perfectly consistent with a single pattern of usage.

5.1.2 Semantic alternatives and hyperintensions

A semantically underdetermined predicate like *athlete* denotes a single property, i.e. a mapping from worlds to individuals to truth values, only when construed truth-conditionally in some perfectly precise way by speakers. The resolution of which property *athlete* denotes is handled by a parameter, as in a relativist semantics. But this parameter is of wider application than those that the above-mentioned accounts employ, and does something distinct in

relative to a perspective are in some sense circularly defined.

spirit. It does not pertain specially to matters of opinion or experience, and as will be shown in §5.2, it derives the behavior of the latter sorts of discourse as a special case. What it does is resolve the semantic underdetermination of linguistic expressions generally, by providing a single way of construing the truth-conditional content of those expressions.

In other words, the point of this parameter is, given an expression of the language, to assign to it a truth-conditional denotation. This is precisely the traditional role of a *model*, which includes an interpretation function. And so the parameter in question is the model: at a model, a linguistic expression receives an intension. Let $[\![\alpha]\!]$ be the *meaning* of an expression α. This meaning is a function from models $\mathcal{M}$ to intensions, so that $[\![\alpha]\!]^{\mathcal{M}} = [\![\alpha]\!](\mathcal{M})$ is the intension or truth-conditional content of α at $\mathcal{M}$, and as usual $[\![\alpha]\!]^{\mathcal{M},w} = [\![\alpha]\!](\mathcal{M})(w)$ is the extension of α at $\mathcal{M}$ and w. A meaning, as in the sense of Lewis (1970), is more fine-grained than an intension, since expressions with distinct meanings can be cointensive at the same model.

The extension of *athlete* is then recast as follows to incorporate model-sensitivity. The metalanguage predicate '$athlete'_{\mathcal{M}}$,' for any $\mathcal{M}$, is of type $\langle s, \langle e, t\rangle\rangle$, i.e. the intension of a property of individuals. '$athlete'$' can in turn be taken to be type $\langle m, \langle s, \langle e, t\rangle\rangle\rangle$, i.e. a function from models to worlds to individuals to truth values, with a symbol denoting a model appearing as a subscript to denote its first argument.[8]

(364) $[\![\text{athlete}]\!]^{\mathcal{M},w} = \lambda x_e.athlete'_{\mathcal{M}}(w)(x)$

And so *athlete*, relative to a model $\mathcal{M}$, denotes whichever property $athlete'$ maps $\mathcal{M}$ to. Semantic underdetermination therefore corresponds to model-sensitivity in an expression's denotation. In what follows, only model-sensitivity that is presently relevant is included in denotations, e.g. for those underdetermined predicates like *athlete* being analyzed. Elsewhere, where semantic underdetermination isn't relevant to the point, model-sensitivity that

8. It might seem odd to have a semantic type for models – are they in a domain subsumed under a higher-order model? This is a terminological matter, but the idea is to show that semantics practiced with this metasemantic sensitivity is not 'model-theoretic' at all: it is simply describing the behavior of speakers in the world. The domains appealed to are therefore not part of a model.

would exist in a complete account is idealized away. However, given the comments in §5.1.1, that nearly all lexical predicates in a natural language are semantically underdetermined, it follows that nearly all of them are actually model-sensitive in some way.

This raises the question of what, in concrete empirical terms, a model is supposed to be. Putting the pieces together: in §5.1.1, it was claimed that the behavior of speakers determines the truth-conditional content of an expression; but here, that role is instantiated by a model; it follows that a model is a *semantically relevant pattern of behavior* that a speaker is capable of enacting. Roughly, a speaker disposed to treat an expression α as denoting some truth-conditional content P is enacting a pattern of behavior (model) $\mathcal{M}$, such that P is the intension of α at $\mathcal{M}$. And so, for the denotation in (364): *athlete*, as construed by a pattern of behavior $\mathcal{M}$, denotes whichever property P that $\mathcal{M}$ treats *athlete* as denoting: and at w, this property is true of any individual x that has P.

The models that are available in a linguistic community are tightly constrained, and it is part of speakers' lexical semantic competence not to employ models that deviate wildly from the norm. And so while formally $athlete'$ is in principle capable of mapping some model to any property at all, this fact only takes into account the skeletal, type-theoretic and formal constraints of the compositional semantics. In addition to this, a robust lexical semantics narrows down the set of employed models in a linguistic community to a small cluster, which decides within some vague range what it is possible to construe *athlete* as denoting without being incompetent. Within this range, there is the sort of variation in the employment of models hinted at in §5.1.1.[9]

9. This notion of semantic competence is descriptive, not normative: it requires observing which patterns of behavior speakers actually employ. It also may require describing norms that speakers in fact obey, e.g. when they take the meaning of an expression to be governed by some authority that they in principle attempt to align with, even where actual use of the expression may be out of sync with it. That the lexical semantics' role is to narrow the space of models employed in a linguistic community answers the criticisms of Lepore (1983), to the effect that model-theoretic semantics cannot provide a specific enough account of the meanings employed in an actual language. Montague (1970) makes speculative reference to 'the actual model' of a language: it should be clear from the above that this is not the right target, but rather the description of speakers' actual behavior, which employs a vague cluster of models: cf. Lewis' (1969: 201) idea of 'a tight cluster of very similar possible languages' (though this is slightly misleading, since the present casts this as similar truth-conditional treatments of the same language). At the brass tacks, what remains is to describe

Since a model assigns, among other things, a precise property to each lexical predicate, it is a maximally determinate pattern of behavior: the behavior of a speaker who really did employ a single model $\mathcal{M}$ would treat each expression of the language with no semantic underdetermination at all. This is not in general how speakers behave – they employ a vague range of candidates, and are more or less consistently disposed to treating predicates as denoting certain properties. To capture this, the notion of a speaker's *semantic alternatives* can be introduced. The set of x's semantic alternatives at w, $Sem_{x,w}$, is the set of patterns of behavior that are consistent with x's behavior at w.

(365) $Sem_{x,w} := \{\mathcal{M} : \mathcal{M}$ is consistent with x's semantically relevant behavior at $w\}$

In the same way that speakers have multiple doxastic alternatives when they are not completely opinionated about the state of the world, speakers have multiple semantic alternatives when their behavior is not perfectly determined as to how it construes the truth-conditional content of linguistic expressions.

Semantic alternatives do not reflect metalinguistic beliefs, e.g. as to a speaker's opinion of what a word means (though they can in principle be adopted on the basis of such beliefs). They do not reflect beliefs about the world of any kind. Rather, they reflect the behaviors speakers habitually enact with respect to truth-conditional language, *given* that they have certain beliefs about the world. In so doing, speakers behaviorally construe linguistic expressions as having certain truth-conditional content.

To characterize this, a notion of taking the meaning of an expression to be true is required. Let Φ be an object of type $\langle m, \langle s, t\rangle\rangle$, viz. a mapping from patterns of behavior to traditional propositions. A speaker takes Φ to be true just in case that speaker behaves ingenuously with respect to truth-sensitive linguistic norms in a way that requires them to commit to its truth. Thus if α's meaning is Φ, then a speaker who treats Φ as true will have no truth-related qualms about ingenuously making assertions using α, and will not not reject

---

the behavioral microvariation: this is what the employment of models amounts to.

or retract assertions made using α on grounds that what α says is false, and so on. Treating a meaning as true is a pragmatic notion: it is intelligible only with respect to the systematic use of linguistic expressions.

The nature of semantic alternatives can then be defined in terms of a speaker's truth-sensitive linguistic behavior, combined with their beliefs about the world. The latter are defined independently of linguistic usage, as usual in terms of doxastic alternatives. Speakers then take a truth-apt meaning Φ to be true just in case all their doxastic and semantic alternatives verify it. Conversely, they take such a Φ to be false where all their doxastic and semantic alternatives falsify it.

(366) a. x takes Φ to be true at w iff $\forall w' \in Dox_{x,w}, \mathcal{M} \in Sem_{x,w}[\Phi(\mathcal{M})(w')]$

b. x takes Φ to be false at w iff $\forall w' \in Dox_{x,w}, \mathcal{M} \in Sem_{x,w}[\neg\Phi(\mathcal{M})(w')]$

These equivalences are taken to be *constitutive* of what semantic alternatives are: if, given a certain set of beliefs about the world, a speaker is inclined to treat some Φ as true (false), then that speaker's semantic alternatives are such that (366) holds. Semantic alternatives, in other words, track how speakers are disposed to treat the truth and falsity of the meanings of linguistic expressions according to pragmatic norms, given their beliefs about the world. This is just what it is to behaviorally construe the meanings of expressions as having a certain truth-conditional content.

To take an example, recall the debate from §5.1.1, over whether Secretariat is an athlete. The debate hinged in large part on whether non-humans could be athletes, and thus whether the predicate *athlete* could truly apply to non-humans like horses. Suppose that Alfonse believes all the true things about Secretariat relevant to the debate, including that he competes in athletic events, i.e. that Secretariat has the property P_1.[10]

(367) $\forall w' \in Dox_{a,w}[P_1(w')(s)]$

10. This notation is an idealization: there is no one property of competing in athletic events. This property has been described with more English words, themselves semantically underdetermined, and so which pick out no one property *qua* function from worlds to individuals to truth values.

Suppose further that Alfonse's semantic alternatives at w are fixed, so that *athlete* denotes that same single property P_1 for all of them. This means that Alfonse's behavior treats *athlete* as uniformly denoting the property of competing in athletic events – it makes no reference to being human, and so Alfonse takes the predicate to be applicable to Secretariat, given his beliefs.[11]

(368) $\forall \mathcal{M} \in Sem_{a,w}[[\![\text{athlete}]\!]^{\mathcal{M}} = P_1]$

From (367), (368), and the definition in (366), it follows that Alfonse treats the meaning of *Secretariat is an athlete* as expressing a truth.

(369) $\forall w' \in Dox_{a,w}, \mathcal{M} \in Sem_{a,w}[athlete'_{\mathcal{M}}(w')(s)]$
$= \forall w' \in Dox_{a,w}[P_1(w')(s)]$
$= 1$

If instead Alfonse's beliefs about the world were such that he believed that Secretariat did not compete in athletic events (did not have P_1), then given the same semantic alternatives, Alfonse would instead treat the meaning of *Secretariat is an athlete* as expressing a falsehood, as the reader can confirm.

Speakers' behavior does not always treat a truth-apt meaning consistently as true or false, even where their relevant beliefs about the world are entirely decided: in such cases, a speaker is in a state of *semantic indecision* with respect to a meaning. Suppose for instance that Bethany's semantic alternatives assign at least two possible meanings to *athlete*: her behavior is consistent with *athlete* denoting P_1, as Alfonse's alternatives were decided on, or with it denoting the property of having P_1 along with P_2, the property of being human.

(370) a. $\exists \mathcal{M} \in Sem_{b,w}[[\![\text{athlete}]\!]^{\mathcal{M}} = P_1]$
b. $\exists \mathcal{M}' \in Sem_{b,w}[[\![\text{athlete}]\!]^{\mathcal{M}'} = \lambda w'_s.\lambda x_e.P_1(w')(x) \wedge P_2(w')(x)]$

11. This too is an idealization, in that an actual speaker will seldom have their behavior settle on construing the meaning of a predicate as a single property. There may be some rare exceptions, e.g. with *even*, as predicated of numbers.

In other words, Bethany's semantically relevant behavior is undecided as to whether *athlete* denotes the property of just competing in athletic events, or the stricter property of both competing in athletic events and being human.

With respect to many individuals, Bethany's beliefs regarding these properties combined with her semantic alternatives still decide her truth-conditional behaviors: she takes any individual that lacks P_1 not to be an athlete, and any individual that has both P_1 and P_2 to be an athlete. But Secretariat, who has P_1 but not P_2, appears as a 'borderline case' with respect to her behavior. Supposing that she truly believes that Secretariat has P_1 but not P_2 (371), she treats the meaning of *Secretariat is an athlete* as neither true nor false (372).

(371) $\forall w' \in Dox_{b,w}[P_1(w')(s) \wedge \neg P_2(w')(s)]$

(372) a. $\forall w' \in Dox_{b,w}, \mathcal{M} \in Sem_{b,w}[athlete'_{\mathcal{M}}(w')(s)]$
only if $\forall w' \in Dox_{b,w}[P_1(w')(s) \wedge P_2(w')(s)]$
$= 0$

b. $\forall w' \in Dox_{b,w}, \mathcal{M} \in Sem_{b,w}[\neg athlete'_{\mathcal{M}}(w')(s)]$
only if $\forall w' \in Dox_{b,w}[\neg P_1(w')(s)]$
$= 0$

And so where Bethany's beliefs and semantically relevant behaviors are like this, she may behave inconsistently as to whether Secretariat is an athlete, taking the relevant expressions to be true or false in turn, or simply remain indifferent and not commit to the truth or falsity of the relevant truth-apt expressions: cf. Caie (2014: 65ff.) on the 'interval of silence.' This notion of indecision is a gradable one, with speakers being more or less consistent in their treatments of truth and falsity: the notions defined in (366) reflect limit cases, where the speakers' behavior is consistent enough that they can be reported truly as construing a certain meaning as expressing a truth or a falsehood.[12]

12. The machinery proposed here already allows for the formal characterization of a certain amount of this gradability. Let $Sem_{x,w,\alpha}$ be the semantic alternatives for the predicate α that x has at w, such that $Sem_{x,w,\alpha} = \{P : \exists \mathcal{M} \in Sem_{x,w}[[\![\alpha]\!]^{\mathcal{M}} = P]\}$. Then $Sem_{x,w,\alpha}$ constitutes an ordering source on

This notion of taking a meaning to be true can be abbreviated: '$believe''(w)(\Phi)(x)$' is to be read, 'x takes Φ to be true at w.' Belief reports then work by having a doxastic verb like *think* compose with the meaning of a subordinate clause Φ and an agent x, and reporting that x takes Φ to be true in this sense. And so *think* composes not with propositions, but with functions from patterns of behavior to propositions.

(373) $[\![\text{think}]\!]^{\mathcal{M},w} = \lambda\Phi_{m,st}\lambda x_e.believe''(w)(\Phi)(x)$

And the composition for *Alfonse thinks that Secretariat is an athlete* is as follows. It is true just in case, given what Alfonse believes about the world, his semantically relevant behaviors are such that he treats the meaning of *that Secretariat is an athlete* to be true (which in turn implicates him in all the truth-relevant practices mentioned above).[13]

(374) $[\![\text{Alfonse [thinks [that [Secretariat [is [an athlete]]]]]}]\!]^{\mathcal{M},w} =$
$[\![\text{think}]\!]^{\mathcal{M},w}([\![\text{that Secretariat is an athlete}]\!])([\![\text{Alfonse}]\!]^{\mathcal{M},w})$
$= believe''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.athlete'_{\mathcal{M}'}(w')(s))(a)$
$= \forall w' \in Dox_{a,w}, \mathcal{M}' \in Sem_{a,w}[athlete'_{\mathcal{M}'}(w')(s)]$

Belief reports are therefore sensitive both to an agent's doxastic and semantic alternatives. They compose with functions from models to intensions, and so the compositional semantics

individuals, ranking them according to which of these properties they have at a world. So $\leq_{x,w,\alpha,w'}$ can be a preorder on individuals at w' according to x's alternatives for α at w. '$y \leq_{x,w,\alpha,w'} z$' can be read, 'y is as canonical an α at w' as z, according to x's semantic alternatives at w.' Then $y \leq_{x,w,\alpha,w'} z$ iff$_{def}$ $\forall P \in Sem_{x,w,\alpha}[P(w')(z) \rightarrow P(w')(y)]$. That is, y is as canonical an α as z at w', where y at w' has all the properties imposed by x's semantic alternatives at w that z has, and possibly more. This would then correspond to x at w being more disposed to treat y as an α than z at w', even in the case where both y and z are 'borderline cases' of an α according to these alternatives.

13. It follows from this sort of treatment of belief reports that agents can sometimes, where their semantic alternatives are very strange, be reported as believing things that don't align at all with the normal way of treating the meanings of words. Thus if Alfonse is just learning English and thinks that *pig* is used to refer to horses, and calls Secretariat a pig on this basis, the report *Alfonse thinks that Secretariat is a pig* comes out true on the hyperintensional reading, despite the fact that Alfonse's beliefs about the world attribute no pig-like qualities to Secretariat – this is apparently the right result. There is also plausibly a *false* reading of this report, on which one might think that no, Alfonse doesn't really think Secretariat is a pig, but is just using a word wrong: here, the proposition is read '*de re*,' outside the scope of the hyperintensional attitude, and so is concerned only with what Alfonse thinks about the world. Presumably a scopal mechanism will generate both readings. *De re* readings are possibly required for beliefs attributed no non-linguistic agents, who plausibly cannot have genuinely hyperintensional attitudes.

is capable of quantifying over models in these contexts. Belief reports are, in other words, *hyperintensional*: they operate at a finer grain of meaning than merely intensional contexts. This fact has been long known (cf. Cresswell 1975), and the ability to quantify over models, i.e. patterns of behavior, is one way of formally realizing it. Where models $\mathcal{M}$ are of type m, the type $\langle m, \langle s, \tau \rangle\rangle$ is therefore the *hyperintensional type* of an expression of extensional type τ, and $[\![\alpha]\!]$, the meaning of an expression α *simpliciter*, is its hyperintension.

In giving a full treatment of these attitudes, and the linguistic phenomena surrounding them, a hyperintensional fragment of the grammar is required, which subsumes its merely intensional counterpart as a special case that offers meanings relative to a single model. Where hyperintensional concerns are irrelevant, as they often are to the compositional semanticist, the effect of models can be ignored, just as intensional concerns are ignored when confining attention to extensional phenomena.[14]

In fact, all attitude verbs that imply an assessment of truth on the part of the subject are hyperintensional, and so can be given a treatment analogous to *think*'s. This includes factives like *realize* (375a),[15] as well as verbs that require an alignment or lack of alignment of beliefs between agents, like *agree* (375b) and *disagree* (375c). Denotations for these expressions can be provided by hyperintensionalizing the denotations offered for them previously (cf. (91)), and importing the present notion of belief, which is sensitive to semantic alternatives.

