

Beauty, Art and Testimony:
Subjectivity and Objectivity in Aesthetics

by

Erica Diane Klempler

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Hannah Ginsborg, Co-chair

Professor Barry Stroud, Co-chair

Professor John MacFarlane

Professor Line Mikkelsen

Spring 2018

Copyright © 2018 by Erica Diane Klempner

All rights reserved

Abstract

Beauty, Art and Testimony: Subjectivity and Objectivity in Aesthetics

by

Erica Diane Klemptner

Doctor of Philosophy in Philosophy

University of California, Berkeley

Professor Hannah Ginsborg and Professor Barry Stroud, Co-chairs

We acquire beliefs on the basis of what others tell us all the time. If you tell me that your house is painted red, chances are that I will simply believe you without question; and if someone asks me what color your house is, I will simply tell her that it is red. Yet we do not seem to accept others' testimony about beauty and art in the same way. If you tell me that the Taj Mahal is beautiful, or that *Middlemarch* is a great novel, it would be strange for me simply to adopt your view, even if I have a lot of confidence in your judgment. In order for me to be in a position to believe or to claim that the Taj Mahal is beautiful, or that *Middlemarch* is a great novel, I must experience these things myself—by going to Agra, or by reading the book.

My dissertation explains this apparent resistance to aesthetic testimony. I argue that aesthetic claims carry what is known in linguistics as 'evidential' information, to the effect that they are made on the basis of the subject's firsthand experience. I locate aesthetic claims amongst a slew of other kinds of claims—including personal taste claims, perceptual appearance claims and moral claims—that all carry evidential information of this kind. These claims are loosely connected by having content that is essentially subjective or perspectival in some way, even when, as in the aesthetic and moral cases, there is also some claim to objectivity in play. I offer a detailed explanation in the aesthetic case—specifically for claims about beauty—of how their content comprises both subjective and objective elements. I argue that the evidential information carried by aesthetic claims and the other 'perspectival' claims I mention is communicated as conversational implicature, in spite of initial appearances to the contrary. Our apparent resistance to aesthetic testimony thus turns out to be an artifact of the evidential implications of aesthetic language: we often do accept aesthetic testimony, but we must describe our beliefs in a way that does not convey misleading evidential information.

To my family, to PG Tips, and to Stanley Chen,
with whom I began this conversation a long time ago.

TABLE OF CONTENTS

| | | |
|----------|---|-----------|
| 1 | Introduction to the problem | 1 |
| 2 | Previous attempts to solve the problem | 4 |
| 2.1 | Emotivism (and other expressivist views) | 4 |
| 2.2 | Subjectivism | 5 |
| 2.3 | An epistemic problem? | 11 |
| 2.3.1 | Lack of ‘full understanding’ | 11 |
| 2.3.2 | Is aesthetic testimony unreliable? | 15 |
| 2.4 | Other solutions | 19 |
| 2.4.1 | Aaron Meskin’s ‘reverse error theory’ | 19 |
| 2.4.2 | Robert Hopkins and ‘Unusability’ | 23 |
| 3 | Evidentiality | 28 |
| 3.1 | Conveying evidentiality | 28 |
| 3.2 | <i>Must</i> | 29 |
| 3.3 | <i>Must</i> , aesthetic claims and evidentiality | 39 |
| 4 | Multiple dimensions of content | 46 |
| 4.1 | Conversational implicature | 46 |
| 4.2 | Presupposition and conventional implicature | 51 |
| 4.3 | The ‘Direct Evidentials’ and conversational implicature | 55 |
| 5 | Content and calculability: some non-aesthetic examples | 57 |
| 5.1 | Presumption of evidential source | 57 |
| 5.2 | The role of perspective: some examples with ‘looks’ (and one or two other things) | 60 |
| 5.3 | The role of perspective: personal taste | 68 |
| 6 | The content of beauty claims: first pass | 71 |
| 6.1 | The normativity of beauty | 72 |
| 6.2 | An objectivist view for beauty? | 77 |
| 6.3 | Critical standards | 79 |
| 6.4 | A problem with the objectivist view | 81 |
| 7 | The content of beauty claims: a hybrid model | 86 |
| 7.1 | Judgments based on pleasure: tastiness | 88 |
| 7.2 | Judgments based on pleasure: beauty | 90 |
| 7.3 | The structure of beauty claims | 94 |

| | | |
|----------|--|-----|
| 7.4 | Negative judgments of beauty | 105 |
| 7.5 | Disagreement about beauty | 107 |
| 7.6 | Calculability for aesthetic claims | 111 |
| 7.7 | Wrapping up aesthetic content and calculability | 116 |
| 8 | Cancelability | 118 |
| 8.1 | Extremely robust conversational implicatures | 119 |
| 8.2 | Aesthetic and other Direct Evidential implicatures | 124 |
| 8.3 | Aesthetic implicature and aesthetic belief | 128 |
| | References | 130 |

1. INTRODUCTION TO THE PROBLEM

Many people think that there is something wrong with acquiring aesthetic beliefs on the basis of the testimony of others. If a friend tells me that *Middlemarch* is a great novel, but I have never read it, it would be strange for me simply to adopt my friend's view, even if I have a lot of confidence in her judgment. In order for me to be in a position to believe that *Middlemarch* is a great novel, I must read it myself. And, it's thought, because I am not in a position to believe that *Middlemarch* is a great novel if I haven't read it myself, I am also not in a position to claim that *Middlemarch* is a great novel. The most I will say on the basis of testimony is that I've heard *Middlemarch* is a great novel, or that it's supposed to be a great novel, or that my friend Margot says it's a great novel. I will not simply say, "*Middlemarch* is a great novel."

Kant holds this view for judgments of beauty; he maintains, for instance, "that the approval of others provides no valid proof for the judging of beauty, that others may perhaps see and observe for [someone], and that what many have seen in one way what he believes himself to have seen otherwise, may serve him as a sufficient ground of proof for a theoretical, hence a logical judgment, but that what has pleased others can never serve as the ground of an aesthetic judgment." (Kant 2000: §33, 5:284.)¹ As Kant here notes, the phenomenon at issue does not extend to many other kinds of beliefs. If there is something wrong with acquiring beliefs about beauty, and other aesthetic beliefs, on the basis of testimony, this makes aesthetic beliefs very different from other kinds of beliefs, even other kinds of perceptual beliefs, as we often acquire such beliefs on the basis of testimony. If Margot tells me that her house is red, I may easily believe that her house is red. As Kant also suggests here, I may do so even if I have previously seen her house myself and on that basis thought it was orange (possibly because I saw it under yellow streetlights). Beliefs based on testimony often, in fact, amount to knowledge. I come to know that Margot lives in a red house when she tells me so. Moreover, I can on this basis tell someone else that Margot's house is red. I would not have to qualify my claim by saying, "Apparently, Margot's house is red," or "Margot told me that her house is red."

The comparison with color seems apt because both color beliefs and aesthetic beliefs are acquired, in the first instance, on the basis of a kind of subjective response. For aesthetic beliefs, this is most obvious in the case of judgments of beauty (Kant's focus), which are based on a particular kind of pleasure felt in the experience of the beautiful thing. However, the connection with subjective response is arguably there with other kinds of aesthetic belief as well—e.g. that a novel is good or a vase graceful—and in any case whatever issue there is concerning aesthetic testimony has generally been taken to apply to any kind of aesthetic belief. The comparison with color raises the question of why the aesthetic domain is different.

Note that the issue does not seem to be that, for some reason, unqualified aesthetic beliefs acquired on the basis of testimony never amount to knowledge. The issue, if real, is prior to this: it is that there is something strange about acquiring such beliefs at all. I do think the issue is real, in that there is a genuine strangeness here that needs an explanation. We do not think of ourselves as having unqualified testimony-based aesthetic beliefs, and we do not make sincere, unqualified aesthetic claims on the basis of testimony. This is presumably not just an accident: there is presumably some norm that governs this aspect of our aesthetic practice. If the norm is epistemic—if, that is, testimony is for some reason an insufficient epistemic ground for aesthetic beliefs, barring testimony-based aesthetic beliefs from fulfilling the conditions necessary for

¹ Henceforth, Kant's *Critique of Judgment* will be referred to as *CJ*.

knowledge—then this norm has been embodied by that practice, to the extent that we do not acquire unqualified aesthetic beliefs on the basis of testimony at all.

In this work I will investigate the problem of aesthetic testimony: why, that is, we do not acquire unqualified aesthetic beliefs on the basis of testimony, or make unqualified testimony-based aesthetic claims. The question has to date received many kinds of answer. The phenomenon is important to Kant, as I've indicated. There has not been agreement even on what kind of explanation we are looking for: an epistemological explanation, or some other kind. I will do a survey of the kinds of answer that have been offered, or are at least possible, in order to try to locate where the answer lies.