(375) a. $[\![\text{realize}]\!]^{\mathcal{M},w} = \lambda\Phi_{m,st} : \Phi(\mathcal{M})(w).\lambda x_e.\mathit{believe}''(w)(\Phi)(x)$

b. $[\![\text{agree}]\!]^{\mathcal{M},w} = \lambda x_e.\lambda\Phi_{m,st}.\lambda y_e.\mathit{believe}''(w)(\Phi)(x) \wedge \mathit{believe}''(w)(\Phi)(y)$

14. The hyperintensional nature of the grammar is already present implicitly in merely intensional fragments that semanticists typically offer since, as Heim & Kratzer (1998: §2.1.2) note for intensionality, a method must be chosen for specifying a semantic value, and this *shows* the hyperintensional nature of predicates like *horse* via the use of symbols like '*horse*'' – the semanticist must construe this behaviorally as denoting some range of properties or other, though an actual property is never specified. But leaving the hyperintensional machinery transparent often leads to theoretical blindness, e.g. thinking that there really is some one property there being denoted (cf. §5.1.1). The hyperintensionality of attitudes, especially subjective attitudes (cf. §5.3.1) cannot be treated fully without explicit appeal to hyperintensions, though as the last four chapters show, an intensional fragment is good enough for most purposes, including pragmatic ones involving differences in truth assessment (cf. §1.3.2).

15. As noted in Ch. 1, fn. 21, this sort of denotation is ultimately inadequate for factives like *realize*, but for reasons orthogonal to the matter at hand.

c. $[\![\text{disagree}]\!]^{\mathcal{M},w} = \lambda x_e.\lambda\Phi_{m,st}.\lambda y_e.believe''(w)(\Phi)(x) \wedge$
$believe''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.\neg\Phi(\mathcal{M}')(w'))(y)$

And so *realize* means the same as *think*, except that it presupposes that the hyperintension of the embedded clause has an intension true at the world of evaluation. The speaker of an attitude report involving *realize* becomes committed to this, as does anyone who accepts that report. This in turn happens felicitously, given the characterization of semantic alternatives, just in case those uttering or accepting the report verify the hyperintension in all of their semantic and doxastic alternatives, i.e. just in case they believe the hyperintension.[16]

Presupposition, and commitment to truth through utterance and acceptance, thus works unsurprisingly: commitment to truth is hyperintensional in general.[17] *Agree-* and *disagree-*reports are then just true where the two agents in question have aligning, or opposing, beliefs toward the hyperintension of the complement clause respectively: in the first case, they both believe it, while in the second, the first believes it while the second believes its negation.

5.1.3 Descriptive and behavioral attitudes

Whether agents hold hyperintensional attitudes depends not only on their beliefs about the world, but also on their behavioral dispositions toward meanings. The commitments that these attitudes impose on their agents can be described at both these levels.

16. In §1.1.3, an important question was how this acceptance of the presupposition of a factive could be linked to commitments to (dispositions to) having certain experiences, where experiential predications are concerned. This will be made clear in §5.2.1, where it is shown that due to the semantics of experiential predicates, one's semantic alternatives are typically constrained by the experiences one is disposed to have: and so to accept a hyperintension as true is typically to commit oneself to having those dispositions.

17. The same is true of conversational dynamics generally: rejection, retraction, etc. work as expected given the way that hyperintensions are evaluated according to one's semantic alternatives. A typical story of informational exchange in conversation along the lines of Stalnaker (1978) could be given here, involving hyperintensional update of a common ground that, in virtue of speakers' semantic alternatives, simultaneously tracks discourse commitments in terms of beliefs about the world and commitments to treating the meaning of words in a certain way. The text will forego this, since similar treatments of this phenomenon are already provided e.g. in Barker (2002, 2013) and Fleisher (2013): cf. also Stephenson (2007a: §5.1); Lasersohn (2017: §9.2); and Coppock (2018: §4), for comments on making the Stalnakerian framework multidimensional. The move from their frameworks to the present one is trivial. A traditional Slalnakerian dynamic pragmatics is thus a (merely) intensional fragment of its hyperintensional counterpart: it models informational exchange in an ideal situation in which all conversational participants' semantic alternatives perfectly align.

An agent's commitments to how the world is in virtue of holding a belief can be recovered by factoring out semantic alternatives. Let $\Phi_{x,w}$ be the *descriptive content* of a hyperintension Φ of type $\langle m, \langle s, t\rangle\rangle$, relative to x at w. This is the proposition of type $\langle s, t\rangle$ that Φ expresses, as treated by x's behavior at w, i.e. what the truth-conditional content of Φ is according to x's semantic alternatives.[18]

(376) $\Phi_{x,w} := \lambda w'_s.\forall \mathcal{M} \in Sem_{x,w}[\Phi(\mathcal{M})(w')]$

Where x believes Φ at w, the descriptive content of this belief is just $\Phi_{x,w}$. *Descriptive error*, where the agent takes the world to be a way that it isn't, results where the descriptive content of the belief is a proposition false at the world of belief.

(377) x makes a descriptive error in believing Φ at w iff:

a. $believe''(w)(\Phi)(x)$

b. $\neg\Phi_{x,w}(w)$

Since belief reports are hyperintensional contexts, it isn't possible to tell whether an agent holds a reported belief in descriptive error, unless one knows what that agent's semantic alternatives are, since only then can one know the descriptive content of that belief.

To illustrate, suppose as in §5.1.2 that at w, Secretariat has P_1 but not P_2: he competes in athletic events, but is not human (378a). Suppose further that Alfonse believes both these things truly of Secretariat (378b).

(378) a. $P_1(w)(s) \wedge \neg P_2(w)(s)$

b. $\forall w' \in Dox_{a,w}[P_1(w')(s) \wedge \neg P_2(w')(s)]$

Alfonse does not make a descriptive error in holding these beliefs, since the belief state represented in (378b) does not exclude w from his doxastic alternatives. However, a report of

18. To accommodate x's semantic indecision at w, (376) can include a restriction on the function's domain, such that it maps to true or false only those worlds w' such that $\Phi(\mathcal{M})(w')$ has a constant value for all $\mathcal{M} \in Sem_{x,w}$. Then if e.g. x is semantically indecisive at w regarding *athlete* as in §5.1.2, the value of the descriptive content of *Secretariat is an athlete* relative to x at w is undefined on worlds at which Secretariat competes in athletic events but is not human.

his beliefs about Secretariat may use a semantically underdetermined predicate, like *athlete*. And given these beliefs about the world, whether Alfonse can be truly reported as believing that Secretariat is an athlete depends on his semantic alternatives.

This can be shown by considering two cases. Suppose in the first case that Alfonse's semantic alternatives are completely decided regarding *athlete*, so that he invariantly treats the predicate as denoting P_1 (379a). Suppose in the second that his alternatives are again completely decided, but so that he invariantly treats the predicate as denoting the property of having both P_1 and P_2 (379b).

(379) a. $\forall \mathcal{M} \in Sem_{a,w}[[\![\text{athlete}]\!]^{\mathcal{M}} = P_1]$

b. $\forall \mathcal{M} \in Sem_{a,w}[[\![\text{athlete}]\!]^{\mathcal{M}} = \lambda w'_s.\lambda x_e.P_1(w')(x) \wedge P_2(w')(x)]$

In the first case, Alfonse's behavior treats the truth-conditional content of *athlete* as the property of competing in athletic events, while in the second, his behavior treats this truth-conditional content as the property of both competing in athletic events and being human. Given this and the denotation in (373), the following belief reports turn out true at w, with respect to the first situation and the second respectively (ignoring $\mathcal{M}$, which doesn't affect the denotation).

(380) a. **Situation I: $Sem_{a,w}$ as in (379a), $Dox_{a,w}$ as in (378b)**

$[\![\text{Alfonse thinks that Secretariat is an athlete}]\!]^{\mathcal{M},w}$

$= [\![\text{think}]\!]^{\mathcal{M},w}([\![\text{that Secretariat is an athlete}]\!])([\![\text{Alfonse}]\!]^{\mathcal{M},w})$

$= believe''(w)(\lambda \mathcal{M}'_m.\lambda w'_s.athlete'_{\mathcal{M}'}(w')(s))(a)$

$= \forall w' \in Dox_{a,w}, \mathcal{M}' \in Sem_{a,w}[athlete'_{\mathcal{M}'}(w')(s)]$

$= \forall w' \in Dox_{a,w}[P_1(w')(s)]$

$= 1$

b. **Situation II: $Sem_{a,w}$ as in (379b), $Dox_{a,w}$ as in (378b)**

$[\![\text{Alfonse thinks that Secretariat is not an athlete}]\!]^{\mathcal{M},w}$

$= [\![\text{think}]\!]^{\mathcal{M},w}(\lambda \mathcal{M}'_m.\lambda w'_s.[\![\text{not}]\!]^{\mathcal{M}',w'}([\![\text{that Secretariat is an athlete}]\!]^{\mathcal{M}',w'}))$

$$
\begin{aligned}
&([\![\text{Alfonse}]\!]^{\mathcal{M},w}) \\
&= \mathit{believe}''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.\neg \mathit{athlete}'_{\mathcal{M}'}(w')(s))(a) \\
&= \forall w' \in \mathit{Dox}_{a,w}, \mathcal{M}' \in \mathit{Sem}_{a,w}[\neg \mathit{athlete}'_{\mathcal{M}'}(w')(s)] \\
&= \forall w' \in \mathit{Dox}_{a,w}[\neg[P_1(w')(s) \wedge P_2(w')(s)]] \\
&= 1
\end{aligned}
$$

In the first case, Alfonse thinks that Secretariat is an athlete, and in the second, he thinks that Secretariat is not an athlete. This difference in beliefs turns on no difference in Alfonse's beliefs about the world – only his non-descriptive behavioral dispositions differ across the cases. Where Φ is the hyperintension of *that Secretariat is an athlete*, the descriptive contents of his belief in both situations are as follows.

(381) a. **Situation I**: $\Phi_{a,w} = \lambda w'_s.P_1(w')(s)$

b. **Situation II**: $\Phi_{a,w} = \lambda w'_s.\neg[P_1(w')(s) \wedge P_2(w')(s)]$

These propositions are reflected by Alfonse's doxastic alternatives on the second-to-last lines of (380a) and (380b). Since w falsifies neither of these, there is no descriptive error in either case, despite the hyperintensions believed being contradictory, in that there is no model and world relative to which both are true.

One can also characterize to what extent attitude reports that compare beliefs across agents are true in virtue of behavioral dispositions as well as beliefs about the world. Take a dispute between Alfonse and Bethany as to whether Secretariat is an athlete, where Secretariat's properties are again as in (378a), and both Alfonse and Bethany's beliefs about the world are as in (378b), so that these beliefs commit them to nothing false about Secretariat. Suppose then that Alfonse's semantic alternatives are as in (379a), while Bethany's are as in (379b). It follows, by parity of reasoning with (380), that Alfonse thinks that Secretariat is an athlete, and Bethany thinks that he isn't. And so by (375b), they are truly reported as disagreeing (with either of them as agent).

(382) $[\![\text{Bethany disagrees with Alfonse that Secretariat is an athlete}]\!]^{\mathcal{M},w}$

$= [\![\text{disagree}]\!]^{\mathcal{M},w}([\![\text{with Alfonse}]\!]^{\mathcal{M},w})([\![\text{that Secretariat is an athlete}]\!])([\![\text{Bethany}]\!]^{\mathcal{M},w})$

$= believe''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.athlete'_{\mathcal{M}'}(w')(s))(a) \wedge$

$believe''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.\neg athlete'_{\mathcal{M}'}(w')(s))(b)$

$= \forall w' \in Dox_{a,w}, \mathcal{M}' \in Sem_{a,w}[athlete'_{\mathcal{M}'}(w')(s)] \wedge$

$\forall w' \in Dox_{b,w}, \mathcal{M}' \in Sem_{b,w}[\neg athlete'_{\mathcal{M}'}(w')(s)]$

$= \forall w' \in Dox_{a,w}[P_1(w')(s)] \wedge \forall w' \in Dox_{b,w}[\neg[P_1(w')(s) \wedge P_2(w')(s)]]$

$= 1$

Alfonse and Bethany disagree, but only in virtue of their treatment of the meaning of *athlete*, not due to any property that they take Secretariat to have. Bethany's criterion for being an athlete is stricter than Alfonse's: she takes athletes only to be human, and so thinks that Secretariat isn't one.[19]

This disagreement is *descriptively compatible*, which occurs when the agents hold beliefs constituting the disagreement whose descriptive contents are truth-conditionally compatible (384). Otherwise, the disagreement is *descriptively incompatible* (385). Where 'ϕ' is a metavariable over proposition-denoting clauses of English:

(383) x disagrees with y that ϕ at w iff:

$believe''(w)([\![\phi]\!])(y) \wedge believe''(w)(\lambda\mathcal{M}_m.\lambda w'_s.\neg[\![\phi]\!]^{\mathcal{M},w'})(x)$

(384) x disagrees descriptively compatibly with y that ϕ at w iff:

a. x disagrees with y that ϕ at w

b. $\exists w'[[\![\phi]\!]_{y,w}(w') \wedge \neg[\![\phi]\!]_{x,w}(w')]$

(385) x disagrees descriptively incompatibly with y that ϕ at w iff:

a. x disagrees with y that ϕ at w

b. $\neg\exists w'[[\![\phi]\!]_{y,w}(w') \wedge \neg[\![\phi]\!]_{x,w}(w')]$

And so in the above scenario, Bethany disagrees descriptively compatibly with Alfonse

19. This is a formalization of what Sundell (2011: §3.4) calls 'character disagreement.'

that Secretariat is an athlete.[20] In a different scenario, the disagreement might be descriptively incompatible: this would happen where, for instance, both Alfonse and Bethany's semantic alternatives uniformly treat the denotation of *athlete* as P_1, but Bethany believes that Secretariat has P_1, while Alfonse believes that he doesn't, perhaps because he mistakenly thinks Secretariat is just a show horse.

Disagreements can also be characterized in terms of whether one of the agents they coordinate has made a descriptive error in holding the beliefs that constitute the disagreement. It is possible for a disagreement to be *descriptively faultless*, i.e. to take place in the absence of any descriptive error.

(386) x disagrees descriptively faultlessly with y that ϕ at w iff:

a. x disagrees with y that ϕ at w

b. $[\![\phi]\!]_{y,w}(w) \wedge \neg[\![\phi]\!]_{x,w}(w)$

The above disagreement between Alfonse and Bethany is descriptively faultless. While MacFarlane (2014: §6.7) has detailed the numerous dimensions along which a disagreement might be 'faultless' (cf. Ch. 1, fn. 38), and expressed skepticism about the usefulness of the notion of 'faultless disagreement,' this descriptive faultlessness is apparently what e.g. Kölbel (2003) and Lasersohn (2017) are after, and it is unproblematically characterizable once beliefs are hyperintensionalized.[21] Its occurrence doesn't depend on special expressions like 'predicates of personal taste' (cf. Lasersohn 2005) or 'discretionary predicates' (cf. Coppock 2018), but is ubiquitous in virtue of semantic underdetermination, in concert with differing behaviors of speakers.

20. Likewise, Alfonse disagrees descriptively compatibly with Bethany that Secretariat is not an athlete, and so on. More colloquial relations could then be defined out of those in the text, such as 'x and y disagree with each other over whether ϕ,' in a straightforward way.

21. The beliefs expressed by agents in a descriptively faultless disagreements are also 'mutually true' (cf. e.g. Stojanovic 2007: 692), in the restricted sense that the beliefs constituting the disagreement have descriptive contents true at the world of belief. They are *not* mutually true in that no model and world can verify both, meaning that speakers can't assent to both expressed beliefs simultaneously, insofar as their semantic behavior is consistent (cf. MacFarlane 2014: 134, on 'faultless$_t$ disagreement'). Doing so empties the agent's semantic alternatives.

The relations among these formal properties might be characterized or dissected further. For instance, from the above it follows that if a disagreement is descriptively incompatible, then it is not descriptively faultless, and that a disagreement is descriptively compatible only if the sets of semantic alternatives of the two agents are non-identical. Analogous notions are also available for characterizing other attitudes besides disagreement: e.g., an agreement might be descriptively incompatible, or non-fautless even where one of the two agents makes no descriptive error.[22]

5.2 Application to experiential predicates

The metasemantic results from §5.1 can be applied to the case of experiential predicates, to show how speakers' truth-conditional behavior with respect to these predicates is constrained by the experiences that they are disposed to have. On a par with the denotation for *athlete* provided in §5.1.2, the denotation for an experiential predicate like *tasty* is as follows, introducing model-sensitivity (but not experiencer-sensitivty) into the determination of its extension, in the same way as for non-experiential predicates. Thus sort of bare denotation provided in §1.3 receives an update using the metasemantic machinery, as follows (for the generic, positive form of the predicate).

(387) $[\![\text{tasty}]\!]^{\mathcal{M},w} = \lambda x_e.tasty'_{\mathcal{M}}(w)(x)$

Given the experiential compositional semantics provided in Chapters 2 and 3, what this amounts to, and how it relates to the status of bare experiential properties, is spelled out in this section.

§5.2.1 recasts autocentric evaluation of experiential predicates using the metasemantic machinery from §5.1, and shows how experiential dispositions can act as metasemantic con-

22. An example of an agreement with both these properties is one in which Bethany is as in the disagreement example given in the text, but Alfonse mistakenly believes that Secretariat has neither P_1 nor P_2. If Alfonse's semantic alternatives then uniformly treat *athlete* as denoting the property of having P_1 but not P_2, then Alfonse agrees with Bethany that Secretariat is not an athlete, and this despite the fact that Alfonse has made a descriptive error in his belief that Secretariat is not an athlete, while Bethany hasn't, and despite the fact that Alfonse's and Bethany's descriptive commitments constituting this agreement are incompatible.

straints on truth-conditional behavior. §5.2.2 then offers an explanation, both in terms of the experiential compositional semantics and in terms of the pragmatics of direct evidentiality, as to why these metasemantic constraints obtain. §5.2.3 then demonstrates how predicates that are evaluated autocentrically, along with semantically underdetermined predicates generally, give rise to a host of metasemantic effects that allow for the expression of opinion. Finally, it addresses the metaphysical status of bare experiential properties, showing that speakers do not commit to the existence of these in virtue of assessing autocentrically, despite experiential predicates not being relativized in their extensions to experiencers.

5.2.1 Experiences as metasemantic constraints

The way speakers treat the truth-conditional content of experiential predicates is typically constrained by the experiences that they are disposed to have. When evaluating such predicates autocentrically, speakers treat them as denoting the property that a stimulus has when it is generically disposed to stimulate a certain sort of experience in themselves.

This behavior is tracked by speakers' semantic alternatives. These alternatives can be defined in terms of speakers' experiential dispositions, and so where speakers evaluate autocentrically, their dispositions to experience act as metasemantic constraints on their semantic alternatives. Autocentric speakers' semantic alternatives align and misalign in their treatment of experiential predicates to the extent that the corresponding experiential dispositions align and misalign.

To characterize these behaviors, a speech community can be taken to be *perfectly autocentric*. This means that for any experiential predicate in the community's language, all speakers of that community treat its truth-conditional content as being the sort of self-directed dispositional property described above. This is an idealization, as shown by §4.2.1: speakers are only autocentric insofar as they take themselves not to be inaccurate perceivers, and they may have any number of reasons for not evaluating autocentrically.[23]

23. This is not an idealization in ignoring exocentric readings of experiential predicates. Recall from §2.3.2

It is difficult to characterize a speaker's treatment of an experiential predicate where autocentricity is abandoned, since speakers taking themselves to be perceptually inaccurate can do just about anything. But autocentric evaluation is a robust default, and it is formally tractable. It is defined in terms of semantic alternatives as follows, where ϵ is the positive, generic form of a gradable experiential predicate, and k_ϵ is the experiential kind associated with ϵ. The same holds for non-gradable experiential predicates as in §3.1.2, *mutatis mutandis*.

(388) x is perfectly autocentric with respect to ϵ at w iff:

$$\forall \mathcal{M} \in Sem_{x,w}[[\![\epsilon]\!]^{\mathcal{M}} = \lambda w'_s.\lambda y_e.0_{k_\epsilon} <_{k_\epsilon} \delta'''_{k_\epsilon}(w')(x,w)(y)]$$

The denotation in (388) makes use of a new symbol: '$\delta'''_{k_\epsilon}(w')(x,w)(y)$' is to be read, 'the degree of the experiential kind associated with ϵ that y is disposed to produce at w', according to the experiential standard set by x at w.' The positive form then functions normally, as in §2.3.1: the predicate is true of an individual just in case this degree that the individual is disposed to produce is greater than the zero-degree of the experiential kind. In all, what (388) says is then that x's semantic alternatives at w are completely decided, and treat the truth-conditional denotation of ϵ uniformly as the property of being disposed to produce a non-zero degree of the relevant experiential kind in x.[24]

This dispositional notion can then be expanded as follows, where $\rho_{k,x,w}$ is an accessibility

---

that exocentric predicates are structurally distinct from their 'properly' experiential counterparts, and deal only with reports of actualized experiential states. The metasemantic question of how to evaluate these doesn't arise in the same way. §5.2.2 shows how autocentricity is traceable to GEN in the compositional semantics, which exocentric readings lack.

24. There are two simplifications in (388). First, the property characterized is a self-directed *de re* one – it should strictly be *de se*. In a framework in the spirit of Lewis (1979) and Chierchia (1989), where *de se* attitudes are those directed towards self-ascribed properties, 'x' following the equals sign would be an individual variable abstracted over by a property-forming operator higher in the clause (cf. Pearson 2013: §3.5), or tracked by an intensional parameter acting as the center (cf. Stephenson 2007a, 2010). This can be idealized away by supposing that every speaker y self-ascribes the property $\lambda w_s.\lambda z_e.z = y$, and no other identity property – where each speaker is perfectly decided as to who they are, the distinction between *de se* and self-directed *de re* attitudes collapses. A full incorporation of *de se* attitudes would allow for autocentricity even in the face of uncertainty of self-identity. Second, this denotation ignores genericity not having to do with the experiencer. Where this is not ignored, speakers will typically have their semantic alternatives for experiential predicates not perfectly decided, in virtue of indecsion as to which of their own experiences are ideally typical.

relation on worlds governed by the experiential kind k, like that appealed to in §2.3.1, but which is anchored to the experiential dispositions of x at w'.

(389) $\delta'''_k(w)(x, w')(y) := \iota d[\forall w'' : w\rho_{k,x,w'}w''[d = \delta''_k(w'')(x)(y)]]$

And so what the δ'''-term denotes is that degree of the experiential kind k, that the stimulus y produces in x, in all worlds ideally normal for the production of k in x as it exists in w'. As in §2.3.1, the relation $\rho_{k,x,w'}$ can be thought of as enforcing a modal base an ordering source, as follows (cf. Ch. 2, fn. 32).

(390) $w\rho_{k,x,w'}w''$ iff:

a. w'' preserves (i) all intrinsic properties holding at w of potential stimuli, relevant for their production of k; and (ii) all intrinsic properties holding of x at w', relevant for the experience of k;

b. w'''s conditions, given these properties, are ideally normal for the production of k by potential stimuli in x.[25]

Further autocentric properties of a speech community can then be defined derivatively: for instance, speakers might be considered perfectly autocentric *simpliciter* where they are perfectly autocentric with respect to every experiential predicate, and a speech community as a whole might be called perfectly autocentric where each of its members is.

The experiential predication *Licorice is tasty* then expresses that licorice is tasty *simpliciter*, and in a perfectly autocentric community the hyperintension that it denotes will be believed, asserted, denied, etc. according to the descriptive commitments of each agent, as

25. Presumably, these ideally normal conditions entail that x has had the actual experience from the stimulus in question, allowing the δ''-term in (389) to be defined. The reason that w and w' are distinct is that in certain modal constructions, it is apparently possible for the world determining the experiencer's dispositions to be stuck to the autocentric evaluator's actual circumstances (hence why the relativists invoke the world of assessment, not the world of evaluation: cf. Ch. 1, fn. 33). An example of this is an autocentric speaker asking herself, *If I were a horse, would hay be tasty?* One answer is 'yes,' but it is apparently possible also to answer 'no,' where the questioner doesn't *actually* like hay: the idea is that hay would remain disgusting, but of course as a horse the questioner would like it: cf. MacFarlane (2014: §7.2.11). The 'yes' response is handled by the exocentric reading, where the question is effectively 'would hay be tasty *to me*?' This shifts the experiencer with the modal as desired.

constrained by their experiences. That is to say, the sentence itself relativizes to no experiencer whatsoever, but the truth-conditional content of the sentence as treated by the agents' behavior takes its truth to be tracked by their own dispositions to experience. And so where again *tasty* is taken to be the positive, generic form of the predicate:

(391) $[\![\text{Licorice is tasty}]\!]^{\mathcal{M},w} = tasty'_{\mathcal{M}}(w)(l)$

Suppose then that Alfonse is perfectly autocentric with respect to *tasty* at w (392a). If he believes about the world that licorice is disposed to produce gustatory pleasure *in him* (392b),[26] then the relevant belief report comes out true (392c), where the denotation is simplified to pertain to the experiential kind GUS-PL (gustatory pleasure).

(392) a. $\forall \mathcal{M} \in Sem_{a,w}[[\![\text{tasty}]\!]^{\mathcal{M}} = \lambda w'_s.\lambda x_e.0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)]$

b. $\forall w' \in Dox_{a,w}[0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)]$

c. $[\![\text{Alfonse thinks that licorice is tasty}]\!]^{\mathcal{M},w}$
$= [\![\text{think}]\!]^{\mathcal{M},w}([\![\text{that licorice is tasty}]\!])([\![\text{Alfonse}]\!]^{\mathcal{M},w})$
$= believe''(w)(\lambda \mathcal{M}_m.\lambda w'_s.tasty'_{\mathcal{M}}(w')(l))(a)$
$= \forall w' \in Dox_{a,w}, \mathcal{M} \in Sem_{a,w}[tasty'_{\mathcal{M}}(w)(l)]$
$= \forall w' \in Dox_{a,w}[0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)]$
$= 1$

And so in general the descriptive commitments that agents in a perfectly autocentric community make in virtue of believing an experiential hyperintension are descriptive commitments about their own dispositions to experience. Agents can then agree or disagree regarding these hyperintensions in the ordinary way. Thus if Bethany is also perfectly autocentric, and takes licorice to be disposed to produce gustatory pleasure *in her*, then in thinking that licorice is tasty, she is descriptively committed to just this, and she and

26. In a *think*-report, the acquaintance inference will imply that Alfonse believes this because he has tasted licorice and liked it: cf. §4.1.4.

Alfonse agree, and literally believe *the same thing*, despite the fact that their descriptive commitments in virtue of believing this differ.

Some consequences follow from this, which are expressible using the terminology introduced in §5.1.3. First, agents' beliefs in the meaning of experiential predications, where those agents are perfectly autocentric, are in descriptive error just in case the agents are mistaken about their own experiential dispositions with respect to the stimulus. Where these beliefs are held due to direct experience with the stimulus, this entails that descriptive error only occurs where agents have forgotten what their experience was, are unaware that the stimulus' properties have relevantly changed, have inappropriately generalized from atypical encounters with it, etc.[27]

Second, a disagreement between perfectly autocentric speakers over an experiential predication is descriptively faultless just in case those speakers' experiential dispositions with respect to the stimulus relevantly diverge. This means that such disagreements are *not* descriptively faultless where these experiential dispositions converge. If $\delta'''_{\text{GUS-PL}}(w)(a,w)(l) = \delta'''_{\text{GUS-PL}}(w)(b,w)(l)$, and if at w Alfonse thinks licorice is tasty while Bethany doesn't, then one of them is making a descriptive error. This is as it should be, since among an 'experiential cohort,' one can be wrong about such matters (in this specific sense), and to the extent that speakers' experiential dispositions align, they convey *information about the world* to

27. 'Epistemic privilege' (cf. Lasersohn 2005: 655; Pearson 2013: 146) about one's own experiential dispositions doesn't amount to infallibility, since speakers can very well be mistaken about their own dispositions, even after having the relevant experiences (cf. Ch. 1, fn. 38). An impression of (attenuated) immunity from error often comes from focusing on cases where the agent has already had the relevant experience, but this need not be the case. Using a predicate like *believe*, which implies a lack of direct evidence (cf. §4.1.4), makes clearer that descriptive error in autocentric experiential beliefs is not only possible, but not unusual, e.g. with *Alfonse believes that licorice is tasty*. There is nothing strange about Alfonse being descriptively mistaken in such a belief: he tries licorice, and finds to his disappointment that he dislikes it. A potentially unappealing feature of Coppock's (2018) outlook semantics is that it hard-codes certain propositions as 'discretionary' and thus immune from any possibilty of descriptive fault (cf. *ibid.* 136: 'So no agent could *ever* be at fault for expressing either one of these sentences.' [*The chili is tasty* or *The chili is not tasty*]). Kneer (m.s.) adduces evidence that speakers do actually take each other to be at descriptive fault when such self-directed mistakes are made. Of course, an agent does not believe something false *simpliciter* just in virtue of being in descriptive error: if Bethany likes licorice, she may take Alfonse to have been right in his prior belief anyway, without committing a descriptive error herself. But Coppock's use of intrinsic parameters to handle the language of opinion prevents capturing these nuances, unless it is supplemented by independent norms on belief, as in Kölbel (2002: §6.4).

each other in virtue of asserting that experiential predicates hold of things, viz. information about what experiences a stimulus is *in fact* disposed to produce.[28]

Experiential predicates do not intrinsically encode matters of opinion as opposed to matters of fact. Like all predicates, they convey information about the world as tracked by speakers' semantic behaviors. But they are prone to a misalignment of semantic alternatives due to their experiential semantics (on which more in §5.2.2). The descriptive commitments that speakers make in virtue of the experiential constraints on semantic alternatives in turn make such expressions suited for expressing those features of speakers in virtue of which a misalignment happens, viz. their varying experiential dispositions. Hence the fact that speakers express opinions pertaining to their own experiences using these predicates is a systematic epiphenomenon.[29]

Descriptively faultless disagreement is also an epiphenomenon of these metasemantic constraints – it happens ubiquitously, but is especially pronounced where experiential predicates are concerned, insofar as speakers' dispositions to experience diverge. It therefore plays no special theoretical role, but because the literature has been occupied with it, it is worth working through an example involving experiential predicates.

Suppose at w that licorice is disposed to produce a non-zero degree of gustatory pleasure in Alfonse (393a), but a zero-degree in Bethany (393b). Suppose further that Alfonse and Bethany have both tasted licorice, had the corresponding reaction, and now rightly believe

28. This is again an idealization in ignoring the effect of genericity upon descriptive faultlessness. Even if Alfonse and Bethany have the exact same experiential reactions to licorice and are perfectly autocentric, they might descriptively faultlessly disagree over whether it is tasty, if their semantic alternatives are misaligned due to their treating different sorts of experiences as ideally normal. They may also take the same experiences of gustatory pleasure to be ranked differently according to the preorder on qualities from §2.1.1, though this is irrelevant for evaluation of the positive form, *modulo* pragmatic imprecision.

29. In the special case where a perfectly autocentric group of speakers are entirely uniform in their experiential dispositions, their semantic alternatives align entirely with respect to every experiential predicate. This means that their disagreements regarding such predicates are always descriptively incompatible, and relative to this group, the experiential predicate simply denotes some single property *simpliciter*, and the opinion-expressing function of the predicate to a large extent disappears. But it still has an experiential semantics, and interacts with the compositional machinery in the same way, and speakers grasp that *if* someone's experiential dispositions were to change, then so would evaluation of the predicate. This is unlike with color predicates like *red*, which despite being tracked by speakers' experiences, do not have an experiential semantics, and simply denote some non-experiential property of individuals.

these things (394). Finally, let both of them be perfectly autocentric with respect to (the positive, generic form of) *tasty* (395).

(393) a. $0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w)(a,w)(l)$

b. $0_{\text{GUS-PL}} =_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w)(b,w)(l)$

(394) a. $\forall w' \in Dox_{a,w}[0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)]$

b. $\forall w' \in Dox_{b,w}[0_{\text{GUS-PL}} =_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(b,w)(l)]$

(395) a. $\forall \mathcal{M} \in Sem_{a,w}[[\![\text{tasty}]\!]^{\mathcal{M}} = \lambda w'_s.\lambda x_e.0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)]$

b. $\forall \mathcal{M} \in Sem_{b,w}[[\![\text{tasty}]\!]^{\mathcal{M}} = \lambda w'_s.\lambda x_e.0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(b,w)(l)]$

The *disagree*-report then comes out true (396). But the descriptive content of Bethany's and Alfonse's beliefs are as in (397a)-(397b), and since these are both verified by the world of belief w, the disagreement is descriptively faultless (397c).

(396) $[\![\text{Bethany disagrees with Alfonse that licorice is tasty}]\!]^{\mathcal{M},w}$

$= [\![\text{disagree}]\!]^{\mathcal{M},w}([\![\text{with Alfonse}]\!]^{\mathcal{M},w})([\![\text{that licorice is tasty}]\!])([\![\text{Bethany}]\!]^{\mathcal{M},w})$

$= believe''(w)(\lambda \mathcal{M}'_m.\lambda w'_s.tasty'_{\mathcal{M}'}(w')(l))(a) \wedge$

$believe''(w)(\lambda \mathcal{M}'_m.\lambda w'_s.\neg tasty'_{\mathcal{M}'}(w')(l))(b)$

$= \forall w' \in Dox_{a,w}, \mathcal{M} \in Sem_{a,w}[tasty'_{\mathcal{M}}(w')(l)] \wedge$

$\forall w' \in Dox_{b,w}, \mathcal{M} \in Sem_{b,w}[\neg tasty'_{\mathcal{M}}(w')(l)]$

$= \forall w' \in Dox_{a,w}[0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l) \wedge$

$\forall w' \in Dox_{b,w}[0_{\text{GUS-PL}} =_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(b,w)(l)$

$= 1$

(397) a. $\Phi_{a,w} = \lambda w'_s.0_{\text{GUS-PL}} <_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(a,w)(l)$

b. $\Phi_{b,w} = \lambda w'_s.0_{\text{GUS-PL}} =_{\text{GUS-PL}} \delta'''_{\text{GUS-PL}}(w')(b,w)(l)$

c. $\Phi_{a,w}(w) \wedge \Phi_{b,w}(w)$

5.2.2 The origin of autocentricity

A treatment of the metasemantics of predicates generally allows a characterization of speakers' behavior when they evaluate experiential predicates autocentrically. But still the question remains: why are experiential predicates evaluated autocentrically? This ought to be explained in terms of their special experiential semantics, as laid out in Chapter 2 and 3.

There are two levels on which the question can be answered. First, the source of autocentricity can be traced in the compositional semantics, to see how the relevant model-sensitivity is introduced. Second, a reason can be given why this truth-conditional component gives rise to certain self-directed metasemantic behaviors. Both sorts of answers will be given in turn.

To trace autocentricity in the compositional semantics, experiential predicates must be decomposed once again, in the manner of Chapters 2 and 3. §5.2.1 showed the end result of the autocentric evaluation of a positive, generic form an an experiential predicate, but with the new metasemantic machinery involving models in place, this result can be built back up from its constituent pieces. This is done here using the exemplar *tasty*, and analogous compositions hold for other sorts of experiential predicates.

To begin with the final denotation for *tasty* given in its positive generic form, from §3.2.2 (cf. (276)): this is composed using the relational noun *taste*, the adjectival suffix *-y*, the silent degree morpheme POS, and finally the generic component GEN, in virtue of which individual-level predicates have dispositional readings. This results in the following denotation, in the prior non-metasemantic fragment of the grammar.

(398) $[\![\text{GEN } [\text{POS tasty}]]\!]^w = \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$

That is, this form of *tasty* is true of an individual just in case its taste stimulates a non-zero degree of pleasure, *simpliciter*. Recall (cf. §2.3.3, (220b)) that δ'_k is a function of type $\langle s, \langle e, d\rangle\rangle$, that maps a world w and an individual x to the degree of the experiential kind k that x is disposed to stimulate *simpliciter* at w. The problem opening this chapter is

that one can't adequately track speakers' use of the experiential predicate by characterizing whether an individual produces a non-zero degree of such a kind *simpliciter*. This is in turn traceable to the fact that it is inadequate to characterize $\delta'_k(w)(x)$ *simpliciter*, for arbitrary k, w, x.

And so the semantics at this level of description ought not to specify to what degree the taste of an individual stimulates pleasure *simpliciter*. This is the general issue for experiential predicates: a mapping needs to be determined from individuals (stimuli) to degrees of the appropriate experiential kind (cf. Glanzberg 2007; Lasersohn 2008; Kennedy 2013; Fleisher 2013). It is here that speakers' behaviors, and so the metasemantics, become crucial – in evaluating autocentrically, speakers behave in a way to determine such a function, one based on their own dispositions to experience. And so in recasting the generic, positive form of *tasy* into the grammar with its metasemantic component, it is at this mapping that model-sensitivity is introduced.