It has been noticed that there seems to be a problem with testimony in other areas of discourse as well—moral discourse in particular, although I believe the issue extends to evaluative discourse generally and to some other areas that are not obviously evaluative at all. There are of course significant differences between the aesthetic domain and these other areas—for instance, a moral claim is not a perceptual claim—but there are interesting parallels to be drawn between these areas, and I will eventually bring them into the discussion.

The problem with aesthetic testimony is generated by this supposition about aesthetic content:

Experience Independence

When someone holds an unqualified aesthetic belief or makes an unqualified aesthetic assertion about something, the belief or assertion has propositional content that is independent of any particular aesthetic response that that person has had to that thing.

An analogous claim about color would surely be correct, and *Experience Independence* is itself very plausible: when I believe that something is beautiful, it does not seem to me that its being beautiful is dependent on my having experienced it. But if this is the case, it is unclear why a speaker must herself have had an aesthetic response to something in order to have an unqualified aesthetic belief about it, or to make an unqualified claim about it. “Margot’s house is red” has a content that is independent of my seeing her house to be red, and I can legitimately believe it, or assert it, without having seen her house myself; but it is apparently different for “Margot’s house is beautiful.”

There are of course significant differences between ‘beautiful’ and ‘red.’ One difference is that ‘beautiful’ is an evaluative term and ‘red’ is not. ‘Beautiful’ is also, arguably, a *normative* evaluative term. Kant has certainly argued this: according to him, my claim that X is beautiful does not reflect the kind of (merely) causal relationship between objects and human response that is arguably behind the use of both color predicates and predicates of ‘personal taste’ (what Kant calls ‘the agreeable’). My claim is not, for instance, that X reliably causes pleasure in the majority of human subjects, or that X reliably causes pleasure in *me*; it instead seems to reflect a normative relationship: in calling X beautiful, I apparently indicate that X *merits* the pleasure—that such pleasure is appropriate to X. But there is nothing about a normative claim that makes it obviously unsuitable for testimony: if I claim that X merits a certain kind of pleasure, why can’t you know this on the basis of my say-so, and make the claim yourself accordingly? In fact, beauty’s normative element, assuming it has one, seems to render its problem with testimony

more puzzling than if it were merely a matter of personal taste, as normative claims are in some sense valid for everyone.²

The general challenge of understanding normative discourse—most visibly in ethics—has led to expressivist theories that in any case deny suppositions like *Experience Independence*. Such theories deny that what look like assertions in the normative domain at issue have *propositional* content at all (at least, at bottom): these utterances do not in the first instance express beliefs, but rather certain noncognitive states that the speaker holds. We could cook up an aesthetic version of expressivism to account for the phenomenon of aesthetic testimony by saying that what appear to be simple, unqualified aesthetic assertions are in fact nonpropositional expressions of the speaker's aesthetic experience. I will look very briefly at the prospects of such a solution, before turning to discussion of a form of subjectivism that also denies *Experience Independence*: in this case not by denying that simple aesthetic assertions have propositional content, but by insisting that their content is in fact dependent on the speaker's aesthetic experience. This will give us a sense of the difficulties involved in trying to deny *Experience Independence*, before we turn to other ways to solve the problem with aesthetic testimony that accept *Experience Independence*.

I will eventually conclude that the norms responsible for the phenomena at issue are *linguistic* norms governing the conveyance by unqualified aesthetic claims of evidential information. That is, while the primary content of a typical unqualified aesthetic claim ascribes an aesthetic property to an object, where this primary content is in accordance with *Experience Independence*, such a claim also carries evidential information to the effect that the speaker bases her claim on direct grounds, which in the aesthetic case is personal experience of the object in question. The fact that unqualified aesthetic claims carry this direct evidential signal is in turn responsible for our unwillingness to use the indicative sentences involved in making such claims as the complements of *that*-clauses in belief ascriptions, when the believer does not have personal experience of the object in question. The appearance that we never accept aesthetic testimony is thus an artifact of the evidential implications of aesthetic language: if I am told by an excellent judge that *Middlemarch* is a great novel, but I have never read it, I cannot felicitously describe myself as believing that *Middlemarch* is a great novel; however, this does not mean that I do not believe what I have been told. I must simply describe the aesthetic belief I have acquired differently.

² I introduce the term 'normative' here and will continue to use it throughout, as it conveniently indicates the connection with merit and appreciative response to merit that I will be discussing: it is natural to express this aspect of judgments of beauty by saying that everyone ought to feel pleasure in a beautiful thing. The nature of this 'ought', though, is unclear. I will return to this issue later.

2. PREVIOUS ATTEMPTS TO SOLVE THE PROBLEM

2.1 Emotivism (and other expressivist views)

Emotivism is the simplest form of expressivism. An emotivist view of aesthetic utterances, if we were to suppose one, would deny, as I said, that such utterances have propositional content. The fact that we do not make unqualified aesthetic claims on the basis of testimony would reflect a difficulty with the corresponding beliefs, but this would be because, strictly speaking, there would be no such thing as aesthetic beliefs (and hence no real aesthetic ‘claims’ either). We therefore wouldn’t acquire aesthetic beliefs on the basis of testimony—or indeed in any other way. For S to say (e.g.) that X is beautiful would not be for S to make a claim at all; S’s utterance wouldn’t express a belief, but rather the pleasure that X gives her (in something like the way saying “wow!” does). As “X is beautiful” would serve to express the speaker’s pleasure in X, the utterance would clearly be inappropriate unless the speaker herself experienced X and took pleasure in it. It could not be appropriately uttered, then, by someone who has merely been told that X is beautiful.

There are well-known insuperable problems with this simple emotivist view. It does a good job of accounting for the connection between a normative claim and feeling: in this case, between the beauty claim and the feeling of pleasure, without which it seems the beauty claim cannot legitimately be made. But it cannot account for the features of truth-conditional discourse—the sense that a claim ascribes a property to an object, the use of a term in non-declarative contexts (“Is X beautiful?”) or declarative but embedded contexts (e.g. “If X is beautiful, I will buy it”), the possibility of substantial disagreement, etc.—which the aesthetic domain seems to have, in common with uncontroversially objective areas of discourse.

Contemporary expressivists like Simon Blackburn and Alan Gibbard have moved far from the simple emotivist idea by offering sophisticated accounts that attempt to build up that idea—that basic evaluative sentences express non-belief states like feelings and attitudes, rather than beliefs—into a theory that has, or least mimics, all the trappings of assertoric discourse. How far their sophisticated kinds of expressivism succeed in general, or could succeed in the aesthetic case, I will not undertake to argue here. Critics of such accounts often claim either that their attempts do not work—that an evaluative claim is just that, a *claim*, which cannot be reduced to or otherwise explained as generated from the expression of a noncognitive state of mind—or that, if they do work, this is because they have taken on so much of the machinery of the cognitivist accounts which they initially set out to oppose, that they ultimately fail to distinguish themselves from them.

In any case, contemporary expressivism, unlike the simpler emotivism, does not provide a swift and easy answer to the problem of testimony. Emotivism very clearly solves the problem by claiming that evaluative utterances have purely nonpropositional, expressive content that serves to express the speaker’s feelings or attitudes. Sophisticated expressivism must complicate the relation between the content of the utterance and the speaker’s noncognitive states of mind in order to imbue evaluative utterances with assertoric character. Whether or not any particular theory of this kind can do a good job of accounting for the phenomenon of testimony in the aesthetic or other normative areas will depend on how the theory treats this relation in its details: on what the content of the utterance ultimately turns out to be. I suspect (although I will not argue at length here) that such theories have no particular advantage over others in dealing with the phenomenon of aesthetic testimony; in fact, Gibbard himself has already been cited by others

as saying that another's trustworthy testimony as to the credibility of some norm (understood in Gibbard's expressivist way) must favor a person's own acceptance of it.³ A sophisticated expressivism such as Gibbard's, adapted for the aesthetic domain, would presumably trade not in aesthetic *experiences* but in the aesthetic equivalent of moral 'norm acceptance': i.e. the acceptance of certain aesthetic standards. One may of course have a certain standard without ever having had the opportunity to exercise it, so this kind of expressivism would have the same problem relating aesthetic content to aesthetic experience as any other view.⁴ I will not pursue this avenue further.⁵

2.2 Subjectivism

We might allow that aesthetic beliefs and assertions have propositional content but still deny *Experience Independence* by supposing that their content is dependent on the speaker's response. This would give us some kind of subjectivist theory.