(399) $[\![\text{GEN [POS tasty]}]\!]^{\mathcal{M},w}$
$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M},\text{PLEASURE}}(w)(\iota y[taste''(w)(x)(y)])$

The denotation of *tasty* in this form, relative to a behavioral disposition $\mathcal{M}$, is the property that an individual has when its taste is disposed to stimulate a non-zero degree of pleasure, as determined by the mapping of individuals to degrees of pleasure determined by $\mathcal{M}$. So δ'_k has been recast as a function of type $\langle m, \langle w, \langle e, d\rangle\rangle\rangle$: it takes a pattern of behavior, a world, and a stimulus into a degree (where $\mathcal{M}$ is denoted with a subscript, per usual). This denotation is just an alternate way of writing '$tasty'_{\mathcal{M}}(w)$.'

This function is what a speaker's autocentric behavior determines: the degree of an experiential kind that a stimulus is disposed to produce *simpliciter* is treated as equal to the degree of the experiential kind that the stimulus is disposed to produce in the speaker (at the world that the speaker occupies). This means that the positive form tracks whether a stimulus is disposed to produce a non-zero degree in that speaker, but the same pattern arises for any degree construction, since the speaker's behavior construes truth-conditionally

how tasty a stimulus is *simpliciter*. And so the pattern is as follows.

(400) x is perfectly autocentric with respect to *tasty* at w iff:
$\forall y, w', \mathcal{M} \in Sem_{x,w}[\delta'_{\mathcal{M},\text{PLEASURE}}(w')(\iota z[taste''(w')(y)(z)) = \delta'''_{\text{PLEASURE}}(w')(w,x)(\iota z[taste''(w')(y)(z))]$

In other words, the autocentric speaker's semantic alternatives are fixed, to treat the degree of pleasure that a stimulus' taste is disposed to produce *simpliciter* as equal to the degree of pleasure that its taste is disposed to produce in themselves. This can be generalized and expressed more succinctly as follows, where k is an experiential kind and y a stimulus.

(401) x is perfectly autocentric with respect to k and y at w iff:
$\forall w', \mathcal{M} \in Sem_{x,w}[\delta'_{\mathcal{M},k}(w')(y) = \delta'''_{k}(w')(w,x)(y)]$

This can be characterized at different levels as well: speakers can be perfectly autocentric with respect to an experiential kind k just in case they are perfectly autocentric with respect to k and every stimulus y; they can be perfectly autocentric *simpliciter* where they are perfectly autocentric with respect to every experiential kind, and so on.

With this equivalence in mind, the predicate can be decomposed to trace the source of the model-sensitivity. Beginning by composing *taste* with *-y*:

(402) a. $[\![\text{taste}]\!]^{\mathcal{M},w} = \lambda x_e.\lambda y_e.taste''(w)(x)(y)$

b. $[\![\text{-y}]\!]^{\mathcal{M},w} = \lambda R_{e,et}.\lambda x_e.\delta_{\text{PLEASURE}}(\sigma y[stim''(w)(\iota z[R(x)(z)])(y)])$

c. $[\![\text{tasty}]\!]^{\mathcal{M},w} = [\![\text{-y}]\!]^{\mathcal{M},w}([\![\text{taste}]\!]^{\mathcal{M},w})$
$= \lambda x_e.\delta_{\text{PLEASURE}}(\sigma y[stim''(w)(\iota z[taste''(w)(x)(z)])(y)])$

The result is, as in §3.2.2, a measure function that maps individuals x to the degree of pleasure instantiated by the maximal experience stimulated by the taste of x. There is no model-sensitivity in the denotation: this is as it should be, since thus far the denotation just pertains to what sort of experience the taste of an individual has produced *simpliciter*, and

at a world, this is a determinate matter, independent of construal by speaker behavior.[30]

Next, degree morphology needs to compose with the adjective to form a property-denoting expression. The positive form (as it pertains to lower-closed predicates: cf. §2.2.2) arises as usual through composition with POS.

(403) a. $[\![\text{POS}]\!]^{\mathcal{M},w} = \lambda G_{s,et}.\lambda x_e.0_G <_G G(w)(x)$

b. $[\![\text{POS tasty}]\!]^{\mathcal{M},w} = [\![\text{POS}]\!]^{\mathcal{M},w}(\lambda w_s.[\![\text{tasty}]\!]^{\mathcal{M},w})$

$= \lambda x_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta_{\text{PLEASURE}}(\sigma y[stim''(w)(\iota z[taste''(w)(x)(z)])(y)])$

The positive form denotes the property that an individual x has just in case the maximal experience that the taste of x has stimulated instantiates a non-zero degree of pleasure. This is the denotation for the positive, exocentric form of the predicate. Still, there is no model-sensitivity: all that is reported by a use of this predicate is that some experiential state is instantiated. There may be questions concerning how to narrow the domain restriction on the experience in question (cf. §2.3.2), to localize it to a contextually relevant experiencer, but these have to do with setting parameters at the context of utterance, and not with the metasemantic behavior of autocentricity.

It is only when GEN composes with the predicate that model-sensitivity arises: this is because all generic quantification involves felicity conditions, which are required to determine which situations are ideally normal and so count towards the application of a predicate. But speakers generally vary on what these felicity conditions are, and this is not a feature of experiential predicates specifically, but of generic quantification as a whole (cf. §2.3.1). Only once the felicity conditions are set can an accessibility relation for the generic be fixed, and so this accessibility relation is construed by different speakers in different ways, according to what they take via their semantic behavior to be ideally normal. And so it is here on this accessibility relation that model-sensitivity is introduced.

30. Again, this is an idealization that ignores the fact that speakers might differ as to the preorder on qualities, and so take the very same experience of gustatory pleasure to instantiate different degrees. But this issue does not pertain to the metasemantics relevant for autocentricity.

(404) $[\![\text{GEN}]\!]^{\mathcal{M},w} = \lambda P_{s,et}.\lambda x_e.\forall w' : w\rho_{\mathcal{M},P}w'[P(w')(x)]$

$\rho_{P,\mathcal{M}}$ is that accessibility relation on worlds relevant for the ideally normal application of P, as construed by the pattern of behavior $\mathcal{M}$. Differences among speaker behavior regarding this accessibility relation will cause the sorts of behavioral disagreements outlined in §5.1.3: for instance, even where two speakers are fully aware of exactly how fast Alfonse can move in which situations (and where the cutoff point for a comparison class in speed for the positive form of *fast* is agreed on, relative to a single event), they might disagree as to whether Alfonse is (generically) fast, since one speaker takes only speed at running the mile to be indicative of ideal normality, while another takes only speed at running the 100-yard dash to be indicative of ideal normality.

Applying this to the case of the experiential predicate yields the following result. Where δ'-terms are redefined as in (406) to incorporate model-sensitivity (and the degree of a kind that a stimulus produces is taken as usual to converge at all ideally normal worlds), the denotation in (405) is identical to the one in (399).

(405) $[\![\text{GEN [POS tasty]}]\!]^{\mathcal{M},w} = [\![\text{GEN}]\!]^{\mathcal{M},w}(\lambda w'_s.[\![\text{POS tasty}]\!]^{\mathcal{M},w})$
$= \lambda x_e.\forall w' : w\rho_{\mathcal{M},\text{PLEASURE}}w'$
$[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta_{\text{PLEASURE}}(\sigma y[stim''(w')(\iota z[taste''(w')(x)(z)])(y)])]$

(406) $\delta'_{\mathcal{M},k}(w)(x) := \iota d[\forall w' : w\rho_{\mathcal{M},k}w'[d = \delta_k(\sigma y[stim''(w')(x)(y)])]]$

And so the positive, generic form of *tasty* denotes that property true of an individual x just in case at all accessible worlds ideally normal for the production of pleasure, as this accessibility relation is construed by a pattern of behavior $\mathcal{M}$, the degree of pleasure instantiated by the maximal experience that the taste of x stimulates is non-zero. Recall from §2.3.1 that ρ_P corresponded to ρ_k, where P is an experiential property pertaining to the experiential kind k: the same holds here, with the model-sensitivity introduced, and so it is $\rho_{\mathcal{M},\text{PLEASURE}}$ that is relevant for the generic form of *tasty*.

With the denotation in (405), the usual metasemantic variation due to genericity arises:

even two speakers assessing autocentrically with the exact same experiential dispositions may differ in the matter of which of these experiencers is indicative of ideal normality. Cf. §2.3.1 for examples with *frightening*: for *tasty*, relevant concerns will be e.g. which instance of a kind counts as having an ideally normal taste, which condition of the experiencer is ideally normal (e.g., whether the palette is clean or not), what the ideally normal season is for tasting a certain sort of food, etc.

But autocentric speakers do something else in construing the generic accessibility relation: they take it to be a relation pertaining to ideally normal conditions on their *own* experiences. In other words, their pattern of behavior construes the accessibility relation as being one like in §5.2.1, one anchored to the experiential status of an individual at a world, where they themselves are the individual in question. And so autocentricity with respect to an experiential kind can be recast as follows.

(407) x is perfectly autocentric with respect to k at w iff:

$\forall \mathcal{M} \in Sem_{x,w}[\rho_{k,\mathcal{M}} = \rho_{k,x,w}]$

Where $\rho_{k,x,w}$ is characterized as in (390): more finely-grained or loosely-grained notions of perfect autocentricity can then be defined as usual.[31] With this, the first question as to the origin of autocentricity is answered: it arises as a result of the generic quantification inherent to dispositional uses of individual-level predicates. Given the portrayal in §2.3.3, this is unsurprising: it was precisely with the addition of GEN that these predicates became 'genuinely' experiential or evaluative (cf. (238a)). This autocentric construal of the generic accessibility relation carries over to degree constructions besides the positive form in the expected way, since the restricted accessibility relation determines how a speaker construes e.g. *how tasty* an individual is *simpliciter*.

31. Though some of these might be a little harder to express with notation. For example, if Alfonse at w is perfectly autocentric with respect to fear (and totally decided with respect to other generic concerns), except with horror movies, which he takes himself to have no taste for, then for all $\mathcal{M} \in Sem_{a,w}$, $\rho_{\mathcal{M},\text{FEAR}}$ is not exactly $\rho_{\text{FEAR},w,a}$, but a closely related accessibility relation that matches $\rho_{\text{FEAR},a,w}$ with respect to the properties of every non-horror-movie individual. The way movies are taken to be in virtue of being ideally normal with respect to stimulating fear will be determined in some other way, depending on what Alfonse decides to do.

Denotations for the other sorts of experiential predicates explored in this work then behave in the expected way, for example:

(408) $[\![\text{GEN [POS frightening]}]\!]^{\mathcal{M},w}$

$= \lambda x_e.\forall w' : w\rho_{\mathcal{M},\text{FEAR}}w'[0_{\text{FEAR}} <_{\text{FEAR}} \delta_{\text{FEAR}}(\sigma y[stim''(w')(x)(y)])]$

$= \lambda x_e.0_{\text{FEAR}} <_{\text{FEAR}} \delta'_{\mathcal{M},\text{FEAR}}(w)(x)$

$= \lambda x_e.frightening'_{\mathcal{M}}(w)(x)$

(409) $[\![\text{GEN [looks wooden]}]\!]^{\mathcal{M},w}$

$= \lambda x_e.\forall w' : w\rho_{\mathcal{M},\text{VIS}}w'$

$[evid''(w')(\lambda w'_s.wooden'(w')(x))(\sigma y[stim''(w')(\iota z[look''(w')(x)(z)](y)])]$

(410) $[\![\text{GEN [POS [sounds funny]]}]\!]^{\mathcal{M},w}$

$= \lambda x_e.\forall w' : w\rho_{\mathcal{M},\text{HUMOR}}w'$

$[0_{\text{HUMOR}} <_{\text{HUMOR}} \delta_{\text{HUMOR}}(\sigma y[stim''(w')(\iota z[sound''(w')(x)(z)])(y)])]$

$= \lambda x_e.0_{\text{HUMOR}} <_{\text{HUMOR}} \delta'_{\mathcal{M},\text{HUMOR}}(w)(\iota y[sound''(w)(x)(y)])$

This leaves the second question: why do speakers construe genericity with experiential predicates in this self-directed way? To answer this, the results concerning the nature of direct evidentiality and its relation to experiential semantics from Chapter 4 need to be revisited. The notions of perceptual alternatives (411a) and verification of a proposition by direct evidence (411b) can be imported unchanged from §4.1.3 (cf. (308) and (309)).

(411) a. $Per_{x,w} := \{w' : w'$ is compatible with x's direct perceptions at $w\}$

b. x's direct evidence at w verifies ϕ iff:

$\forall w' \in Per_{x,w}[\phi(w')]$

The contribution of a direct evidential DIR must be modified slightly to incorporate model-sensitivity. The results from §4.1.3 are preserved if the contribution of the direct evidential is again treated as intrinsically autocentric: it commits speakers to having their own perceptual alternatives verify the at-issue proposition, relative to the relevant model

that is consistent with their own actual experiences. Let $\mathcal{M}_{x,w}$ be the model like $\mathcal{M}$ but anchored to x's experiences at w, defined as follows.

(412) $\mathcal{M}_{x,w}$:= that model just like $\mathcal{M}$, except that:
$\forall y, k, w'[\delta'_{\mathcal{M}_{x,w},k}(w')(y) = \delta''_k(w)(x)(y)]$

That is, $\mathcal{M}_{x,w}$ differs from $\mathcal{M}$ only in treating all dispositions to stimulate experiences *simpliciter* as actualized experiences in x at w. This notion does the exact same formal job as experiential alternatives from earlier in this work, which it now supersedes. For simplicity of exposition, the definition here is confined to dealing with experiential kinds, though it could be extended to dealing with experiences generally as experiential alternatives were in Chapter 3. The direct evidential is then recast as follows.

(413) $[\![\text{DIR}]\!]^{\mathcal{M},c,w} = \lambda\Phi_{m,st}.\langle\Phi(\mathcal{M})(w), \forall w' \in Per_{s_c,w}[\Phi(\mathcal{M}_{s_c,w})(w')]\rangle$

It commits the speaker s_c at the context of utterance to the truth of the not-at-issue proposition that s_c's perceptual alternatives verify the at-issue proposition relative to the speaker's own experiences. A commitment to direct evidence in uttering *Licorice is tasty* then preserves the desired result as follows, where the not-at-issue proposition in (414b) reduces to (415) by the equivalence in (412), and the collapse of direct evidence into experience (cf. §4.1.3).

(414) a. $[\![\text{licorice [is [GEN [POS tasty]]]}]\!]^{\mathcal{M},c,w}$
$= 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M},\text{PLEASURE}}(w)(\iota y[taste''(w)(l)(y)])$

b. $[\![\text{DIR [licorice [is [GEN [POS tasty]]]]}]\!]^{\mathcal{M},c,w}$
$= \langle 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M},\text{PLEASURE}}(w)(\iota y[taste''(w)(l)(y)]),$
$\forall w' \in Per_{s_c,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M}_{s_c,w},\text{PLEASURE}}(w')(\iota y[taste''(w'')(l)(y)])]\rangle$

(415) $0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(s_c)(\iota y[taste''(w)(l)(y)])$

And so the acquaintance inference is derived as usual. With this updated characterization of direct evidentiality in mind, a new characterization can be given for what it means to

perceive accurately, and it turns out that to dispositionally treat oneself as an accurate perceiver with respect to the denotation of experiential predicates just is to assess those predicates autocentrically. Recall that in §4.2.1, an accurate perceiver was defined as one whose direct perceptions tracked the truth of propositions. This notion also must be updated within the new metasemantic machinery, and it too must be model-sensitive.

(416) x is an accurate perceiver according to $\mathcal{M}$ with respect to Φ at w iff:
$\Phi(\mathcal{M})(w) \wedge \forall w' \in Per_{x,w}[\Phi(\mathcal{M}_{x,w})(w')]$

And so an individual x accurately perceives that Φ (according to $\mathcal{M}$) where Φ is true and x perceives that it is.

Recall from §4.2.1 that there are robust cross-linguistic generalizations associating direct evidence with a commitment to truth: in general, speakers must commit to propositions asserted on the basis of direct evidence, and direct evidence is treated as fundamentally indicative of truth in a way that other evidential sources are not. From this it was concluded that speakers are bound by pragmatic norms on evidentiality to take their own experiences to be accurate. With the present metasemantic machinery implemented, this notion has become more nuanced: to be an accurate perceiver is a model-relative notion, and this means that one is not an accurate perceiver *simpliciter*, but rather is construed as one by a pattern of semantically relevant behavior.

And so whether one speaker treats another as an accurate perceiver can be characterized by quantifying over semantic alternatives. For x to treat y as an accurate perceiver of Φ is for all of x's semantic alternatives to be such that Φ is actually true, and y perceives that it is.

(417) x treats y as an accurate perceiver with respect to Φ at w iff:
$\forall \mathcal{M} \in Sem_{x,w}[\Phi(\mathcal{M})(w) \wedge \forall w' \in Per_{y,w}[\Phi(\mathcal{M}_{y,w})(w')]]$

Given the fact that speakers are pragmatically compelled to take their own direct evidence as accurate, the interesting case is where $x = y$, i.e. where an individual treats herself as

an accurate perceiver with respect to a hyperintension. This happens just where according to that speaker's semantic alternatives, the hyperintension is true, and further, that same speaker perceives that it is.

For x not to treat x as an accurate perceiver with respect to whether licorice is tasty, for example, might be for x to have semantic alternatives $\mathcal{M}$ determining that both:

(418) a. $0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta'_{\mathcal{M},\text{PLEASURE}}(w)(\iota y[taste''(w)(l)(y)])$

b. $0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(x)(\iota y[taste''(w)(l)(y)])$

To be x in this situation would be to have the taste of licorice stimulate pleasure in oneself, but *not* to take this fact to verify the hyperintension of *Licorice is tasty*. In this case, one's direct evidence is out of step with one's semantic behaviors, and so one is disobeying the pragmatic norm, and refusing to treat one's own perceptions as a reliable guide to truth.

By contrast, to treat oneself as an accurate perceiver is to take whatever degree of an experiential kind an individual stimulates in oneself to be, as treated by one's own semantic alternatives, the degree that this individual stimulates *simpliciter*. If this notion is dispositionalized, then it is equivalent to evaluating autocentrically. Recall the definition of perfectly autocentric assessment from (401) in §5.2.2.