Note, though, that it would have to be a subjectivist theory of an unusual kind. Subjectivism is a form of contextualism, and there are plenty of contemporary contextualist theories of, say, personal taste—a domain of discourse which is at the least closely related to the aesthetic domain. In its simplest form, such a theory would say that for S to say that X is (e.g.) tasty is for S to claim that X tastes good to her. Similarly, we could say that for S to say that X is (e.g.) beautiful would be for S to claim that X gives her a certain kind of pleasure. But, put like this, the theory does not actually solve the problem with aesthetic testimony. So put, S's claim is that X is *such as* to give her pleasure; it is *disposed to* give her pleasure—it is not (for anything claimed by S) something that *is* giving her pleasure or *has* given her pleasure. And I can often know that something is the kind of thing that gives me pleasure without having been pleased by it (yet), if I have been told that someone else has been pleased by it and I know that we share similar tastes.

Contemporary semantic theories for personal taste generally do not focus on the testimony issue, even though it is a problem for personal taste claims just as much as for aesthetic claims.

³ E.g. Aaron Meskin (2004: 71, fn. 21). The relevant passage is in Gibbard (1990: 180-181): "When conditions are right and someone else finds a norm independently credible, I must take that as favoring my own accepting the norm... . In short, suppose under good conditions for judgment that others find a norm independently credible. Then that must favor the norm in my own eyes. I must accept norms that say to treat that fact as weighing with me in favor of accepting the norm." In Gibbard (2003), he talks not of norm acceptance but of one's commitment to a certain kind of contingency plan; however, I see no reason why this shift would change his attitude to the rationality of giving weight to others' views.

⁴ It has been claimed, for instance, that Gibbard-style expressivism is a very close cousin of certain relativist views. (E.g. in John MacFarlane 2014: §7.3.) I will discuss the relation between relativism and the problem of aesthetic testimony later (cf. §2.3.2). In general I will take it that any differences between a sophisticated Gibbard-style expressivism and relativism are not relevant to my discussion: i.e. I will discuss both contextualism and relativism in what follows but leave expressivism to one side.

⁵ I will mention here, however, that there has recently been a spate of views in metaethics, classified as 'hybrid' expressivism, that develop the expressivist idea in a different way. They are 'hybrid' theories because they take moral utterances to express both beliefs *and* noncognitive states. The attraction of hybrid theories is the possibility of getting the best of both worlds: the advantages of both cognitivism and noncognitivism. They take moral utterances to express more than one kind of content: propositional content and expressive content. As my eventual diagnosis of the problem of aesthetic testimony suggests a kind of dual-content view for aesthetic utterances, my view has something in common with hybrid expressivist views. They are quite different in the details, however—and not just because they deal with different normative domains. (See e.g. Barker 2000, Copp 2001, Finlay 2005, Ridge 2006, Boisvert 2008, Bar-On and Chrisman 2009, Schroeder 2009.)

Contextualist theories typically relativize personal taste claims to particular judges or standards of taste, and it is therefore natural to understand such claims as attributing dispositional properties to objects: the disposition (say) to taste good to judge S or according to S's standard of taste.⁶ Such theories thus leave the testimony issue hanging. But a subjectivist theory that is specifically designed to deal with the problem of testimony in the way we are supposing must shift from an analysis in terms of the reciprocal dispositions of object and speaker—the disposition of X to please S, and the disposition of S to be pleased by X—to an analysis in terms of S's *actual* response to X. We must say something like this: for S to say that X is beautiful is for S to claim that X *is* giving her, or perhaps has given her, a certain kind of pleasure.

If we take it that typical claims of personal taste are to be relativized to the speaker or her standard of taste, contextualist views cannot allow for the kind of robust disagreement between someone who says e.g. “X is tasty” and someone else who says “X is not tasty,” that involves the assertion and denial of a common content. Margot claims that X tastes good to her (her, Margot), and Margaret claims that X does not taste good to *her* (her, Margaret). Any disagreement between Margot and Margaret must be a disagreement in attitude of a different kind—just as, if I say I like the vase and my husband says he hates it, we have not disagreed about the truth of a disputed proposition but we have disagreed about what to put in the dining room. (We may argue with each other about this, each trying to change the other's attitude, in ways that are similar to arguments about the truth of disputed propositions.)⁷

But for (what I will call) the *experiential* subjectivist view currently on the table, we have a prior, and more intractable, problem. It is in fact a problem that overlaps significantly with the problems of the simple emotivist view. The problem is that, if we tie the content of beauty claims to a subject's actual experience, we are left with no obvious way to account for the content of utterances involving ‘beautiful’ that are made in the absence of any experience of the object in question. Contextualist (and, for that matter, relativist) views that tie content to judges or standards do not have this problem. The judge or standard deemed relevant to the truth of any particular claim is taken to shift with context, but on such views the judge need not have actually exercised her judgment; the standard need not have actually been deployed. Such views allow for the properties in question (tastiness, beauty) to be like other dispositional properties.

If we are not careful, our experiential subjectivism may not even be able to account for negative beauty claims: claims that an object is *not* beautiful. For a more conventional contextualism that trades in the dispositions related to judges or standards of taste rather than experiences, the claim that X is not beautiful is unproblematic: if someone who says “X is beautiful” claims that X is such as to give her a certain kind of pleasure, someone who says “X is not beautiful” claims that X is not such as to give her that kind of pleasure. But the kind of experiential subjectivism here at issue, taken in one way, fares worse with negation even than the emotivist view. Emotivism has simply given up on compositional semantics for the terms to which it is supposed to apply, and undertakes to specify the meanings of sentences using those terms piecemeal. An emotivist may specify that “X is beautiful” expresses the speaker's pleasure in X, and that “X is not beautiful” expresses the speaker's lack of pleasure in X (her

⁶ Cappelen and Hawthorne (2009: 102-104), in defending contextualism for personal taste, suggest that such claims attribute judge-related dispositional properties to objects.

⁷ MacFarlane (2014: Ch. 6) calls this kind of disagreement ‘practical noncotenability’, because, although Margot and Margaret have made claims with compatible contents, they have expressed the noncotenable practical attitudes of liking and disliking X, respectively.

indifference or dislike). But giving up ordinary compositional semantics is not an attractive option for contextualist views, which trade in propositions: “X is not beautiful” must presumably express the negation of the proposition expressed by “X is beautiful.” But if the claim “X is beautiful” asserts the existence of a certain relation between the object and the speaker—namely, that the speaker is experiencing (or has experienced) the object pleurably—then the claim “X is not beautiful” must assert the lack of this relation. But “X is not beautiful” is not plausibly rendered as the claim that the speaker is not experiencing X (or hasn’t experienced X) pleurably, as this claim is true of any object the speaker simply hasn’t experienced at all.

We may do better if we reconfigure the content of the positive beauty claim in this way: maybe, in saying that X is beautiful, Margot makes a kind of *demonstrative* claim: that *this* experience of X is pleurable (or that *that* one was pleurable). In denying that X is beautiful, then, Margaret would be denying that her experience of X is (or was) pleurable: “This/that experience is/was not pleurable.” This move would thus yield the right result (at least, in this respect) for both positive and negative claims: the positive claim would be made on the basis of a pleurable experience of the object, and the negative claim would be made on the basis of an experience of the object that is not pleurable, and Margot and Margaret would once again be said to disagree in the practical way described above. But we would still be at a loss to account for the content of utterances using ‘beautiful’ in non-declarative and embedded contexts generally that are made in the absence of any experience of the object at all.

The problem of accounting for the use of ‘beautiful’ in this range of contexts has cropped up in commentary on Kant, who arguably presents a subjectivist view somewhat similar to what I have been discussing here, though his emphasis on the normativity of beauty gives his view greater complexity. Hannah Ginsborg (1998) argues for an interpretation of Kant along these lines; her idea is to inject reference to the subject’s actual pleasure in a beautiful thing into the normative claim that such pleasure is appropriate. What we thus get is this: “My judgment of beauty claims, not just that the object merits a certain kind of pleasure, but that it merits the very feeling of pleasure that I experience in it.” (465) The content of the judgment is demonstrative in the way just suggested, with the added normative element: i.e. that *this* pleasure that I take in the object is merited by it. On this suggestion, I clearly cannot make a legitimate beauty claim without myself having experienced pleasure in the beautiful thing; at the same time, however, given the reference to merit, my claim is that a feeling of pleasure like mine should be felt in response to the object by anyone in the same circumstances: “[I]f I feel this kind of pleasure when a given object is presented to my senses, I am at the same time aware that everyone else who perceives the object ought to share my state of mind. I am aware, that is to say, that all other perceivers of the object ought to experience the same pleasure as I do.” (463)

Ginsborg’s interpretation of Kant thus satisfies two significant strands of thought about beauty, both of which are important to Kant: what he calls the ‘autonomy’ of taste—i.e. more or less my topic here, the apparent requirement to make judgments of beauty, acquire beliefs about beauty, only on the basis of one’s own experience—and what he takes to be beauty’s normative character. Yet the added normative element threatens to upset our account of negative beauty claims and how people making positive and negative claims about the same object might be said to disagree. If, in judging that X is beautiful, Margot claims that her experience of X is a merited pleasure, then, in judging that X is not beautiful, Margaret presumably claims that her own experience of X is not a merited pleasure. There is a kind of disagreement here, and intuitively the disagreement is in fact more substantial than the kind of practical disagreement in attitude

described above, as is presumably required for disagreement about a normative matter: here Margot claims that she is appropriately responding to a normative demand, and Margaret is not responding to the normative demand that Margot claims is there. But while Margaret's supposed judgment—"My experience of X is not a merited pleasure"—makes it appropriate for her to *refrain* from claiming that X is beautiful, it is not a plausible stand-in for the judgment that X is not beautiful. If beauty has the normative character attributed to it here, the judgment that X is not beautiful surely must claim (at least in part) that X does not carry a normative demand for experiencing pleasure in it, and Margaret's claim about her own experience of X is too weak for that.⁸

Accounting even for negative judgments of beauty is thus once again a problem, on this Kantian view. Kant himself says very little about negative judgments of beauty; in fact, he says not much at all about how to understand judgments expressed by sentences employing 'beautiful' in embedded contexts generally: a problem I left hanging for any experiential subjectivist view, in which the content of beauty claims is tied to actual speaker experience. One remark Kant does make concerns generalizations, e.g. the judgment that roses are beautiful. (This is his own example: *CJ* §8, 5:215.) This cannot be a genuine aesthetic judgment, he says, as all aesthetic judgments are 'singular.' This is because they are based on the judge's feeling of pleasure in the object the judgment concerns, in a way I am taking, at Ginsborg's suggestion, to amount to a claim about the judgment's content. Kant calls the generalization "an aesthetically grounded logical judgment": 'aesthetically grounded,' because it is an inference from some number of genuine aesthetic judgments about individual roses; 'logical' because, unlike those judgments about individual roses, what is attributed to roses in the generalization cannot be a relation to the speaker's actual experience of them, but has to be an ordinary property, as with any objective cognitive judgment. Kant says that in a genuine aesthetic judgment, the feeling of pleasure "accompanies the representation of the object and serves it instead of a predicate" (*CJ* §36, 5:288). In the aesthetically grounded logical judgment, what is predicated must be different: an ordinary concept. In saying that roses are beautiful, the speaker of course means to assert something of roses generally: of more roses than those she has actually experienced.

Commentators (including Ginsborg) have used Kant's meager remarks here to speculate about what Kant's aesthetically grounded concept could be, and how it could help with the problem of embedded contexts, including negation.⁹ The natural suggestion is that, in claiming that roses are beautiful, a speaker claims that roses in general are *such as to be* rightly judged aesthetically to be beautiful—i.e. such as to merit genuine, aesthetic, singular judgments of beauty.¹⁰ This derivative concept of beauty, predicating an object's disposition to merit

⁸ Ginsborg (forthcoming) argues that the kind of disagreement illustrated here by Margot and Margaret, in which one person claims an object to be beautiful and the other refrains from judging it to be beautiful, is robust enough to account for the difference, acknowledged by Kant, between disagreement about the agreeable and normative disagreement about the beautiful. She notes that for Kant, the latter but not the former involves quarrel or argument (Kant's *Streit*), and suggests that the disagreement between Margot and Margaret constitutes a quarrel in a way that a disagreement about the agreeable does not.

⁹ E.g. Miles Rind (2003), Hannah Ginsborg (2014: 109-110).

¹⁰ Rind's actual suggested formula is, "capable of being judged with pleasure in mere reflection" (2003: 69). I take it he has put it like this to guarantee that the concept of the beautiful is rightly applied only to objects that (as it were) really are beautiful, not also to objects that some people may mistakenly judge to be so. I have tried to reflect the same concern with the use of 'rightly' above, as I do not want to get too far into the underlying details of Kant's account. Ginsborg (2014: 109) says that the concept at issue could be "a derivative concept of beauty which applies to just those things to which a judgement of beauty, on its original construal, is appropriate."

judgments of beauty (in the basic, genuinely aesthetic sense), could then be used to account for the use of ‘beautiful’ in embedded contexts generally. In the case of negation, we could say that “X is not beautiful” expresses the judgment that X does not merit a judgment of beauty (in the basic, aesthetic sense), or, in other words, that X does not carry a normative demand to take pleasure in it of a kind appropriate to beauty. This is the judgment that eluded Margaret above.

However, there is a fundamental problem with this suggestion: it does not give us the right answer regarding such judgments and testimony. We have been on the lookout for something for ‘beautiful’ to mean when it clearly cannot predicate a relation between an object and the actual feeling of pleasure a speaker takes in it. We need this for cases in which ‘beautiful’ is used in utterances made in the absence of any experience of the object at all: in questions, say, or conditional statements. But we also need it for the examples of judgments we have just been looking at—negation and generalization—which *are* made on the basis of experience, but which clearly cannot incorporate that experience into their content in the same way as the positive, singular judgment. Neither the negation nor the generalization can be *about* the speaker’s firsthand experience of the object or objects in question, on this experiential, normative view, but both examples do in fact require that the judgments be made *on the basis of* such experience. The phenomenon concerning testimony here at issue is as pressing for them as for the positive, singular judgment. “X is not beautiful” must be asserted on the basis of the speaker’s experience of X, just like “X is beautiful”; “Roses are beautiful” clearly can’t be asserted on the basis of the speaker’s firsthand experience of all roses, or the generality of roses, but it does require that the speaker have firsthand experience of some roses, on which her inductive generalization is based.

The current suggestion does not give us the right answer regarding such judgments and testimony, because to judge, and to go on to claim, that Xs in general are such as to merit judgments of beauty (in the basic, aesthetic sense), or that X does not merit a judgment of beauty (in the basic, aesthetic sense), *are* things that I should be able to do on the basis of testimony. Even if the judgments under consideration are that Xs merit judgments of beauty *made by me*, or that X does not merit a judgment of beauty made by me, these would still appear to be judgments I could legitimately make on the basis of testimony, as I may know I share a taste in Xs with somebody else, and be told by her that Xs are beautiful; or I may trust the word of a judge who tells me that X is simply not worth it.¹¹

Miles Rind (one of the commentators running with Kant’s remarks about roses to suggest the derivative concept of beauty) presumably agrees with this assessment of how the proposed derivative concept fares with respect to testimony, as he takes Kant’s denial that genuine (singular) judgments of beauty employ a concept of the beautiful, to be the only option for explaining why such judgments must be based on firsthand experience:

[U]nless I have actually found a thing beautiful, in the sense of being struck by its beauty, I am in no position to affirm without qualification that it is so. ... This could not be the case if in

¹¹ Note that in Ginsborg (forthcoming), mentioned earlier, she concedes the problem described here concerning the current suggestion for “X is not beautiful”: “The problem is that in introducing the concept of *calling for a judgment of beauty* we have introduced something which looks very much like an objective property. ... [I]t seems to be a property which an object can have independently from an individual’s response to that object. ... But ... ordinary usage also prohibits asserting “This is not beautiful” if one hasn’t seen the object oneself. So if “This is not beautiful” is understood as making the second order judgment that the thing does not call for a judgment of beauty, rather than as simply expressing one’s absence of pleasure, then we have the question of why not only one’s own lack of pleasure, but also someone else’s lack of pleasure, could not serve as a ground for that judgment.” (15)

judgements of the form 'X is beautiful', when properly made, the word 'beautiful' expressed a concept; for if it did so, then testimony or description could in principle yield sufficient evidence to justify a judgement of beauty apart from any exercise of taste on the part of the judging person. (2003: 71)

However, Rind does not discuss the implication regarding testimony and the use of 'beautiful' in the examples we have looked at. Possibly he thinks that such judgments (e.g. "Xs are beautiful", "X is not beautiful"), unlike positive, singular judgments, *can* be transmitted through testimony. Possibly Kant would have taken the same stance on "aesthetically grounded logical judgments."

I have gone into some detail about this kind of subjectivism, exemplified in what may be its most plausible form by Ginsborg's interpretation of Kant, in order to highlight the difficulties of this kind of solution. This kind of subjectivism faces significant challenges, just as emotivism does. As we have seen, it is not impossible to assign content to judgments involving the beautiful that is independent of the subject's actual experience, when necessary. But the fact that, when we do this, we *still* have the problem of testimony to deal with, makes this kind of subjectivism lose much of its appeal. If negative judgments of beauty and generalizations are accounted for in the way outlined here, we have the same problem with testimony for *them* that we had in the first place, for the simple, positive judgment. But my consideration of this kind of subjectivism was originally motivated by the testimony issue. I will return to Kant later, but here we should note that his account of the experience of beauty, and the experiential judgment made by someone who experiences beauty firsthand, does not easily lend itself to a semantics for 'beautiful'.