(419) x is perfectly autocentric with respect to k and y at w iff:
$\forall w', \mathcal{M} \in Sem_{x,w}[\delta'_{\mathcal{M},k}(w')(y) = \delta'''_k(w')(x,w)(y)]$

What this says is that a speaker treats the degree of k that y is disposed to stimulate *simpliciter* as equivalent to the degree of k that y is disposed to produce in x as it is at w. To deviate from autocentric evaluation would therefore be to entertain the possibility that one is disposed to have an experience that does not track the amount of an experiential kind the stimulus is disposed to produce *simpliciter*. But per the above, this is exactly to entertain the possibility that one is not disposed to be an accurate perceiver – that one might receive direct evidence for an experiential property that one does not thereby commit to the stimulus having. To take oneself to be relevantly disposed to perceive accurately with

respect to an arbitrary proposition pertaining to the stimulus y and kind k is therefore to evaluate autocentrically with respect to y and k. Since presumption of accurate perception is default across the board, so is autocentric evaluation.

In other words, speakers are compelled pragmatically to take their own direct perceptions as indicative of truth, and since to have direct evidence that an individual has an experiential property *simpliciter* is just to have a certain experience stimulated in oneself by that individual, it follows that to treat one's own perceptions as dispositionally accurate just is to take the conditions of application for an experiential predicate to be tracked by one's own disposition to experience. But this is just what it is to evaluate autocentrically. Where speakers deviates from autocentric evaluation, they must commit to breaking this norm, and treat themselves as a dispositionally inaccurate perceiver with respect to the relevant stimulus and kind. But notice that this implies no commitment to descriptive error in of itself – only a mismatch between the evidence one has, and one's truth-conditional behavior.

This is where the explanation ends concerning experiential predicates for the present work. But this explanation rests on two deeper principles, which themselves can be questioned. First, why is direct evidentiality intrinsically autocentric? Second, why is there a robust expectation that speakers treat their own direct evidence as accurate? These plausibly have much deeper linguistic and cognitive explanations, of which the author is ignorant; but empirically they are apparently true, and from their combination with experiential semantics comes the default autocentric evaluation of experiential predicates.

5.2.3 Metasemantic competence and metaphysical commitments

Since speakers' linguistic acts and attitudes are hyperintensional, they simultaneously coordinate how speakers take the world to be, and the behaviors that speakers have towards the meanings of expressions. Above some idealized scenarios were represented, in which only one of these factors might be at stake, but for the most part they are both simultaneously in play to varying degrees, and this allows speakers both a great range of expressive power in

the attitudes they express linguistically, as well as the possibility for a lack of understanding stemming from imperfect knowledge of what other speakers' attitudes are, or lack of commitment to what exactly one's own attitudes are.

And so while speakers can be said to have a general *metasemantic competence*, in that they understand implicitly that expressions can be semantically underdetermined, and verbal as well as descriptive truth-conditional attitudes are possible because of this, there is always the threat of *metasemantic blindness*, where speakers are unaware of exactly why their attitudes converge or diverge, for what combination of descriptive or behavioral reasons. Linguistic exchanges trafficking in truth-conditional content simultaneously entail commitments among both these dimensions, without differentiating them explicitly. Where blindness occurs, it is possible for speakers to 'talk past each other:' they may not know exactly what someone else thinks about the world, or what exactly they are agreeing or disagreeing with them about, just in virtue of the fact that they understand an interlocutor to commit to the truth of a hyperintension.[32]

But the extent to which expressed attitudes are reflective of descriptive versus behavioral concerns can often be inferred from general linguistic norms. For traditionally non-evaluative predicates, it can typically be taken for granted that speakers' semantic alternatives, in virtue of their lexical semantic competence (cf. §5.1.2) relevantly converge, and so the expression of truth-conditional attitudes usually pertains to descriptive matters, except in 'borderline' cases. But in those cases, as with the debate over Secretariat in §5.1, descriptive issues are backgrounded in favor of behavioral ones, where what is at stake is the construal of the meaning of *athlete*, not the properties that Secretariat has. And so there is simply no

32. Metasemantic blindness does not have methodological problems, like semantic blindness (cf. fn. 2): it does not require any linguistic incompetence on the part of native speakers, but rather exists in virtue of speakers being ignorant, given their linguistic competence, of exactly what sorts of semantic and doxastic alternatives their interlocutors have. But this is just to say, unsurprisingly, that speakers are not omniscient, since it would be incredible if they were able to have such knowledge in general. Blindness is likely ubiquitous in realms of discourse where meanings don't have the benefit of everyday use to settle a robust lexical competence in speakers – disagreement on theoretical matters is probably robustly blind, while there is unlikely to be blindness in asking where the bus stop is. Where blindness is uncovered in the wild, there may be the expression of metasemantic competence *par excellence*: the speakers realize it was all 'just semantics.'

strict divide between evaluative and non-evaluative predicates at all, and evaluative discourse makes use of no special linguistic mechanisms – though evaluativity might still be a legitimate gradable and situational notion, defined in terms of the tendency for semantic alternatives to diverge.

Misalignment in semantic alternatives, which is at the root of the expression of attitudes having behavioral as opposed to merely descriptive import, can occur for more or less systematic reasons. At its most banal, this misalignment happens just in virtue of the fact that it would be a miracle if speakers' truth-conditional behaviors somehow all perfectly converged – and so some misalignment is just behavioral noise. But much of it is not, since semantic alternatives often misalign due to independently operating pragmatic norms (cf. §5.2.2), and where this does not happen, speakers can adopt differing alignments strategically in order to express matters of normative import. In both cases, the expression of 'opinion,' over and above mere truth-conditional content, becomes possible as a result of semantic underdetermination.

In the case of norm-driven misaligment, as with the autocentric evaluation of predicates, speakers are 'forced' to diverge in their semantic alternatives insofar as they follow norms of direct evidence and their experiential dispositions differ. The commitment to the truth of experiential hyperintensions therefore reliably signals what speakers' experiential dispositions are, insofar as they are norm-following, knowledgeable, and honest. Since other speakers are implicitly aware that this is happening, they may take such attitudes as an 'expression of opinion' – they reflect the speaker's 'tastes,' which their descriptive content truly targets. Where that expression of taste is at odds with one's own, one is free to acknowledge that while there is no descriptive fault in that speaker's commitments, it imposes a behavioral commitment that, insofar as one is also norm-following and honest, one cannot adopt oneself and be semantically consistent. Insofar as there are no normative stakes to this behavioral misalignment, the speakers can 'live and let live' – but as will be shown shortly, normative

concerns may fuel dispute nonetheless.[33]

When misalignment is not norm-driven, as when speakers disagree as to what counts as an athlete, the signal is less reliable. But nonetheless, the adoption of certain semantic alternatives may be purposeful, and therefore reflective of speaker opinion. If Alfonse loves horses, and wants to treat them with dignity, he may insist that Secretariat is an athlete, because to behaviorally construe the word *athlete* as applying to Secretariat is to include Secretariat in whatever honors are conventionally bestowed on individuals in using the underdetermined predicate (such as being included in lists of great sports figures). Alfonse thus signals his non-descriptive attitude that he treats racehorses with this dignity, merely by choosing to use a word in a certain way, and insofar as he pressures others to have a similar belief, he pressures them to behave in the same way.

So it is that behavioral disputes are not merely the result of blindness, but may be purposeful, and have serious, non-descriptive stakes. If a legal document uses the word *athlete* to confer certain rights on competitors, then whether Secretariat is an athlete may determine whether he gets those rights. And so the dispute over whether he is an athlete has normative consequences: consistent semantic behavior that obeys the legal document and takes *athlete* to apply to Secretariat has tangible effects. And so what is at stake is not what properties Secretariat has, but the behavioral construal of a word, which may entail treating Secretariat in a certain way. Whether a disagreement is 'real' or substantive is orthogonal to whether it is descriptive: just as there are behavioral disputes with pressing importance, there are descriptive disputes that don't matter.

With the above said, it should be clear how to answer the question of what the point is of using autocentrically-evaluated predicates in truth-conditional practices of assertion,

33. There is therefore an important sense in which expressions of autocentric attitudes are not 'opinions,' in that they are not freely adopted. One cannot in general control what experiences one is disposed to have – and neither are the pragmatic norms governing direct evidence chosen. And so insofar as one is norm-following, having certain experiential dispositions 'forces' one to believe certain things – not the mark of a freely-adopted opinion. The word *opinion* is of course itself semantically underdetermined, and so it is important to keep these varying dimensions in mind, and to appeal to them directly where possible, rather than speaking in terms of the vague language of 'opinion.'

agreement and disagreement, and so on (cf. Lasersohn 2005: §7.3, 2017: §9.3; MacFarlane 2007: §6). At some level, this question is ill-formed – what is the point of doing anything? Linguistic practices don't have intrinsic purposes, and given what they are, speakers can do whatever they want with them. But the present treatment offers an open-ended range of candidate rationales for why experiential predicates are used the way they are.[34]

The use of experiential predications, given autocentric norms, automatically conveys information about the experiential dispositions of speakers, and so can be used to inform other speakers about these dispositions. Their triggering of acquaintance inferences (cf. §4.1) also express information about what experiences speakers have actually had. Further, to the extent that semantic alternatives align, they provide factual information about the world, viz. what experiences individuals are disposed to stimulate, in a perfectly straightforward way (cf. §5.2.1 & fn. 29). Where Alfonse asserts *Licorice is tasty*, and Bethany believes him descriptively correctly, she learns something factual about the taste of licorice, viz. how it is disposed to affect her, even where both she and Alfonse evaluate autocentrically.[35]

But what the above-cited authors are particularly concerned with are cases where semantic alternatives knowingly misalign, and especially the fact that dispute about experiential matters can still potentially persist in these circumstances. There are an infinity of normative reasons why this might be, paralleling the case described above with *athlete* – what is

34. And so while there is nothing objectionable in principle about authors offering reasons for why experiential predicates are used a certain way, as with Lasersohn's (2017: 211) 'practical advantage' account, or Coppock's (2018: 154) notion of 'seeking a common outlook,' these accounts can be misleading in that they unnecessarily privilege one sort of thing speakers *can* do with these predicates, as if this rationale were suggested by the linguistic mechanism itself. This isn't so – speakers can just as easily use experiential predicates to make their disagreements clear, or berate others and persist in snobbery, as they can to seek a common opinion. What speakers use these predicates for is as open-ended as human rationality, or irrationality.

35. It is another question why Bethany might believe Alfonse on the basis of his assertion – where she is aware of mutual autocentric evaluation, she may only do so on the basis of presuming that Alfonse's tastes track her own with respect to the tastiness of licorice. It doesn't particularly matter why she believes him – still, she learns something from his assertion, so long as she really is disposed to like licorice's taste. Note that Bethany, and anyone else, regardless of whether they share experiential dispositions with her, can also use what Alfonse asserts in truth-conditional reasoning: *If licorice is tasty, I should buy some*; *Licorice is tasty*. Conclusion: *I should buy some*. This is because committing to the truth of the hyperintensions of the first two sentences commits one to the truth of the third, insofar as one is semantically consistent, *regardless* of what one's semantic alternatives decide for *tasty* – and semantic consistency is a prerequisite for reasoning from premises to conclusions.

taken as *tasty simpliciter* according to speakers' semantic behaviors may determine what is accorded the benefits of being described with that word. And as with the cases in §4.2.1 and Anthony (2016), to refuse to legitimate certain autocentric evaluations is to insist that certain experiential reactions are defective – but now it is clear why this defectiveness may entail a normative refusal to let speakers 'get away with' having those reactions uncriticized, without entailing attribution of a descriptive error to them.[36]

MacFarlane (2011a: 445-446 & fn. 5; 2014: §1.1) worries that treating experiential predicates as not relativized to an experiencer in their meaning commits autocentric speakers to a combination of systematic error and chauvinism. This worry depends on the idea, articulated in §5.1.1, that such a lack of relativization requires that some one property be truth-conditionally denoted by such predicates *simpliciter*. If autocentricity causes speakers to systematically diverge in their evaluations in virtue of their differing experiential dispositions, it follows that many autocentric speakers must not only be mistaken about which individuals have which experiential properties, but also chauvinistic in that they take their own experiences to track these properties, even in the face of overwhelming disagreement.

But this worry disappears once it's understood that nearly all predicates are semantically underdetermined, and do not in virtue of their meaning denote a single property *simpliciter* (cf. §5.1.1). Further, since experiential predicates are metasemantically constrained by speakers' experiential dispositions in autocentric evaluations (cf. §5.2.1), it follows that in virtue of evaluating autocentrically, speakers do not commit to anything more than the existence of self-directed dispositional properties, regarding what is disposed to stimulate

36. Because the 'substance' of disagreement has so often been assumed to lie in descriptive dispute, the relativist literature, like that of the blindness- and error-theorists described in §5.1, is rife with worries that discourse about experiential matters is somehow defective. MacFarlane (2007: 30): "From lofty philosophical heights, the language games we play with words like 'funny' and 'likely' may seem irrational. But that is no reason to deny that we do play these games, or that they have a social purpose." Cf. Stephenson (2007a: 493). But it should be clear from the above that non-descriptive experiential disputes, like all behavioral disputes with potential normative consequences, may be no game, and not even seemingly irrational. That is not to deny that there is room for genuinely blind or irrational disputes, as there must be, since they are real – but the focus on descriptive content of attitudes has resulted in theoretical blindness to the varieties of disagreement.

experience in them, and their descriptive commitments pertain only to these properties. There is therefore no metaphysical commitment to experiential properties *simpliciter* in virtue of autocentric evaluation of bare experiential predicates, nor is there any chauvinism regarding one's own experiences tracking matters of fact in a privileged way. This is not to say that the semantics of these predicates rule out realism or chauvinism, only that they don't systematically follow from the semantic machinery: and so likewise the sort of error theory in virtue of semantic behavior that Hirvonen (2016) espouses is mistaken.

It is also possible in principle, given the semantics of experiential predicates, for an autocentric speaker (or a non-autocentric one) to be a realist about experiential properties *simplicitier*, and to take their own application of experiential predicates to track these – this would only require taking certain experiential dispositions to be somehow privileged in tracking certain other facts about the world, whatever these might be, so that the relevant experiencer-anchored properties coincide with them. But what this might entail, and whether speakers do this, is a matter of folk metaphysics, not of semantics, and speakers' linguistic competence doesn't hinge on these sorts of beliefs.[37]

One might then be tempted to say that it is only autocentricity in concert with a folk metaphysical belief in experiential properties *simpliciter* that yields chauvinism, but even this is not right. Chauvinism about matters of experience and metaphysical realism about experiential properties are orthogonal: the former pertains to how uncompromising a speaker is with respect to the legitimacy of their own autocentric experiential predications over and against those of others. A food critic might be so about tastiness, even where entirely anti-realistic about experiential properties, because of the normative stakes involved in judging the quality of food, or out of a desire to align the treatment of experiential predicates in

37. There is some preliminary evidence to suggest that most speakers of American English are in fact not realists about the properties denoted by experiential predicates: cf. Cova & Pain (2012). The methodology is questionable (cf. Hirvonen 2014: §15.2 for criticisms), and it isn't entirely clear what speakers are being asked as far as its relevance to belief in experiential properties *simpliciter* goes, or how to make them understand what they are being asked, even if it were clear. But insofar as what empirical tests for intuitions probe is a metaphysical matter, it doesn't strictly pertain to the semantics of experiential predicates.

the language with a certain class of favored experiential dispositions. Conversely, robust metaphysical realists about such properties may not have this kind of stringency, and let people say whatever they want regarding experiential properties, because they take the matter to be epistemically open, or simply don't care.[38]

5.3 Semantic sensitivity to metasemantics and experience

§5.1.2 and §5.1.3 demonstrated how the compositional semantics is sensitive to metasemantic concerns, in treating attitude reports as hyperintensional. But there are other constructions that deal with the metasemantic portion of the grammar in more explicit ways, and some of these deal crucially with how metasemantics interacts with direct evidence, and so with speaker experience. The relationship between semantics and metasemantics is therefore not one-directional: the composition must have access to models of the language, and so to possible truth-conditional behaviors of speakers.

Further, as §5.2.2 showed by example, metasemantic and pragmatic facts interact with the compositional semantics in regular ways, to yield phenomena relevant to truth-assessment, like autocentricity. Metasemantic concerns must therefore be traceable to their origin in the compositional semantics, and explained in terms of the model-sensitivity it introduces. §5.3.1 demonstrates how certain expressions marking subjective attitudes benefit from explicit reference to metasemantic concerns in the semantics. §5.3.2 concludes this work by summarizing how this chapter's proposals provide a unified characterization of linguistic subjectivity, and offers some suggestions as how to locate and treat such subjectivity where it appears.

38. In fact there is even less connection between belief and attitude here, because a realist about experiential properties could very well, in light of the fact that autocentric speakers only commit to the existence of self-directed dispositional experiential properties, take experiential predicates not even to denote these properties in the first place.

5.3.1 Subjective attitudes and dative marking

This section treats subjective attitudes, *dativus iudicantis*, and overt experience phrases (cf. §2.3.3) with the new metasemantic machinery. It further demonstrates how the verb *find* and overt experiencer phrases are sensitive to speakers' experiences in a special metasemantic way.

Recall from §1.1.1 the basic data concerning the subjective attitude verb *consider*. It is a doxastic verb that can embed a small clause, and has differing behavior depending on the choice of predicate in that clause. Where the predicate is evaluative, the verb's contribution is difficult to distinguish from that of *think* (420a); where the predicate is not evaluative, its use implies that specifically evaluative, as opposed to merely descriptive, beliefs verify the complement clause (420b); and where it is difficult to construe how evaluative beliefs are relevant for the truth of the complement clause, its use is often simply infelicitous (420c).

(420) a. Alfonse considers licorice tasty.

b. Alfonse considers the table wooden.

c. ?Alfonse considers Bethany six feet tall.

Consider can be treated as having the same asserted truth-conditional content as *think*, in line with its use to report beliefs. But it carries a metasemantic presupposition, that there is some model or other, relative to which the believer's doxastic alternatives *falsify* the hyperintension of the embedded clause.