Rind expresses in the quote above, however, what is the general worry if we reject both emotivism and this kind of experiential subjectivism: if the word 'beautiful' corresponds to a 'concept', in Kant's sense, then use of the word in basic, singular judgments would seem to attribute an objective property to objects, and there would be no apparent reason why such judgments couldn't be made on the basis of testimony.¹² We can see from contemporary contextualist and relativist views that, in fact, the problem does not just concern objective properties, although the normative element in Kant's account of beauty makes objectivism an attractive option, if we jettison the account's semantic reliance on the subject's experience. But contemporary contextualist and relativist views of personal taste also generally accept *Experience Independence*, in that they take the content of claims of personal taste to be related to a judge or her standard of taste rather than to any particular manifestation in her experience of that standard. This shift puts to rest the structural difficulties with experiential subjectivism. But the problem with testimony is then also a problem for such views: I may surely know, after all, at least in some cases, that something measures up to my standard in a certain way, even though I haven't experienced it myself. It is as odd, however, to say "That's delicious" on the basis of testimony, as it is to say "That's beautiful."

Nevertheless, I will turn now to possible solutions to the problem of testimony that accept *Experience Independence*.

¹² Kant's denial that 'beautiful' expresses a concept in basic, singular judgments, and his claim that such judgments are not 'logical', do not imply that he is himself an expressivist, although they might seem to do. In saying that 'beautiful' does not express a concept, his intent is to deny that judgments of beauty predicate any concept to objects independently of the pleasure taken in a beautiful object, but this is consistent with non-expressivist views, as we have seen.

2.3 An epistemic problem?

We have just seen an interpretation of Kant that would explain why he, at least, would have insisted on the illegitimacy of transmitting aesthetic belief through testimony: on Ginsborg's interpretation, Kant thinks not just that a judgment of beauty is made on the basis of a judging subject's pleasure, but that any judgment of beauty contains an ineliminable reference to the judging subject's pleasure. If this is the case, my belief that X is beautiful has as part of its content my pleasure in X, and I cannot feel this without experiencing X. I thus cannot, on this view, acquire the belief without the experience.

However, the issue with testimony is pressing for anyone who accepts *Experience Independence*, as it is not at all clear why beliefs with such apparently unproblematic content should behave differently from myriad other beliefs that have no general problem with testimony. It seems often to have been assumed by those who agree that there is some problem with aesthetic testimony (and who also assume *Experience Independence*), that the problem is an epistemic problem—an assumption which is not surprising, I suppose, as testimony is obviously an epistemological issue. Thus Aaron Meskin (2004: 66-67) says the following about how the problem has been conceived:

[W]hy does testimony in the aesthetic realm ... seem to lack the epistemic value that testimony has in other domains? Not only does it seem as if we rarely (if ever) accept aesthetic testimony as a basis for our own aesthetic judgments, but it also may seem that it would be illegitimate to do so. It does not seem that we can gain much in the way of aesthetic knowledge from others, nor even much in the way of justified aesthetic belief. In fact, it is often claimed that it is simply *impossible* to make an aesthetic judgment of an object solely on the basis of others' testimony—not even an unjustified one. At least this is how the situation has appeared to many—perhaps most—aestheticians, going back at least as far as Kant in his *Critique of Judgment*.

Ginsborg's interpretation of Kant does have as a consequence that making aesthetic judgments on the basis of testimony is impossible. If we accept *Experience Independence*, however, and look for an epistemic solution to the problem of testimony, it is hard to find one that is convincing. As far as I can see there are two candidates: one based on the claim that aesthetic characterizations of an object cannot be fully understood by anyone who lacks experience of that object, and one based on the supposed unreliability of aesthetic testimony.¹³ I will look briefly at each.

2.3.1 Lack of 'full understanding'

Philip Pettit (1983) has put forward the first idea. He likens how it is with aesthetic testimony to this kind of situation: someone I trust says, "He is fair-haired" in my hearing, but I cannot see to whom he is referring. "I might be said to know that the assertion 'He is fair-haired', on the lips of my informant, expressed a truth, but knowledge that such an assertion is true may not involve knowledge of the truth expressed. ... Because testimony does not enable

¹³ There is another epistemological issue concerning aesthetic testimony that has been discussed in the literature but that I will not discuss here: this is the question, familiar from Frank Sibley's work, of whether the presence of aesthetic properties can be inferred from a description of non-aesthetic properties. I am ignoring this issue because even if such inferences never go through, and we thus cannot acquire aesthetic beliefs by inference from testimony as to an object's *non-aesthetic* properties, this would not explain the illegitimacy of acquiring them directly from testimony as to an object's *aesthetic* properties. (See Sibley 1959.)

me fully to understand what is expressed by [the utterance], as this is asserted by my informant, so it does not give me a title to full knowledge of what is expressed by [it].” (28) We might add (which Pettit does not do explicitly) that even though I know that my informant said something that is true, I cannot even be said to fully *believe* that “he” is fair-haired (assuming, as Pettit seems to suggest, that ‘full belief’ is *de re* belief), nor can I felicitously assert, “He is fair-haired.”

What is it about aesthetic characterizations, then, that makes them fully understandable only by someone who has perceived the object in question? Pettit posits that the answer can be found in the phenomenon of ambiguous Gestalts, and asks us to consider a certain figure that looks sometimes like the letter ‘B’, sometimes like the numeral ‘13’, depending on whether it is seen as part of a string of letters (A B C) or a string of numerals (12 13 14). If the background ‘reference class’ is other letters, the figure is seen as a ‘B’; if it is numbers, the figure is seen as ‘13’. This example suggests to Pettit that, generally,

for any property which an object can display in perception, the object displays that property only in so far as it is positioned in an appropriate class: that is, only in so far as the perceiver knows what the relevant contrasts are. The pictorial property of redness will be displayed only in so far as the bearer is positioned by reference to the colour paradigms or, allowing for denseness, the colour spectrum. The aesthetic property of sadness will be displayed only in so far as the bearer is positioned by reference to certain parallel contrasts. (30-1)

Pettit thinks that there is an important difference, however, between redness and sadness (which he uses as an exemplar of aesthetic properties). The ‘positioning’ for redness is “by reference to something given once and for all,” but for sadness it is “by reference to something which may change from case to case. ... It requires only normal information and memory to position an object appropriately for colour; it requires imagination to position it so that it displays a property like sadness.” (31)

This is important because Pettit thinks that, in order to fully understand what is expressed by “X is sad,” for any X, one must fully understand this associated conditional: “X is sad if and only if X is such that it looks sad under standard presentation and under suitable positioning.” (33) And he thinks that to fully understand this conditional, one must know what positioning for X (if any, presumably) makes it look sad. And to know what positioning for X, if any, makes it look sad, one must have seen X.

The case is different for “X is red,” Pettit says. To fully understand “X is red,” one must also fully understand its associated conditional, “X is red if and only if X looks red under standard presentation.” But in this case, there is no need to mention positioning explicitly, as the positioning for redness is the same for all red things, and it can thus “be taken as a further aspect, over and beyond normal sight and normal illumination, of standard presentation” (31). We can fully understand the conditional without seeing X, because we already know what it is for anything to look red, and what standard presentation for redness (including its appropriate positioning) involves.

We can put Pettit’s point like this: full understanding of the sentence that results from a particular application of the predicate ‘is red’ to some object X, does not require seeing X, because Pettit’s requirements for understanding the sentence’s associated conditional—knowing what it is for anything to look red, and knowing what standard presentation and appropriate positioning are for redness—are in fact required simply for understanding the predicate ‘is red.’

The positioning for redness, and, presumably, these other things, are “given once and for all” because knowledge of them is required for knowing the meaning of the predicate. And it is indeed plausible that someone who has not managed to position redness with respect to the rest of the color spectrum—someone who is colorblind, say—does not have a full understanding of the predicate ‘is red’ (although she might have a partial understanding—it need not be as if ‘red’ is a word in a completely foreign language).

In contrast, Pettit essentially claims that the requirements for understanding the associated conditional for “X is sad” are *not* met simply in virtue of understanding the predicate ‘is sad’ (as this is applied to works of art). In particular, we cannot know the appropriate positioning for X’s sadness, Pettit claims, unless we have seen X ourselves: “Only someone looking at a picture and putting it imaginatively through various positionings can understand what that positioning is under which the picture looks sad. One fixes the positioning, one finds the appropriate reference class, only in so far as one succeeds in making the picture display the appearance of sadness.” (33)

In one way, what Pettit says here is clearly too strong, even by his own lights. It is surely unnecessary that someone actually be looking at a picture in order to understand its positioning for sadness; remembering how the picture looks must suffice.¹⁴ But putting this aside, his main point here is still a very strong one: he claims that the requisite knowledge of positioning for a full understanding of “X is sad” depends on actually seeing X as sad. This is a very implausible result. For one thing, it cannot make sense of disagreement about X’s aesthetic properties. If I say X is sad, but you disagree, saying X is tritely sentimental, you cannot, according to Pettit, have fixed X’s positioning for sadness, as you have not succeeded in seeing it as sad. But then you do not even (fully) understand what I have said—and yet you have the nerve to disagree with me. (Compare: I say, “He is fair-haired,” but you don’t know to whom I refer; you nevertheless say, “No, he’s not.”)

We might take Pettit as simply having overstated his requirement. Knowing the appropriate positioning for X’s sadness, in the way required for a full understanding of “X is sad,” might not require having seen X as sad; but it still might require having seen X. Maybe I fully understand “X is sad” only if I have seen X and brought my imagination to bear on it in a way that can be thought of as considering *whether* X is sad. In considering whether X is sad, I would find out what it is for X to look sad or not to look sad, and understand the appropriate positioning of X for this question.

I agree with Pettit to this extent: it is plausible that the imagination is involved in considering whether a picture is sad in a way that does not seem to be the case for considering whether a picture is red. It is also plausible that this has something to do with positioning a picture in the way that Pettit has in mind—i.e. being aware, at least to some extent, of the potential sad-making features of the picture, features which, if different, would have made the picture have a different character—and that an understanding of exactly what the positioning is under which the picture looks sad, or not, can only be achieved by seeing the picture. This understanding is not, then, something achieved simply in virtue of understanding the predicate ‘is sad’ as it is applied to pictures. In contrast, it does seem that we already know how to position an object for detection of redness, simply in virtue of our understanding of ‘is red’: we do not have to see it in order to gain this knowledge.

¹⁴ This is also noted by Malcolm Budd (2003).

Let's assume that we do have to fully understand this conditional—"X is sad if and only if X is such that it looks sad under standard presentation and under suitable positioning"—in order to fully understand "X is sad." Let's also grant that someone who hasn't seen X doesn't know how to position it for detection of sadness. Pettit claims that someone who doesn't know this cannot fully understand the conditional. The problem is that we have been given no good reason to think that this is the case. A full understanding of the conditional surely need not require that we know exactly what the positioning is under which X looks sad, or not; we need only understand that some suitable positioning or other, of a kind that is generally familiar to us given our understanding of the predicate 'is sad,' is being claimed to make X look sad.

Perhaps it might be claimed that we can't know what it is for X to look sad, without seeing X, and that this is why seeing X is necessary for a full understanding of the conditional. This claim is somewhat tempting, as, understood in one way, "knowing what it is for X to look sad" seems to require actually having seen the way in which X does look sad—having seen X as sad. But, as I noted above, having seen X as sad is too strong a requirement for full understanding of the conditional. I can know, in a general way, what it is for something to look sad, if I understand the predicate 'is sad' as this is applied to pictures, and while this knowledge is necessary for understanding Pettit's conditional, it is also possible to have it without having seen X.

Pettit's argument relies on a misleading comparison. He claims, plausibly, that knowledge of certain kinds of positioning is necessary both for seeing an object to be red and seeing a picture to be sad. He claims, plausibly again, that knowledge of the positioning for an object's redness is necessary for fully understanding any sentence of the form "X is red." He concludes that a knowledge of the positioning for an object's sadness is necessary for fully understanding any sentence of the form "X is sad." But the two cases are different. The knowledge of positioning for detecting an object's redness is presumably required for understanding the predicate 'is red' in the first place; but the knowledge of positioning for detecting a particular object's sadness is not required for understanding the predicate 'is sad'. And Pettit has given no reason to think that this knowledge is nevertheless required for a full understanding of "X is sad."

It may be that knowledge of something akin to positioning is required for a full understanding of 'is sad,' but, if so, it is not the kind of positioning that Pettit says must be imaginatively undertaken to detect sadness in a particular picture. It is not obvious what exactly it is to call a picture sad—it is not necessarily that it makes us sad, or that anyone in the picture is sad, or even that the artist intended anything about the picture to be sad—but as an example, without worrying about whether it's actually right, let's say that the sadness of a picture is in some way analogous to a person's expression of sadness—in her look, say, or tone of voice¹⁵—and that this is on some level at least what people are responding to when calling pictures sad. Being 'analogous' to the expression of sadness obviously gives a lot of freedom to possible describers of a picture as sad; but that is what we should expect from aesthetic description, the appropriateness of which is often disputed. In any case, it is plausible that I must have managed

¹⁵ This idea is adapted slightly from Roger Scruton (1974). His example is the sadness in music, and he proposes that "the experience of hearing the music is analogous to the experience of hearing the emotion." (127) He offers this in the context of his 'affective' aesthetic theory, which in many ways is like an expressivist theory. In fact, Pettit cites Scruton's view as an example of the kind of anti-realism to which he is opposed. The anti-realist Scruton is concerned to couch his proposal in terms of the subject's experience, rather than any realistically construed aesthetic property of the music. In my example, I simply refer to the sadness of a picture.

to position sadness with respect to the other emotions in order to understand a description of a *person* as sad, in a way something like what Pettit has in mind for ‘positioning,’ and this knowledge would presumably include knowledge of how expressions of sadness differ from the expressions of other emotions. If I don’t know how *people’s* emotional expressions differ, then, on this proposal of what it means to call a *picture* sad, I will not understand a description of a picture as sad, as opposed to tender, or overly sentimental, or what have you. But this knowledge, which on this account is required for understanding the predicate ‘is sad,’ as this is applied to pictures, does not include (fore)knowledge of the particular way in which a given picture will be found to be analogous to a person’s expression of sadness; and it does not need to.¹⁶

To be sure, if we are told simply that X is sad, there are many things about its sadness that we don’t know—we don’t know, for instance, what makes it sad. But when my friend tells me that her house is red, there are many things about its redness that I don’t know either, and yet I fully understand her claim: I don’t know exactly what shade of red her house is, or how much of it is red. It is true that I cannot *appreciate* a picture’s sadness unless I see it, but I can see no reason for such appreciation to be required for fully understanding a characterization of the picture as sad.¹⁷ I understand perfectly well what my friend is saying when she tells me that a picture, which I have never seen, is sad, although I do not myself see the sadness and I also do not know, unless she tells me more, in what way it is sad. But I do not need to know this, nor do I need to see the picture, to understand what she tells me.

2.3.2 *Is aesthetic testimony unreliable?*

I will turn now to the second possible epistemic solution to the problem of aesthetic testimony that I mentioned earlier, which relies on aesthetic testimony’s supposed unreliability. When we ask people to tell us the color of a thing, we are usually justified in trusting their eyesight and their honesty. But people’s judgments on aesthetic matters vary much more than their judgments of color. Of course, many who deny *Experience Independence* take this as further evidence for their views: if it is common for people to have such different, apparently conflicting, aesthetic opinions about the same thing, isn’t it reasonable to suppose that those opinions merely reflect their different aesthetic responses to that thing, rather than being beliefs,

¹⁶ The difference I have been discussing between ‘red’ and ‘sad’ has much to do with the different uses that can be made of vague expressions. If we both know what the standard of tallness is around here, and I say, “Margot is tall,” you will take me to be making a claim about Margot’s height to the effect that it exceeds a certain value. But I may use “Margot is tall” in a different way, to illustrate what counts as tall around here. Often claims about tallness will perform both functions: they will simultaneously refine the hearer’s knowledge of an object’s height and her knowledge of the standard for height relevant to the conversation. (See e.g. Chris Barker 2002.)

A typical claim about redness will rely on people’s prior knowledge of the standard for redness; but an aesthetic claim about sadness will, in much larger part, serve to sharpen the ‘sadness’ standard (as it were) for the conversation in which it takes place. A speaker typically makes such a claim as a result of a particular experience that makes it compelling for her to make the comparison with sadness and thus to refine the standard in the way that she does. Someone who does not share this experience will obviously not feel similarly compelled; but, as I have been arguing, this does not make her incapable of understanding the claim that is made. She is likewise perfectly capable of understanding the claim “Margot is tall,” when this is used to set the standard for tallness. In both cases, of course, further questions may be asked and answered, to give her further knowledge.

Although I disagree with Pettit’s diagnosis, he does latch onto an important difference here between aesthetic claims and other kinds of claim. I will return to these issues concerning vagueness and aesthetic claims later.