(421) $[\![\text{consider}]\!]^{\mathcal{M},w} = \lambda\Phi_{m,st}.\lambda x_e : \exists\mathcal{M}'[\forall w' \in Dox_{x,w}[\neg\Phi(\mathcal{M}')(w')]].$
$\forall w' \in Dox_{x,w}, \mathcal{M}' \in Sem_{x,w}[\Phi(\mathcal{M}')(w')]$

The asserted content following the domain restriction, on the second line of the denotation, is equivalently denoted by '$believe''(w)(\Phi)(x)$.' The presupposition appears as a domain restriction on the agent x: it requires that, holding the agent's beliefs about the world constant, those very beliefs falsify the hyperintension at some model. In short, what

is presupposed is that the agent's holding the belief according to the asserted content, as opposed to its negation, is non-trivially dependent on the pattern of behavior that the agent adopts with respect to the meaning of the embedded clause.

This presupposition has similarities to the subjective contingency presupposition proposed by Bouchard (2012: §3.5), as well as the counterstance-contingency presupposition offered in Kennedy & Willer (2016) – models distinct from the one relative to which an expression is evaluated can be thought of as 'counterstances,' i.e. alternate ways that linguistic behavior could construe the truth-conditional content of linguistic expressions. With that said, the basic data are derived from this denotation as follows. For (420a), the result is:

(422) $[\![$Alfonse [considers [licorice [GEN [POS tasty]]]]]$]\!]^{\mathcal{M},w}$
$= \forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M},\text{PLEASURE}}(w')(\iota y[taste''(w')(l)(y)])]$,
if $\exists \mathcal{M}'[\forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta'_{\mathcal{M}',\text{PLEASURE}}(w')(\iota y[taste''(w')(l)(y)])]]$;
else undefined

And so *Alfonse considers licorice tasty* asserts that Alfonse thinks that licorice is tasty, and presupposes that his beliefs about the world falsify *that licorice is tasty* relative to some model.[39] The presupposition is met, since whatever relevant beliefs about licorice that Alfonse has regarding its taste (in whom it is or is not disposed to produce pleasure, for instance), a model is available to falsify the hyperintension of the complement clause given those beliefs. Just take any model $\mathcal{M}''$ whose assignment of degrees of pleasure to stimuli is inverse to Alfonse's semantic alternatives, as follows.

(423) For all w'' and z:
$0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta_{\mathcal{M}'',\text{PLEASURE}}(w'')(z)$ iff
$\forall \mathcal{M}' \in Sem_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta_{\mathcal{M}',\text{PLEASURE}}(w'')(z)]$

Then Alfonse's beliefs about the world falsify *that licorice is tasty* relative to $\mathcal{M}''$.

39. Because *consider* is an evidentially neutral doxastic verb, just like *think* it will also imply that Alfonse believes what he does about licorice due to having tasted it, via the acquaintance inference: cf. §4.1.4.

(424) $\forall w' \in Dox_{a,w}[0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta'_{\mathcal{M}'',\text{PLEASURE}}(w')(\iota y[taste''(w')(l)(y)])]$

The models existentially quantified over in the presupposition can be taken to be domain-restricted, so that some model stereotypically 'available' in linguistic practice can act as the falsifier. In the case of models that differently assign degrees of experiential kinds to stimuli, such a falsifying model can always be chosen with evaluative predicates in predicate position, and so *consider*-reports trivially meet the presupposition. Hence they are difficult to distinguish from *think*-reports.[40]

For a predicate like *wooden*, which is traditionally non-evaluative, but like *athlete* (cf. §5.1.1) is semantically underdetermined, the result is as follows (taking *table* not to be underdetermined, for simplicity).[41]

(425) $[\![\text{Alfonse [considers [[the table] wooden]]}]\!]^{\mathcal{M},w}$
$= \forall w' \in Dox_{a,w}, \mathcal{M}' \in Sem_{a,w}[wooden'_{\mathcal{M}}(w')(\iota y[table'(w')(y)])]$,
if $\exists \mathcal{M}'[\forall w \in Dox_{a,w}[\neg wooden'_{\mathcal{M}'}(w')(\iota y[table'(w')(y)])]]$;
else undefined

And so *Alfonse considers the table wooden* asserts that Alfonse thinks that the table is wooden, and it presupposes that at some model (within the relevant restricted domain), Alfonse's beliefs about the world instead determine the table not to be wooden.

40. There is apparently a subtle difference between *consider* and *think*, even in these contexts. Where Alfonse and Bethany have the exact same experiential reactions to licorice, and their opinions are being compared, *Alfonse considers licorice tasty*, in a context where experiential reactions are not at issue, may imply that what Alfonse's belief turns on is which experiences Alfonse takes to be indicative of which degrees of pleasure, or what kind of experience is necessary to have a non-zero degree of pleasure. This would be a case where metasemantic concerns irrelevant to experiential dispositions *per se* are being presupposed to affect the truth of the belief. *Alfonse thinks that licorice is tasty* apparently doesn't presuppose this, though it can of course report Alfonse's belief, which he may have for whatever reason, hinging on experiential matters or not.

41. There are issues here having to do with the clausal structure of *consider*: the reported belief apparently cannot concern whether the subject of the embedded clause counts as a table, but only whether the table counts as wooden. Cf. Kennedy & Willer (2016: §3.2) for discussion, and Sæbø (2009: §4) for related observations with *find*. Plausibly, this is unrelated to the semantics of *consider* itself, but has to do with the subject of this sort of small clause always being read *de re*, so that in this case it is outside the scope of the hyperintensional attitude.

Suppose that Alfonse's beliefs about the world are such that he thinks that the table is made of particle board, and that it is a non-trivial question whether something made of particle board is wooden, even given full knowledge of its material composition. The asserted content of (425) ensures that Alfonse's semantic behaviors, given this belief, treat the hyperintension of *the table is wooden* as true: according to Alfonse, particle board counts as wood. The presupposed content takes there to be some available model relative to which Alfonse's very same beliefs, i.e. that the table is made of particle board, falsify that the table is wooden.

Where the asserted content is true, this occurs just in case, as is plausible for English, there is an available model $\mathcal{M}''$ such that for any x made of particle board at w, $\neg wooden'_{\mathcal{M}''}(w)(x)$. This is a 'stricter' model than those in Alfonse's semantic alternatives, which rules out things made of particle board as *not* wooden. But then, since Alfonse believes that the table is made of particle board, relative to *this* model, according to his beliefs about the world, the table is not wooden.

(426) $\forall w' \in Dox_{a,w}[\neg wooden_{\mathcal{M}''}(w')(\iota y[table'(w')(y)])]$

And so the presupposition is satisfied. Unlike with *tasty*, this presupposition is not trivial, but depends on what Alfonse's descriptive beliefs (tracked by his doxastic alternatives) are. Given his beliefs about the world, his linguistically relevant behavior must treat the table as wooden, in such a way that if he changed those behaviors relevantly, even believing nothing different about the world, he would come to believe that the table was not wooden. Where this is not so, the presupposition is not satisfied.[42]

Suppose instead that Alfonse believes the table to be made of oak. What is then pre-

42. Clear cases of *consider*-reports exhibiting this pattern are those whose predicate is a positive-form scalar adjective along a single dimension, like *tall*. Thus, *Alfonse considers Bethany tall* asserts that Alfonse thinks that Bethany is tall, and presupposes that his beliefs about Bethany's height are such that, were he to treat the vague cutoff for *tall* differently, he would believe her not to be tall, even where his beliefs about the world are held constant. This follows straightforwardly if the vague cutoff contributed by the positive form of relative adjectives has a number of models available that set this cutoff at different heights. Traditional scalar vagueness is therefore just semantic underdetermination that happens to be along an obvious scalar dimension, and is a metasemantic phenomenon *par excellence*.

supposed is that there is a model relative to which Alfonse's belief *that the table is made of oak* falsifies that the table is wooden. But if oak tables are not wooden, it is not clear what is: such a presupposition is difficult to accommodate simultaneously with the truth of the asserted content, unless some reason can be found to existentially quantify over very strange models, e.g. ones that take oak things not to be included in the extension of *wooden* at a world. Since the asserted content and presupposition clash, infelicity results.

In the extreme case, as with (420c), a *consider*-report may result in outright infelicity, where no plausible falsifying model can be found regardless of the agent's descriptive beliefs. This report comes out as follows, making some standard assumptions about the composition: *tall* denotes a measure function, mapping individuals to degrees of height (427a); and *six feet* is a piece of degree morphology, setting the standard of comparison to the degree to which individuals are mapped by the measure function when they are at least six feet along the relevant dimension (427b).

(427) a. $[\![\text{tall}]\!]^{\mathcal{M},w} = \lambda x_e.\delta'_{\text{HEIGHT}}(w)(x)$

b. $[\![\text{six feet}]\!]^{\mathcal{M},w} = \lambda G_{s,et}.\lambda x_e.d_{6ft,G} \leq_G G(w)(x)$

c. $[\![\text{Alfonse [considers [Bethany [[six feet] tall]]]}]\!]^{\mathcal{M},w}$
$= \forall w' \in Dox_{a,w}, \mathcal{M}' \in Sem_{a,w}[d_{6ft,\text{HEIGHT}} \leq_{\text{HEIGHT}} \delta'_{\text{HEIGHT}}(w')(b)]$,
if $\exists \mathcal{M}'[\forall w' \in Dox_{a,w}[d_{6ft,\text{HEIGHT}} \leq_{\text{HEIGHT}} \delta'_{\text{HEIGHT}}(w')(b)]]$;
else undefined

It is difficult to see how this presupposition can be satisfied, given that if Alfonse's beliefs about Bethany's height place her at at least six feet tall, it is unclear what alternate model would cause his beliefs to falsify *that Bethany is six feet tall*. This is represented in the denotation by not including model-sensitivity in any of the constituent lexical items: taking such a denotation seriously, the presupposition is impossible to satisfy, and so the report is infelicitous *simpliciter*.

But there are ways to rescue the presupposition in the right contexts (cf. Ch. 1, fn. 14). Suppose that Bethany reaches above six feet in platform shoes, but is very slightly

less than six feet tall out of them – if she normally wears these shoes, then Alfonse might employ a model that maps Bethany's height to at least six feet. Another model that refuses to make this allowance based on shoes would not do so, and falsify the belief, rescuing the presupposition. And in fact, such an odd case is precisely one in which the *consider*-report is felicitous, due to model-sensitivity on the height mapping tracked by the δ'-term.

These fine-grained variations in the felicity of *consider*-reports therefore generally depend on facts about the world and agents' descriptive beliefs – this means that the existential quantification over models can be seen as context-invariant (though it may itself be subject to metasemantic variation among speakers). If one instead wants these changes in felicity to be genuinely context-dependent, then an indexical domain restriction can be baked into this quantification, allowing cross-contextual leeway (cf. the treatment of exocentric domain restriction in §2.3.2).

This subjective attitude verb *find* has related metasemantic behavior, but following Stephenson (2007b: 61), it can be treated as pertaining to direct evidence in a special way. In brief, *find*-reports assert that their agents have direct evidence for the hyperintension of their embedded clauses, and presuppose that this direct evidence is impossible for any agent other than the subject of the attitude to have. Using the characterization of direct evidence in §5.2.2, the denotation is as follows.

(428) $[\![\text{find}]\!]^{\mathcal{M},w} = \lambda\Phi_{m,st}.\lambda x_e : \neg\exists y, w' \neq x, w[\forall w'' \in Per_{y,w'}[\Phi(\mathcal{M}_{x,w})(w'')]].$
$\forall w'' \in Per_{x,w}[\Phi(\mathcal{M}_{x,w})(w'')]$

Again, the presupposition is encoded as a domain restriction on the agent. It enforces that there is no agent-world pair non-identical with the agent and world to which the direct evidence in question is anchored, at which there is direct evidence verifying the hyperintension of the embedded clause as it anchors to the agent.

The data surrounding *find* are quite complex (cf. Vardomskaya 2018: ch. 4 for an overview), but it is worth working through some simple cases using this denotation with adjectival small clauses. *Find*-reports commit their agents to certain experiential reactions

when embedding experiential predicates (429a); they impose presuppositions of direct experience, that e.g. project out of negation (429b) (cf. Hirvonen 2014: §4.3.2); and they are infelicitous with non-experiential and traditionally non-evaluative predicates (429c).

(429) a. Alfonse finds licorice tasty.

b. Alfonse doesn't find licorice tasty.

↪ Alfonse has tasted licorice.

c. #Alfonse finds the table wooden.

The denotation for (429a) is as follows.

(430) $[\![\text{Alfonse [finds [licorice [GEN [POS tasty]]]]}]\!]^{\mathcal{M},w}$

$= \forall w' \in Per_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M}_{a,w},\text{PLEASURE}}(w')(\iota z[taste''(w')(l)(z)])]$

$= \forall w' \in Per_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w')(a)(\iota z[taste''(w')(l)(z)])]$

$= 0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(a)(\iota z[taste''(w)(l)(z)]),$

if $\neg\exists y, w' \neq a, w$

$[\forall w'' \in Per_{y,w'}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M}_{a,w},\text{PLEASURE}}(w'')(\iota z[taste''(w'')(l)(z)])]];$

else undefined

The asserted content is that licorice's taste has produced pleasure in Alfonse, as desired to capture the basic truth conditions pertaining to experiential reactions. A presupposition is also introduced here, that Alfonse has tasted licorice, due to the definedness conditions on the δ''-term on the fourth line of (430): the sentence is true just in case the experience that the taste of licorice has stimulated in Alfonse instantiates a non-zero degree of pleasure, and so there must be an experience stimulated in Alfonse by the taste of licorice. This projects out of the usual environments, including negation, where the asserted content and direct experience presupposition of (429b) is as follows (ignoring the further presupposition).

(431) $[\![\text{Alfonse [doesn't [find [licorice [GEN [POS tasty]]]]]}]\!]^{\mathcal{M},w}$

$= 0_{\text{PLEASURE}} =_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(a)(\iota z[taste''(w)(l)(z)])$

This presupposes the exact same thing, again due to the definedness conditions of the δ''-term: it is defined where the taste of licorice has produced an experience in Alfonse, and where defined is true just in case the maximal such experience instantiates a zero-degree of pleasure.

Where experiential predicates are embedded, the further presupposition encoded by *find*, that no other individual could have the relevant direct evidence verifying the hyperintension of the embedded clause, is also met. This can be shown by *reductio*.

(432) a. Let there be some $y, w' \neq a, w$, such that:
$\forall w'' \in Per_{y,w'}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M}_{a,w},\text{PLEASURE}}(w')(\iota z[taste''(w')(l)(z)])]$

b. By the equivalence in (412),
$\forall w'' \in Per_{y,w'}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w'')(a)(\iota z[taste''(w'')(l)(z)])]$

c. But then by (320) in §4.1.3,
$y = a \wedge \delta''_{\text{PLEASURE}}(w')(y)(\iota z[taste''(w')(l)(z)])$

d. Since by assumption $y \neq a$, this is false, and so the assumption in (432a) is false. Therefore, the presupposition is satisfied.

And so with experiential predicates generally, the presupposition is always met, since no agent other than the subject of the attitude can have direct evidence that the subject of the attitude had a certain experience. This is not so with predicates like *wooden*, however, in which case the denotation is as follows.

(433) $[\![\text{Alfonse [finds [[the table] wooden]]}]\!]^{\mathcal{M},w}$
$= \forall w'' \in Per_{a,w}[wooden'_{\mathcal{M}_{a,w}}(w')(\iota z[table'(w'')(z)])]$,
if $\neg\exists y, w' \neq a, w[\forall w'' \in Per_{y,w'}[\neg wooden'_{\mathcal{M}_{a,w}}(w')(\iota z[table'(w)(z)])]]$;
else undefined

There is nothing wrong with the asserted content: it is true just in case Alfonse has direct evidence verifying that the table is wooden. This is the sort of thing to which he would be committed if he uttered *The table is wooden* while committing to a direct evidential source.

But the presupposition cannot be satisfied. This is because there are an agent y and world w', such that y's direct evidence at w' verifies the hyperintension of *that the table is wooden* at the model anchored to Alfonse at w. The semantics of *wooden* introduces no sensitivity to direct evidence that makes the anchoring to a, w relevant for direct experience: and so any agent who has the same relevant perceptions of the world that Alfonse has will have the very same direct evidence for the hyperintension that he does. Therefore, the report is infelicitous, since the presupposition of the inherent uniqueness of Alfonse's direct experience is not met.[43]

The infelicity of *find*-reports is in general more categorical than that of *consider*-reports, and tends to hold in virtue of the lexical semantics of predicates, tracking whether they contain some component specially sensitive to direct evidence, in such a way that there can in principle be direct evidence for the relevant hyperintension that only one agent can have. However, where speakers can construe the same model-sensitive contribution as either sensitive to the anchoring of direct evidence to an agent or not, there may arise the sort of fine-grained differences in felicity conditions seen for *consider*-reports.[44]

Pace Stephenson (2007b: 61) and Kennedy & Willer (2016: 922), *find* is not a doxastic verb. It asserts that the agent has direct evidence verifying the hyperintension, which in turn entails belief, where agents treat themselves as accurate perceivers (and so are evaluating

43. This dispels Sæbø's (2009: 334)'s worry that direct experience is insufficient for the felicity of a *find*-report: the additional requirement of this uniqueness, which entails the direct 'unobservability' of the relevant experiential state as relativized to the experiencer, is required. Sæbø's own examples make use of sexual orientation: this account would therefore predict that predicates of sexual orientation are not unobservable in this sense.

44. A plausible example of such a case is with embedding positive-form relative gradable adjectives beneath *find*, as with *Alfonse finds Bethany tall*. There is disagreement in the literature as to whether these sentences are felicitous: Sæbø (2009: §6.1) assumes that they are, while Kennedy (2013: 263 ff.) assumes that they are not. Plausibly, what is happening is that the model-sensitivity in the vague standard contributed by the positive form is being construed as experientially sensitive in the right way by some speakers, and not by others: perhaps how tall one must be to be tall *simpliciter* can be taken as a matter of an individual 'striking' one as tall, and so producing an experiential reaction that only a particular agent can have. In general, where *find*-reports are of uncertain felicity across speakers, this should be a result of speakers differing as to whether the metasemantic issue is one that a particular agent's direct evidence can non-trivially decide. Vardomskaya (2018: ch. 4) addresses a range of constructions that might be construed as relevant to direct experience in this way.

autocentrically: cf. §5.2.2) and are cognizant of their own experiences. Either of these two conditions may fail to be met for whatever reason, in which case the link between direct evidence and belief fails, and so does the commitment to belief in virtue of a *find*-report.

(434) a. Alfonse doesn't find licorice tasty, but he thinks that it is.

b. #Alfonse doesn't consider licorice tasty, but he thinks that it is.

Thus (434a) is ordinarily strange, but not where Alfonse is an autocentric evaluator who has forgotten what licorice tastes like to him, and so holds a descriptively mistaken belief, or where Alfonse is not evaluating autocentrically, and so his opinion of licorice's taste is not tracked by his own experiences (cf. the sorts of cases described in Anthony 2016, and §4.2.1). These circumstances can't save the felicity of the corresponding *consider*-report (434b).

First-person *find*-reports also typically engender Moorean effects accompanied by denial of belief, but again not where autocentric evaluation is abandoned: (435a) might be uttered by a speaker who has seen a movie so many times that they are no longer interested by it, but hold that it does in fact produce interest in ideally normal conditions (that their fatigued experiences no longer track). These concerns are not relevant for *consider*-reports, which are genuinely doxastic, and so produce Moorean effects *simpliciter* (435b).[45]

(435) a. I don't find the movie interesting (anymore), but it is interesting.

b. #I don't consider the movie interesting (anymore), but it is interesting.

This treatment of subjective attitudes can then be imported into the treatment of *dativus iudicantis*. Recall from §2.3.3 that this is the phenomenon of relativizing predicates to individuals using free dative marking, as in (436).

(436) If there's alcohol, that's a party to me.

There is a compelling similarity with *consider* here: a paraphrase might be, *If there's alcohol, I consider that a party.* And so to relativize a predicate to an agent generally can

45. Though note that there is also a non-doxastic, non-stative meaning of *consider* that does not produce these Moorean effects: cf. Vardomskaya (2018: 175).

be seen as invoking the same metasemantic mechanisms as subjective attitudes. The *to* of *dativus iudicantis* is then as follows.

(437) $[\![\text{to}]\!]^{\mathcal{M},w} = \lambda x_e.\lambda\mathcal{P}_{m,\langle s,et\rangle} : \exists\mathcal{M}'[\forall w' \in Dox_{x,w}[\neg\mathcal{P}(\mathcal{M}')(w')]].$
$\lambda y_e.believe''(w)(\lambda\mathcal{M}'_m.\lambda w'_s.\mathcal{P}(\mathcal{M}')(w')(y))(x)$

This is just an importation of the behavior of *consider* into a property modifier: *party to x* is then just a predicate true of individuals that x considers a party.

A similar move then suggests itself for the metasemantic treatment of overt experiencers encoded by *to-headed* PPs, as in *tasty to Alfonse*. At the end of §2.3.3, it was noted that the account given there did not take into account the fact that such experiencer phrases have a limited distribution, not appearing with exocentric predicates, not stacking with other experiencer phrases, not occurring with non-evaluative predicates, and not occurring with predicates embedded beneath *find*. With the metasemantics established, the reason for this can be given: these experiencer phrases encode sensitivity to direct evidence just like *find* does. In other words, *tasty to Alfonse* is a predicate true of those individuals that Alfonse finds tasty.

(438) a. $[\![\text{to}]\!]^{\mathcal{M},w} = \lambda x_e.\lambda\mathcal{P}_{m,\langle s,et\rangle} : \neg\exists y, w' \neq x, w[\forall w' \in Per_{y,w'}[\mathcal{P}(\mathcal{M}_{x,w})(w')]].$
$\lambda y_e.\forall w' \in Per_{x,w}[\mathcal{P}(\mathcal{M}_{x,w})(w')]$

b. $[\![[\text{GEN } [\text{POS tasty}]] \text{ [to Alfonse]}]\!]^{\mathcal{M},w}$
$= \lambda y_e.\forall w' \in Per_{a,w}[0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta'_{\mathcal{M}_{a,w}}(w')(\iota z[taste''(w')(y)(z)])]$
$= \lambda y_e.0_{\text{PLEASURE}} <_{\text{PLEASURE}} \delta''_{\text{PLEASURE}}(w)(a)(\iota z[taste''(w)(y)(z)])$

And so felicitous *find*-embeddability and occurrence with overt experiencer PPs are taken to pattern with each other, and occurrence of an experiencer phrase and embedding beneath *find* are taken to have the same semantic effect. This may ultimately be a simplification, but by and large is apparently true.[46] The reader can confirm that overt experiencers then

46. Compare the similarities between *Alfonse finds the book helpful* and *The book is helpful to Alfonse*, or *Alfonse finds climbing the mountain easy* and *Climbing the mountain is easy to Alfonse*. In each case, the

appear in the right circumstances, due to the relevant presupposition. In particular, they can only adjoin to experiential predicates once they are 'genuinely' experiential or evaluative, viz. above GEN.

5.3.2 The semantic origin and metasemantic nature of linguistic subjectivity

The issues surrounding the semantics of experiential predicates have often been framed in terms of linguistic 'subjectivity' (e.g. in Kennedy 2013; Fleisher 2013; Bylinina 2017; Vardomskaya 2018). The idea is that certain truth-conditional features of language use are not hard-fixed by the linguistic conventions of a speech community as a whole, but may depend on the choices or preferences of individual speakers in that community, leading to licensed variation in truth assessment among those speakers, even where they use the same expressions. Since Kölbel (2002), researchers have suspected that this sort of subjectivity is distinct from context-sensitivity, by which the truth-conditional content of an expressions is sensitive to its context of use, and since Lasersohn (2005), experiential predicates have been thought to be paradigmatically subjective expressions.

In seeking to provide an adequate semantic treatment of experiential predicates, the present work has outlined a way to treat the metasemantics of expressions generally, and this in turn has offered a construal of linguistic subjectivity as a unified phenomenon. Semantic variation among a speaker population, represented in the compositional semantics as model-sensitivity, results from semantic underdetermination. It follows from this underdetermination that speakers are capable of disposing themselves differently towards the meanings of the same expressions, so that they can in principle construe those same expressions as

connection with direct evidence that only a particular agent can have is clear: to have direct evidence that something is helpful is to be helped by it; to have direct evidence that something is easy is to do it with little effort. Thus direct evidence is a far broader category than that pertaining to experiential semantics alone, even if the two interact interesting ways. Note also that the specifically experiential treatment extends to other constructions as well: *Alfonse finds Bethany admirable* looks to say that Alfonse has direct evidence that Bethany is admirable, which is just to say that he admires her (in fact *-able* looks to be a sort of subject-experiencer counterpart to *-ing*).

having distinct truth-conditional contents. This, on the present picture, is what linguistic subjectivity is: a metasemantic phenomenon, and one that a merely semantic portion of the grammar is blind to.

Semantic underdetermination, and therefore linguistic subjectivity, is ubiquitous in natural language. The present work has confronted only those instances of it that were necessary for the explication of a semantics of experiential predicates, but even so has encountered: (i) the general underdetermination of almost all lexical predicates, concerning the fringes of their canonical application; (ii) the lack of fixed norms to provide modal bases for generic quantification; (iii) the underdetermination of norms deciding which experiences evidence the truth of which propositions; and (iv) variation in the size preordering on qualities, including portions of experience. But as soon as it is looked for, semantic underdetermination will be found everywhere: e.g. in vagueness (cf. fn. 42; Richard 2004), in multidimensionality (cf. Sassoon 2013), and in the setting of conversational backgrounds for the evaluation of modals (cf. the references concerning epistemic modality below). Not only is it everywhere, but speakers' metasemantic competence recognizes this (cf. §5.2.3), as does the grammar itself, which has semantic mechanisms specifically devoted to metasemantic concerns (cf. §5.3.1).

The idea that subjectivity pertains in a special way to lexical items that are sensitive to matters of opinion ought to be abandoned. As noted in §5.2.1, linguistic subjectivity is not even primarily about opinion to begin with, nor about the circumstances in which speakers assess semantic contents – both these things piggyback on a deeper and more basic fact about linguistic usage, that it is fundamentally underdetermined with respect to its truth-conditional content. This fact is presumably rooted in non-linguistic facts about the impossibility of coordinating the behaviors on which semantic facts supervene.

Where model-sensitivity simply tracks this undertermination, it might be said to reflect *weak* or *non-convergent* subjectivity. In these cases, the undetermination alone is what drives the subjectivity, though it can be co-opted to express various normative beliefs, or beliefs reflective of individual opinions (cf. §5.2.3). But there are certain cases, where

the model-sensitivity in the compositional semantics interacts with independent linguistic norms, and the result is not mere underdetermination, but a systematic anchoring of the construal of truth-conditional content to speakers. As §5.2.1 and §5.2.2 showed, this occurs with autocentricity in experiential predicates, since the underdetermination accompanying generic quantification interacts with norms on direct evidentiality. This latter sort of self-directed underdetermination might be called *strong* or *divergent* subjectivity. It does not merely leave matters of usage open, but 'forces' them to differ among the population, as speakers differ in ways relevantly tracked by the norm. Another plausible candidate for such strong subjectivity is e.g. the anchoring of the knowledge base for epistemic modality to the evaluator (cf. Egan et al. 2005; Stephenson 2007a; MacFarlane 2011b).[47]

The above, especially §5.2.2, implies some methodological constraints as to how linguistic subjectivity should be studied and modeled. First, wherever it is thought to occur, it must be traced to a concrete source in the compositional semantics. Second, wherever there is thought to be strong subjectivity, it must be shown to be parasitic on weak subjectivity: the source of the weak effect must be found, along with some independently attested norm that interacts with it to produce the strong effect. The existence of subjectivity should never be posited for constructions that do not systematically give rise to the same sort of underdetemrination in virtue of their compositional semantics elsewhere in a predictable way, nor should strong subjectivity ever be appealed to where independent norms interacting with weak subjectivity cannot explain its existence.

Strong subjectivity has the flavor of deontic necessity: insofar speakers are obeying certain pragmatic norms, they must behave semantically in such a way as to construe the truth-conditional content of certain expressions as depending on features about themselves. This

47. Coppock (2018: 126-127) has pointed out that certain Swedish subjective attitude verbs do not felicitously embed epistemic modals. This is a piece of data that cannot be ignored in treating epistemic modality and its relation to metasemantic underdetermination more generally. Whether certain propositions are 'subjective' or not cannot simply be read off their embedding behavior – but if a decent account of subjective attitudes is along the lines suggested in §5.3.1, then the idea that epistemic modals are driven by a distinct mechanism must be taken seriously. Only an evaluation of an account of the two constructions together can bear on this question.

necessity is ultimately a hypothetical one: one can in the end say whatever one pleases, but one must deal with the normative consequences of this (for experiential predicates, presenting oneself as an inaccurate perceiver when refusing to evaluate autocentrically). When obeying the force of this norm, usage of strongly subjective expressions automatically conveys information about speakers (e.g. what experiential dispositions they take themselves to have).

Weak subjectivity has the flavor of deontic possibility: the linguistic conventions constituted by speakers' behaviors leave it open to some extent how to behave semantically, and so it is within a certain window 'up to' speakers how to behave and construe the truth-conditional content of certain expressions. Speakers must make some decision or other here, if they want to get on with it and actually use their language in an intelligible way, and some form of consistency is expected of them; but this only relatively loosely over time and with space for indecision left in. What choices speakers make can then indirectly reflect facts about them, signaling why they chose that way, and not another (and this choice may be purposeful or not).

Where subjectivity is lacking in virtue of the compositional semantics – which above has been represented by lack of model-sensitivity in the denotation of an expression – speakers' behavior is so uniform that deviation from it is not commonplace, and unlikely to be communicatively successful, or possibly even put up with. Using functional items in a way that does not encode their specific formal characteristics – say, trying to use *might* with universal modal force – or using lexical items with an application so far outside the vague norm as to be incomprehensible, might elicit anything from a correction to simple bewilderment. Part of the task of a future metasemantics is to understand how this strict regularity in speaker behavior comes about, and for the linguist interested in subjectivity, where and why it does not.

At some level, all expressions are underdetermined with respect to their truth-conditional usage, but this underdetermination isn't uniform. The formal properties of functional items

might be hardly susceptible to it, while certain contentious lexical items may be so underdetermined that they have no clear stable application in the linguistic community at all. The present chapter has shown how experiential semantics specifically drives a prominent form of strong linguistic subjectivity. Its conclusion is that experiential predicates occupy just one spot on a far wider continuum, but one that is especially salient due to how obviously differences in experience and direct evidence among a speaker population give rise to stable misalignments in construed truth conditions. But this variation is itself implicitly understood by speakers, and takes on a life of its own, which increases the expressive power of the language by allowing it to do more than invoke fixed truth conditions, but using nothing over and above its truth-conditional features.

REFERENCES

[1] Aikhenvald, A. Y. (2004). *Evidentiality*. Oxford: Oxford University Press (Oxford Linguistics).

[2] Akatsuka, N. (1976). Reflexivization: a transformational approach. In M. Shibatani, ed., *Syntax and Semantics 6: The Grammar of Causative Constructions*, 51-116. New York: Academic Press.

[3] Anagnostopoulou, E. (1999). On experiencers. In A. Alexiadou, G. C. Horrocks, & M. Stavrou, eds., *Studies in Greek Syntax*, 67-93. Dodrecht: Kluwer.

[4] Anand, P. & N. Korotkova (2018). Acquaintance content and obviation. In U. Sauerland & S. Solt, eds., *Proceedings of Sinn und Bedeutung* (ZASPiL 60) 22(1): 55-72. Berlin: ZAS.

[5] Anthony, A. (2016). Experience, evaluation and faultless disagreement. *Inquiry* 59(6): 686-722.

[6] Arad, M. (1998). VP-structure and the syntax-lexicon interface. Ph.D. thesis, University College London.

[7] Asudeh, A. & I. Toivonen (2012). Copy raising and perception. *Natural Language and Linguistic Theory* 30(2): 321-380.

[8] Bach, E. (1986). Natural language metaphysics. *Studies in Logic and the Foundations of Mathematics* 114: 573-595.

[9] Baker, C. (2012). Indexical contextualism and the challenges from disagreement. *Philosophical Studies* 157(1): 107-123.

[10] Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy* 25(1): 1-36.

[11] Barker, C. (2013). Negotiating taste. *Inquiry* 56(2-3): 240-257.

[12] Barnes, J. (1984). Evidentials in the Tuyaca verb. *International Journal of American Linguistics* 50(3): 255-271.

[13] Belletti, A. & L. Rizzi (1988). Psych verbs and θ-theory. *Natural Language and Linguistic Theory* 6(3): 291-352.

[14] Bhatt, R. & R. Pancheva (2006). Implicit arguments. In M. Everaert & H. van Riemsdijk, eds., *The Blackwell Companion to Syntax*, vol. II: 554-584. Oxford: Blackwell.

[15] Bittner, M. (2011). Time and modality without tenses or modals. In R. Musan & M. Rathert, eds., *Tense Across Languages*, 147-188. Berlin: de Gruyter.

[16] Bouchard, D-É. (2012). Long-distance degree quantification and the grammar of subjectivity. Ph.D. thesis, McGill University.

[17] Burgess, A. & B. Sherman (2014). Introduction: a plea for the metaphysics of meaning. In A. Burgess & B. Sherman, eds., *Metasemantics: New Essays on the Foundations of Meaning*, 1-16. Oxford: Oxford University Press.

[18] Bylinina, L. (2017). Judge-dependence in degree constructions. *Journal of Semantics* 34(2): 291-331.

[19] Caie, M. (2014). Metasemantics and metaphysical indeterminacy. In A. Burgess & B. Sherman, eds., *Metasemantics: New Essays on the Foundations of Meaning*, 55-96. Oxford: Oxford University Press.

[20] Cappelen, H. & J. Hawthorne (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press.

[21] Cappelen, H. & T. T. Huvenes (2018). Relative truth. In M. Glanzberg, ed., *Oxford Handbook of Truth*, 517-542. Oxford: Oxford University Press.

[22] Carlson, G. N. (1977). Reference to kinds in English. Ph.D. thesis, University of Massachusetts Amherst.

[23] Chang, B. S. & K. Chang (1984). The certainty hierarchy among spoken Tibetan verbs of being. *Bulletin of the Institute of History and Philology, Academica Sinica* 55(4): 603-634.

[24] Chierchia, G. (1984). Topics in the syntax and semantics of infinitives and gerunds. Ph.D. thesis, University of Massachusetts Amherst.

[25] Chierchia, G. (1989). Anaphora and attitudes "de se." In R. Bartsch, J. van Benthem, & P. van Emde Boas, eds., *Semantics and Contextual Expression*, 1-31. Dodrecht: Kluwer.

[26] Chierchia, G. (1995). Individual-level predicates as inherent generics. In G. N. Carlson & F. J. Pelletier, eds., *The Generic Book*, 176-223. Chicago: The University of Chicago Press.

[27] Clapp, L. (2015). A non-alethic approach to faultless disagreement. *Dialectica* 69(4): 517-550.

[28] Collins, J. (2013). The syntax of personal taste. *Philosophical Perspectives* 27(1): 51-103.

[29] Condoravdi, C. & J. M. Gawron (1996). The context-dependency of implicit arguments. In M. Kanazawa, C. Piñón, & H. de Swart, eds., *Quantifiers, Deduction, and Context*, 1-32. Stanford: CSLI Publications.

[30] Coppock, E. (2018). Outlook-based semantics. *Linguistics and Philosophy* 41(2): 125-164.

[31] Cova, F. & N. Pain (2012). Can folk aesthetics ground aesthetic realism? *Monist* 95(2): 241-263.

[32] Cresswell, M. J. (1975). Hyperintensional logic. *Studia Logica* 34(1): 25-38.

[33] DeLancey, S. (2017). Evidentiality in Tibetic. In A. Y. Aikkhenvald, ed., *The Oxford Handbook of Evidentiality*, 580-594. Oxford: Oxford University Press (Oxford Handbooks in Linguistics).

[34] Denwood, P. (1999). *Tibetan*. Amsterdam: John Benjamins (London Oriental and African Language Library 3).

[35] Dowty, D. R. (1979). *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and Montague's PTQ*. Dodrecht: Kluwer.

[36] Egan, A. (2010). Disputing about taste. In R. Feldman & T. A. Warfield, eds., *Disagreement*, 247-286. Oxford: Oxford University Press.

[37] Egan, A., J. Hawthorne, & B. Weatherson (2005). Epistemic modals in context. In G. Preyer & G, Peter, eds., *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 131-170. Oxford: Oxford University Press.