¹⁷ This point is made by Budd (2003).

so often missing the mark, about the thing's properties, properties which are independent of each person's response?¹⁸

There is perhaps much less genuine disagreement about aesthetic matters than it might appear. A lot of aesthetic talk really does seem to concern people's likes and dislikes, and people who simply don't like the same things don't necessarily have conflicting aesthetic beliefs. If I prefer reading 1930's murder mysteries to teenage vampire stories, and you (my seventeen-year-old friend) are the reverse, we don't necessarily think that what we each happen to prefer is better than the other, just as I might prefer biking and you running, or I wearing red and you blue, without each of us elevating our preferences to normative demands. A lot of aesthetic discussion takes place on this level. However, having said that, there is a lot of genuine disagreement as well—such as when I think the painting sad and you, on the contrary, think it tritely sentimental—and it really does seem that there's more aesthetic disagreement than disagreement about color. An objectivist about aesthetic claims must hold that many people's aesthetic beliefs are false—more often than people's color beliefs—and hence that aesthetic testimony is in general less reliable than testimony about color. There is, no doubt, a wide range of possible objectivist views; I count among them a view according to which “X is beautiful,” for instance, attributes to X the (dispositional, yet normative) property of being such as to merit a kind of pleasure.¹⁹

Possible contextualist and relativist views may also take *Experience Independence* to be true. The kind of subjectivism I discussed earlier (§2.2) is a contextualist view for which *Experience Independence* is held to be false: on this view, the content of an utterance of “X is beautiful” is taken to include the speaker's pleasure in X. Other kinds of contextualism, however, may take the relevant context of utterance to include, not the speaker's particular aesthetic experience, but the speaker's standard of aesthetic taste. In a view of this kind, for S to say that X is beautiful could be for S to claim that X gives a certain kind of pleasure to people with standard of taste T (where T is shared by S). For relativist views, the content of “X is beautiful” would not itself include reference to a standard of taste, as in this contextualist view, but would be assessed for truth or falsity relative to a standard of taste, usually either the speaker's or that of the person assessing the claim.²⁰ These views can of course easily accommodate the variability of aesthetic opinion, as there can be many different standards of taste.

¹⁸ Note that *Experience Independence* doesn't say that aesthetic claims ascribe *entirely* response-independent properties to objects; such a thesis would be wholly implausible for aesthetic properties. It doesn't actually say they ascribe properties to objects at all, although presumably adherents of *Experience Independence* would think that they do. Different sorts of views can entail *Experience Independence*, as I am about to go on to illustrate in the main body of the text.

¹⁹ John McDowell arguably puts forward a view of this kind in (1998a) and (1998b). Ginsborg attributes this view to him in (1998).

²⁰ So I have mentioned four kinds of view in this paragraph. What I have called ‘contextualism,’ John MacFarlane (2014) calls ‘standard’ or ‘indexical’ contextualism, as the view takes there to be an implicit indexical element in aesthetic sentences. The subjectivism I discussed earlier takes this to refer to the speaker's aesthetic experience; I mention another possible contextualist view here in which reference is made instead to the speaker's standard of taste. I then mention two different views under the name ‘relativism’; MacFarlane calls the first of these ‘nonindexical contextualism,’ as the speaker's standard of taste is not part of the content of the proposition she expresses, like the referent of an indexical, but part of the context relevant to determining the proposition's truth-value. He reserves the name ‘relativism’ for the last view; the relevant standard of taste here is not necessarily the speaker's—not necessarily part of the context of utterance at all—but the assessor's. (He endorses this kind of view

It is perhaps surprising that these views also have a problem with aesthetic testimony. In the contextualist case, this comes entirely of the shift to dependence on standard of taste, rather than speaker experience. On the view I described, someone can have a standard of taste suitable for making the claim that X is beautiful, without having actually taken pleasure in X. While having some standard of taste or other must depend in general on the kinds of aesthetic reactions one has to things, someone's standards can be identified independently of any particular reaction to any specific object. Thus it seems I can know I share aesthetic standards with someone else—a friend or critic, say, with whom I often agree—even when we haven't experienced exactly the same things: in particular, even when she has seen X and found it beautiful, but I have never set eyes on it. But if I can know other people's tastes to be similar to my own, it is unclear on this view why I cannot in these cases adopt their aesthetic beliefs. A possible epistemic solution to the problem with testimony on this view may try to establish that it is much harder to know when standards of taste agree than it might appear.

In the case of relativism, the possibility of such views shows, in fact, that *Experience Independence* doesn't quite cast its net wide enough²¹: not even the standard of taste, on both of the relativist views I allude to, forms a part of the propositional content of an aesthetic claim, as it does in the contextualist case. So *Experience Independence* is of course held to be true by these views, but this is not by itself what renders them problematic. The problem such views have is that the relativity involved, even though it relates to the context for truth evaluation of an aesthetic claim rather than the content of the claim, still does not make reference in any way to the speaker's specific prior aesthetic experience: the claim is true or false relative to a *standard of taste*, and, once again, this leaves it unclear why I cannot then adopt another's aesthetic belief if I know her standards to be similar to mine.

A possible solution in both the objectivist and contextualist/relativist cases, then, may focus on the difficulty of identifying the right kind of aesthetic informant: for the objectivist, an informant who can be relied on to be correct in her aesthetic judgments; for the contextualist or relativist, one who can be relied on to share your aesthetic standards. The claim may then be that identifying these people is in general too hard to justify adopting others' aesthetic beliefs in the absence of personal experience. We would then suppose there to be a basic understanding of this fact embodied in everyday practice, that prevents us from adopting others' aesthetic beliefs on the basis of their testimony, and from making testimony-based aesthetic claims.

In support of this idea, it is true that we reserve judgment about the value of the aesthetic testimony of informants whose aesthetic standards are completely unknown to us, in a way that we do not with testimony about color or any straightforward matter like the direction of the post office. If someone we do not know at all tells us confidently to look for the red house or to go two blocks north and then turn left, we tend to believe she knows what she's talking about and act accordingly. If someone we do not know at all tells us confidently, however, that some movie we've never heard of is great, we will not, just on that basis, put it on our 'To See' list.

However, the proposed solution is committed to overestimating the amount of unreliability in aesthetic testimony. While we reserve judgment about the value of the aesthetic testimony of informants whose aesthetic standards are unknown to us, we *don't* reserve judgment about the value of the aesthetic testimony of those whose standards we know well. I value my friend's, or

for some areas of discourse, like talk of food's being tasty. As far as I know he has mostly steered clear of applying the view to the aesthetic case, although occasionally his examples involve aesthetic terms (cf. eg. 2014 §4.7).

²¹ I have for simplicity ignored this wrinkle about *Experience Independence* in the rest of this work.

a favorite critic's, judgment when she mentions good movies she's seen or books that she's read, so much so that I may go to see those movies or read those books on her recommendation. We are as willing to take others' aesthetic *advice*, when we trust our informant, as we are to take their advice about the whereabouts of the post office. And we very often give and take such advice, even if we are somewhat more circumspect with it than with post office advice. It seems, then, that the contention that reliable or like-thinking aesthetic informants are too hard to find, cannot be sustained.

There is certainly a puzzle here, though: what does 'valuing' someone's aesthetic judgment amount to (I tried to find, in the last two paragraphs, an innocuous formulation of this), if not taking on her aesthetic beliefs? The giving and taking of aesthetic advice seem to stand in contrast to the intuitions suggesting that we don't acquire unqualified aesthetic beliefs on the basis of testimony. The phenomena related to aesthetic advice do not, however, require straightaway that we regard those intuitions as mistaken or at least distorted. Meskin (2004: 72) offers this possibility for someone wanting to understand aesthetic advice while retaining the contrasting intuitions:

What I am apt to judge on the basis of a favorite reviewer's judgment is not that the film in question is [say] terrible, but that it is likely that I would find the film terrible if I were to watch it, or perhaps merely that I would dislike it a great deal. ... Perhaps I even *know*, on the basis of the reviewer's testimony, that I would find the movie terrible and/or dislike it. But these judgments are not aesthetic judgments ... rather, they are judgments about what my aesthetic judgments would be, or perhaps judgments about what I am likely to like.

I want to put aside the suggestion that my judgment, on the basis of another's aesthetic testimony, concerns merely what I will like or dislike, as I see no reason why we would not at least be entitled to judgments concerning what our aesthetic judgments about a thing *would* be, if we were to see it. This leaves two options: on the basis of another's testimony that the film is terrible, I judge either that I will *probably* find the film terrible, or I judge that I *will* find it terrible (and may even know that I will).

There is no doubt that we make judgments of the first kind on the basis of others' testimony. I might say, for instance, "Oh, do we have to go see that? I'll probably think it's terrible." We even say (and this formulation actually seems more natural), "It'll probably be terrible"—which clearly has a different meaning but which is used in much the same way—and even just, "It's probably terrible." This gives an opening to advocates of the solution which relies on the claim that aesthetic informants are too unreliable: they can say that we often manage to find good informants, but that there is enough uncertainty even with them to make beliefs about what is aesthetically merely probable the best they can give us.