[38] Epstein, S. D. (1984). Quantifier-pro and the LF representation of PRO*ARB*. *Linguistic Inquiry* 15(3): 499-505.

[39] Faller, M. T. (2002). Semantics and pragmatics of evidentials in Cuzco Quechua. Ph.D. thesis, Stanford University.

[40] Fillmore, C. J. (1986). Pragmatically controlled zero anaphora. In V. Nikiforidou, M. VanClay, M. Niepokuj, & D. Feder, eds., *Proceedings of the Twelfth Annual Meeting of the Berkeley Linguistics Society*: 95-107. Berkeley: Berkeley Linguistics Society.

[41] von Fintel, K. & A. Gilles (2010). *Must*...stay...strong! *Natural Language Semantics* 18(4): 351-383.

[42] Fleisher, N. (2013). The dynamics of subjectivity. In T. Snider, ed., *Proceedings of SALT* 23: 276-294.

[43] Francez, I. & A. Koontz-Garboden (2017). *Semantics and Morphosyntactic Variation: Qualities and the Grammar of Property Concepts*. Oxford: Oxford University Press (Oxford Studies in Theoretical Linguistics).

[44] Garrett, E. J. (2001). Evidentiality and assertion in Tibetan. Ph.D. thesis, University of California Los Angeles.

[45] Gawne, L. (2017). Egophoric evidentiality in Bodish languages. In L. Gawne & N. Hill, eds., *Evidential Systems of Tibetan Languages*, 61-94. Berlin: de Gruyter (Trends in Linguistics Studies and Monographs, vol. 302).

[46] Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.

[47] Giannakidou, A. & A. Mari (2016). Epistemic future and epistemic MUST: nonveridicality, evidence, and partial knowledge. In Blaszack, J. A., A. Giannakidou, D. Klimek-Jankowska, & K. Mygdalski, eds., *Mood, Aspect, Modality Revisited: New Answers to Old Questions*, 75-117. Chicago: The University of Chicago Press.

[48] Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies* 136(1): 1-29.

[49] Glanzberg, M. (2014). Explanation and partiality in semantic theory. In A. Burgess & B. Sherman, eds., *Metasemantics: New Essays on the Foundations of Meaning*, 259-292. Oxford: Oxford University Press.

[50] Goodhue, D. (2017). *Must* ϕ is felicitous only if ϕ is not known. *Semantics and Pragmatics* 10, Article 14.

[51] Gutzmann, D. (2016). If expressivism is fun, go for it! In C. Meier & J. van Wijnbergen-Huitnik, eds., *Subjective Meaning: Alternatives to Relativism*, 21-46. Berlin: de Gruyter (Linguistische Arbeiten 559).

[52] de Haan, F. (1999). Evidentiality and epistemic modality: setting boundaries. *Journal of Southwest Linguistics* 18: 83-102.

[53] Heim, I. & A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell (Blackwell Textbooks in Linguistics).

[54] Heycock, C. (1994). *Layers of Predication: The non-lexical syntax of clauses*. New York: Garland.

[55] Hill, N. (2012). "Mirativity" does not exist: *ḥdug* in "Lhasa" Tibetan and other suspects. *Linguistic Typology* 16(3): 389-433.

[56] Hill, N. & L. Gawne (2017). The contribution of Tibetan languages to the study of evidentiality. In L. Gawne & N. Hill, eds., *Evidential Systems of Tibetan Languages*, 1-38. Berlin: de Gruyter (Trends in Linguistics Studies and Monographs, vol. 302).

[57] Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Two Notions*. Ithaca: Cornell University Press.

[58] Hintikka, J. (1969). Semantics for propositional attitudes. In *Models for Modality: Selected Essays*, 87-11. Dodrecht: D. Reidel (Synthese Library).

[59] Hirvonen, S. (2014). Predicates of personal taste and perspective dependence. Ph.D. thesis, University College London.

[60] Hirvonen, S. (2016). Doing without judge dependence. In C. Meier & J. van Wijnbergen-Huitnik, eds., *Subjective Meaning: Alternatives to Relativism*, 47-68. Berlin: de Gruyter (Linguistische Arbeiten 559).

[61] Huvenes. T. T. (2012). Varieties of disagreement and predicates of taste. *Australasian Journal of Philosophy* 90(1): 167-181.

[62] Huvenes, T. T. (2014). Disagreement without error. *Erkenntnis* 79(1S): 143-154.

[63] Iacona, A. (2008). Faultless or disagreement. In M. García-Carpintero & M. Kölbel, eds., *Relative Truth*, 287-295. Oxford: Oxford University Press.

[64] Izvorzki, R. (1997). The present perfect as an epistemic modal. In A. Lawson, ed., *Proceedings of SALT* 9: 222-239.

[65] Jackendoff, R. *Semantic Structures*. Cambridge, MA: The MIT Press (Current Studies in Linguistics, vol. 18).

[66] Katzir, R. & R. Singh (2013). Hurford disjunctions: embedded exhaustification and structural economy. In U. Etxeberria, A. Fălăuş, A. Irurtzun, & B. Leferman, eds. *Proceedings of Sinn und Bedeutung* 18: 201-216.

[67] Kennedy, C. (1997). Projecting the adjective: the syntax and semantics of gradability and comparison. Ph.D. thesis, University of California Santa Cruz.

[68] Kennedy, C. (2013). Two sources of subjectivity: qualitative assessment and dimensional uncertainty. *Inquiry* 56(2-3): 258-277.

[69] Kennedy, C. (2016). Two kinds of subjectivity. In C. Meier & J. van Wijnbergen-Huitnik, eds., *Subjective Meaning: Alternatives to Relativism*, 105-126. Berlin: de Gruyter (Linguistische Arbeiten 559).

[70] Kennedy, C. & L. McNally (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81(2): 345-381.

[71] Kennedy, C. & M. Willer (2016). Subjective attitudes and counterstance contingency. In M. Moroney, C.-R. Little, J. Collard, & D. Burgdorf, eds., *Proceedings of SALT* 26: 913-933.

[72] Klecha, P. (2014). Diagnosing modality in predictive expressions. *Journal of Semantics* 31(3): 443-455.

[73] Kneer, M. (m.s.). Predicates of personal taste: empirical data. M.s., University of Zürich.

[74] Kneer, M., A. Vicente, & D. Zeman (2017). Relativism about predicates of personal taste and perspectival plurality. *Linguistics and Philosophy* 40(1): 37-60.

[75] Kölbel, M. (2002). *Truth Without Objectivity*. London: Routledge (International Library of Philosophy).

[76] Kölbel, M. (2004). Faultless disagreement. In *Proceedings of the Aristotelian Society* 104(1): 53-73.

[77] Kölbel, M. (2009). The evidence for relativism. *Synthese* 166(2): 375-395.

[78] Kratzer, A. (1986). Conditionals. In *Proceedings from the Annual Meeting of the Chicago Linguistic Society* 22(2): 1-15.

[79] Krifka, M., F. J. Pelletier, G. N. Carlson, A. ter Meulen, G. Chierchia, & G. Link (1995). Genericity: an introduction. In G. N. Carlson & F. J. Pelletier, eds., *The Generic Book*, 1-124. Chicago: The University of Chicago Press.

[80] Kuno, S. (1973). *The Structure of the Japanese Language*. Cambridge, MA: The MIT Press (Current Studies in Linguistics 3).

[81] Lakoff, G. (1970). *Irregularity in Syntax*. New York: Holt, Rinehart & Winston (Transatlantic Series in Linguistics).

[82] Landau, I. (2010a). *The Locative Syntax of Experiencers*. Cambridge, MA: The MIT Press (Linguistic Inquiry Monograph 53).

[83] Landau, I. (2010b). The explicit syntax of implicit arguments. *Linguistic Inquiry* 41(3): 357-388.

[84] Landau, I. (2011). Predication vs. aboutness in copy raising. *Natural Language and Linguistic Theory* 29(3): 779-813.

[85] Lasersohn, P. (1995). Sounds like *like*. *Linguistic Analysis* 25(1-2): 70-77.

[86] Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28(6): 643-686.

[87] Lasersohn, P. (2008). Quantification and perspective in relativist semantics. *Philosophical Perspectives* 22(1): 305-337.

[88] Lasersohn, P. (2009). Relative truth, speaker commitment, and control of implicit arguments. *Synthese* 166(2): 359-374.

[89] Lasersohn, P. (2013). Non-world indices and assessment-sensitivity. *Inquiry* 56(2-3): 122-148.

[90] Lasersohn, P. (2017). *Subjectivity and Perspective in Truth-Theoretic Semantics*. Oxford: Oxford University Press (Oxford Linguistics: Oxford Studies in Semantics and Pragmatics 8).

[91] Lee, C. (2013). Psych-predicates: 1st person and evidentiality. *Procedia – Social and Behavioral Sciences* 97: 414-421.

[92] Lepore, E. (1983). What model theoretic semantics cannot do? *Synthese* 54(2): 167-187.

[93] Lewis, D. (1969). *Convention: A Philosophical Study*. Oxford: Blackwell.

[94] Lewis, D. (1979). Attitudes *de dicto* and *de se*. *The Philosophical Review* 88(4): 513-543.

[95] Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Öhman, eds., *Philosophy and Grammar*, 79-100. Dodrecht: Reidel.

[96] Lim, D. & C. Lee (2012). Perspective shifts of Korean evidentials and the effect of contexts. In A. Chereches, ed., *Proceedings of SALT* 22: 26-42.

[97] Link, G. (2002). The logical analysis of plurals and mass terms: a lattice-theoretic approach. In P. Portner & B. H. Partee, eds., *Formal Semantics: The Essential Readings*, 127-146. Oxford: Blackwell. Originally (1983) in R. Bäuerle, C. Schwarze, & A. von Stechow, eds., *Meaning, Use and Interpretation of Language*, 303-329. Berlin: de Gruyter.

[98] López de Sa, D. (2008). Presuppositions of commonality: an indexical relativist account of disagreement. In M. García-Carpintero & M. Kölbel, eds., *Relative Truth*, 297-310. Oxford: Oxford University Press.

[99] Ludlow, P. (2008). Cheap contextualism. *Philosophical Issues* 18(1): 104-129.

[100] MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies* 132(1): 17-31.

[101] MacFarlane, J. (2009). Nonindexical contextualism. *Synthese* 166(2): 231-250.

[102] MacFarlane, J. (2011a). Simplicity made difficult. *Philosophical Studies* 156(3): 441-448.

[103] MacFarlane, J. (2011b). Epistemic modals are assessment-sensitive. In B. Weatherson & A. Egan, eds., *Epistemic Modality*, 108-130. Oxford: Oxford University Press.

[104] MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press (Context and Content).

[105] Maslova, E. (2003). Evidentiality in Yukaghir. In R. M. W. Dixon & A. Y. Aikhenvald, eds., *Word: A Cross-Linguistic Typology*, 219-236. Cambridge: Cambridge University Press.

[106] Matthewson, L. (m.s.). Evidence type, evidence location, evidence strength. M.s., University of British Columbia.

[107] Matthewson, L., H. Davis, & H. Rullmann (2007). Evidentials as epistemic modals: evidence from St'át'imcets. *Linguistic Variation Yearbook* 7: 203-256.

[108] McCloskey, J. & P. Sells (1988). Control and A-chains in modern Irish. *Natural Language and Linguistic Theory* 6(2): 143-189.

[109] McCready, E. (2007). Context shifting in questions and elsewhere. In E. Puig-Waldmuller, ed., *Proceedings of Sinn und Bedeutung* 11:433-447.

[110] McLendon, S. (2003). Evidentials in Eastern Pomo with a comparative survey of the category in other Pomoan languages. In R. M. W. Dixon & A. Y. Aikhenvald, eds., *Word: A Cross-Linguistic Typology*, 101-130. Cambridge: Cambridge University Press.

[111] McNally, L. & I. Stojanovic (2017). Aesthetic adjectives. In J. O. Young, ed., *Semantics of Aesthetic Judgments*, 17-37. Oxford: Oxford University Press.

[112] Milsark, G. (1977). Toward an explanation of certain peculiarities of the existential construction in English. *Linguistic Analysis* 3: 1-30.

[113] Moltmann, F. (2010). Relative truth and the first person. *Philosophical Studies* 150(2): 187-220.

[114] Montague, R. (1970). English as a formal language. In B. Visentini, ed., *Linguaggi Nella Società e Nella Tecnica*, 189-224. Milan: Edizioni di Communità.

[115] Montague, R. (1973). The proper treatment of quantification in ordinary English. In P. Suppes, J. Moravcsik, & J. Hintikka, eds., *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, 221-242. Dodrecht: D. Reidel (Synthese Library 49).

[116] Murray, S. E. (2017). *The Semantics of Evidentials*. Oxford: Oxford University Press (Oxford Studies in Semantics and Pragmatics 9).

[117] Ninan, D. (2014). Taste predicates and the acquaintance inference. In T. Snider, S. D'Antonio, & M. Weigand, eds., *Proceedings of SALT* 24: 290-309.

[118] Oswalt, R. L. (1986). The evidential system of Kashaya. In W. L. Chafe & J. Nichols, eds., *Evidentiality: The Linguistic Encoding of Epistemology*, 29-45. Norwood: Ablex.

[119] Partee, B. H. (2004). Binding implicit variables in quantified contexts. In *Compositionality in Formal Semantics: Selected Papers by Barbara H. Partee*, 259-281. Oxford: Blackwell (Explorations in Semantics). Originally (1989) in C. Wiltshire, R. Graczyk, & B. Music, eds., *Proceedings from the Annual Meeting of the Chicago Linguistic Society* 25: 342-365.

[120] Partee, B. H. & M. Rooth (1983). Generalized conjunction and type ambiguity. In R. Bäuerle, C. Schwarze, & A. von Stechow, eds., *Meaning, Use, and Interpretation of Language*, 361-383. Berlin: de Gruyter.

[121] Pearson, H. (2013). A judge-free semantics for predicates of personal taste. *Journal of Semantics* 30(1): 103-154.

[122] Pérez Carballo, A. (2014). Semantic hermeneutics. In A. Burgess & B. Sherman, eds., *Metasemantics: New Essays on the Foundations of Meaning*, 119-146. Oxford: Oxford University Press.

[123] Pesetsky, D. (1995). *Zero Syntax: Experiencers and Cascades*. Cambridge, MA: The MIT Press (Current Studies in Linguistics, vol. 27).

[124] Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford: Oxford University Press (Oxford Studies in Linguistics).

[125] Rákosi, G. (2006). Dative experiencer predicates in Hungarian. Ph.D. thesis, University of Utrecht.

[126] Richard, M. (2004). Contextualism and relativism. *Philosophical Studies* 119(1-2): 215-242.

[127] Rogers, A. (1972). Another look at flip perception verbs. In P. M. Peranteau, J. N. Levi, & G. C. Phares, eds., *Papers from the 8th Regional Meeting of the Chicago Linguistic Society*, 303-315.

[128] Rudolph, R. (forthcoming). Talking about appearances: the roles of evaluation and experience in disagreement. *Philosophical Studies*.

[129] Sæbø, K. J. (2009). Judgment ascriptions. *Linguistics and Philosophy* 32(4): 327-352.

[130] Sassoon, G. (2010). Restricted quantification over tastes. In M. Aloni, H. Bastiaanse, T. de Jager, & K. Schulz, eds., *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18 2009, Revised Selected Papers*, 163-172.

[131] Sassoon, G. (2013). A typology of multidimensional adjectives. *Journal of Semantics* 30(3): 335-380.

[132] Schafer, K. (2011). Faultless disagreement and aesthetic realism. *Philosophy and Phenomenological Research* 82(2): 265-286.

[133] Schaffer, J. (2011). Perspective in taste predicates and epistemic modals. In A. Egan & B. Weatherson, eds., *Epistemic Modality*, 179-226. Oxford: Oxford University Press.

[134] Snyder, E. (2013). Binding, genericity, and predicates of personal taste. *Inquiry* 56(2-3): 278-306.

[135] Stalnaker, R. (1978). Assertion. *Syntax and Semantics* 9: 315-332.

[136] Stephenson, T. (2007a). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30(4): 487-525.

[137] Stephenson, T. (2007b). Towards a theory of subjective meaning. Ph.D. thesis, Massachusetts Institute of Technology.

[138] Stephenson, T. (2010). Control in centered worlds. *Journal of Semantics* 27(4): 409-436.

[139] Stojanovic, I. (2007). Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy* 30(6): 691-706.

[140] Sun, J. T.-S. (2017). Evidentials and person. In A. Y. Aikhenvald, ed., *The Oxford Handbook of Evidentiality*, 47-63. Oxford: Oxford University Press (Oxford Handbooks in Linguistics).

[141] Sundell, T. (2011). Disagreements about taste. *Philosophical Studies* 155(2): 267-288.

[142] Tournadre, N. & S. Dorje (2003). *Manual of Standard Tibetan: Language and Civilization*. Ithaca: Snow Lion.

[143] Tournadre, N. (2017). A typological sketch of evidential/epistemic categories in the Tibetic languages. In L. Gawne & N. Hill, eds., *Evidential Systems of Tibetan Languages*, 95-129. Berlin: de Gruyter (Trends in Linguistics Studies and Monographs, vol. 302).

[144] Vardomskaya, T. (2018). Sources of subjectivity. Ph.D. thesis, University of Chicago.

[145] Willett, T. (1988). A cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language* 12(1): 51-97.

[146] Williamson, T. (1994). *Vagueness*. London: Routledge (The Problems of Philosophy).

[147] Wolf, L. (2016). Predicates of personal taste and the evidential step. In C. Meier & J. van Wijnbergen-Huitnik, eds., *Subjective Meaning: Alternatives to Relativism*, 69-90. Berlin: de Gruyter (Linguistische Arbeiten 559).

[148] Wyatt, J. (2018). Absolutely tasty: an examination of predicates of personal taste and faultless disagreement. *Inquiry* 61(3): 252-280.

[149] Yoon, Y. (1996). Total and partial predicates and the weak and strong interpretations. *Natural Language Semantics* 4(3): 217-236.

[150] Zeman, D. (2015). Relativism and bound predicates of personal taste: an answer to Schaffer's argument from binding. *Dialectica* 69(2): 155-183.

[151] Zeman, D. (2016). Contextualism and disagreement about taste. In C. Meier & J. van Wijnbergen-Huitnik, eds., *Subjective Meaning: Alternatives to Relativism*, 91-104. Berlin: de Gruyter (Linguistische Arbeiten 559).