The problem with this is that we hear about many aesthetic matters that do not merit this epistemic caution. Before I saw the Grand Canyon, I did not *doubt* others—not even a little bit—when they said that it was breathtaking, even though I would not have said, "The Grand Canyon is breathtaking," and would not have described myself as believing that the Grand Canyon is breathtaking. This brings us to the second option mentioned above. We seem to express this kind of certainty about others' aesthetic testimony by saying, "I will find it breathtaking," "I know I'll find it breathtaking," or—as before, more directly—"It'll be breathtaking," or, in some cases, "It *must* be breathtaking." We don't say it *is* breathtaking

unless we have seen it ourselves, but the example suggests that the reason for this doesn't depend on some kind of completely general epistemic problem with aesthetic testimony.

Moreover, as more than one philosopher has pointed out—at least for objectivist aesthetic views—a general epistemic embargo on aesthetic testimony, on the grounds of its unreliability, would seem to have to apply to one's own views as well as those of others.²² If there were a general epistemic reason to distrust others' aesthetic beliefs, and we take them to be objective, the same reason would have to prevent me from trusting my own: it does not seem, after all, that I am uniquely well placed, amongst all the people of the world, to discern aesthetic truths.

2.4 Other solutions

Many people have recently come to agree that an epistemic solution to the problem of aesthetic testimony looks unpromising. In this section I will look at two alternative solutions, one offered by Meskin and one by Robert Hopkins. Both think that an epistemic solution will not work.

2.4.1 Aaron Meskin's 'reverse error theory'

Meskin (2004) is an objectivist about aesthetic discourse, and would therefore take *Experience Independence* to be true, but he thinks that much resistance to aesthetic testimony is justified, especially testimony about art, because of its general unreliability. He thus agrees up to a point with the second epistemic solution detailed above. However, he does not think all aesthetic testimony merits epistemic caution: he thinks that the testimony of experts in the various genres of art, as well as testimony about nature generally, can be as reliable as testimony on other matters. His explanation of why we nevertheless resist acquiring aesthetic beliefs on the basis of testimony, even when they would be epistemically safe, depends on attributing to us a kind of folk theoretical mistake: we mistakenly take aesthetic testimony's relatively high level of unreliability as evidence that aesthetic matters are not objective, and this leads us to resist testimony that we shouldn't resist. He says:

[S]ome of our resistance to accepting aesthetic testimony may be best explained by reference to relativist and subjectivist strains in the folk conception of beauty and artistic value. ...[T]he epistemic weakness of some aesthetic testimony can be understood as a product of its unreliability. This unreliability, in turn, partly explains the existence of the relativist strain of thinking about beauty. But aesthetic testimony is not always unreliable. So it can, under certain conditions, provide justification. (68)

How exactly does the folk strain of relativist or subjectivist thought lead to this resistance? Meskin claims that this is the result of

a combination of massive unreliability in aesthetic testimony about art and some quite distinctive phenomenology with respect to our own aesthetic judgments. The massive unreliability manifests itself as widespread divergence in aesthetic judgments. The phenomenology of aesthetic experience, on the other hand, is one of immediacy and manifestness. Aesthetic qualities do not seem hidden or difficult to discern. ... The obviousness and immediacy of our aesthetic experiences, in combination with the appearance of widespread disagreement, suggest to us that a sort of relativism or subjectivism about aesthetic properties is true. That is, although

²² E.g. Rind (2003: 72), Robert Hopkins (2011: 141).

claims of beauty have the surface form of claims about objective features of the world, the folk view is that in truth they amount to little more than declarations of how objects seem to us, or how we experience them. (88)

So, because we think that the content of aesthetic claims, and aesthetic beliefs, is more or less as subjectivism or relativism would have it, we are unwilling to acquire the beliefs or make the claims except as the result of our own experience.

Meskin notes that there is an “objectifying strain in folk aesthetic practice” (85), by which he seems to mean that ordinary conversations about aesthetic matters involve argument and, in general, broadly normative claims (cf. his fn. 48), but he thinks that the subjectivist or relativist element in folk thought about that practice is nevertheless dominant. He cites, for instance, the sayings “De gustibus non disputandum est” and “Beauty is in the eye of the beholder” as representative of folk thought (84).

It is certainly plausible that folk aesthetic theory is dominated by this strain. But I think that Meskin does not have an accurate picture of the phenomena he sets out to explain. These are, he says, both psychological and epistemological. The psychological phenomenon is that “we do not accept aesthetic testimony to the same extent that we accept other sorts of testimony”; the epistemological phenomenon is that “aesthetic testimony does not have the epistemic value possessed by other forms of testimony” (68). He accounts straightforwardly for the latter by advertent to aesthetic testimony’s comparative unreliability. But he thinks that the psychological phenomenon is not absolute: that what he calls our resistance to aesthetic testimony surpasses what we would expect just from general recognition of its epistemic shortcomings, but that we do sometimes, on occasion, acquire an aesthetic belief or make an aesthetic claim on the basis of testimony. This is evident in his use of the phrase “to the same extent” in the quote above; he also, for instance, claims elsewhere that “we *rarely* accept aesthetic testimony as a basis for our own judgments” (71—my italics).

If Meskin were right about the nature of the psychological phenomenon, his diagnosis of our resistance might not be implausible. He thinks that the folk conception of the content of aesthetic claims is mistaken. The folk theory, according to him, construes aesthetic content along subjectivist or relativist lines, whereas such content should be construed in an objectivist way. It is like a reverse error theory: in e.g. Mackie’s (1977) error theory concerning moral value, people are right about the content of their moral claims—right to think they make a claim to objective truth—but wrong in thinking any of the claims are actually true. Meskin’s idea is that people are wrong about the content of their aesthetic claims: wrong to think they *don’t* make a claim to objective truth, when in fact they do, and when in fact they can actually be true as well.

Presumably, however, not everyone is mistaken in this way. Meskin himself, for instance, is an aesthetic objectivist, and there are probably others out there too. Such people should not run afoul of the mistake. We would expect them to accept any epistemically respectable aesthetic testimony they come across, and to acquire aesthetic beliefs and make aesthetic claims accordingly. Moreover, Meskin suggests that even the folk who believe in aesthetic relativism should be expected to acquire such beliefs at least occasionally (in line with my discussion in the last section about versions of contextualism or relativism that trade in standards of taste); he says, “If A believes that aesthetic judgments have a hidden relativization (i.e., “That’s beautiful.” really expresses something whose form is closer to “That appears beautiful to people like me.”) then there is good reason for him to ignore the testimony of others, at least in most cases. Where

A does accept the testimony of others this will depend on his making an inference about his relation to the testifier.” (84)²³ We might think of the mistaken folk theory, in any case, as responsible for a *tendency* to resist aesthetic testimony that is not necessarily always active. So maybe sometimes believers in the folk theory would resist aesthetic testimony that there is no good reason to resist, but sometimes they would not.

In a later paper, Meskin (2007) gives some examples of the kinds of aesthetic beliefs he thinks we may acquire on the basis of testimony, some of which may also be warranted:

It is not that we *never* form aesthetic judgments on the basis of the testimony of others. For example, when we are confronted with the testimony of an art critic we know to be reliable, we may well be willing to make an evaluative aesthetic judgment on the basis of what they tell us. And it is not at all implausible that we sometimes form beliefs about the beauty of natural objects and landscapes on the basis of what we are told. That might explain some of our willingness to form vacation plans on the basis of testimony about natural beauty. Moreover, it is plausible that there are some cases (perhaps quite rare) when we are warranted in making aesthetic judgments on the basis of aesthetic testimony. That is, while testimonially-acquired aesthetic justification and knowledge may be rare, they do seem possible. (120-121)

Thinking that we sometimes acquire such beliefs, and that there is nothing particularly strange about doing so, fits well with his explanation of our surplus resistance: if his explanation is correct, we would expect that awareness of our mistake would lead to an end of it. If everybody came to accept the explanation, today’s low rate of acceptance of aesthetic testimony would simply increase to a level reflecting its real epistemic deserts.

Meskin’s examples of the kinds of aesthetic beliefs we may acquire through testimony match his estimation of the kinds of aesthetic testimony that are generally the most reliable epistemically: testimony from art critics who are knowledgeable about their field, and testimony about nature. Presumably, these beliefs are acquired through testimony by those who are not biased by their mistaken meta-aesthetic views, and who judge the testimony’s epistemic value

²³ Meskin in this paper doesn’t try to be very precise about what subjectivism or relativism involve. His example here seems to be a form of what I earlier called (indexical) contextualism, but it is not the kind of contextualism that I labeled ‘subjectivism’ in §2.2, which would incorporate the speaker’s aesthetic experience into the meaning of her aesthetic claim. His example is a form of contextualism that accepts *Experience Independence*, which is why it makes sense for A in the example to accept another’s testimony. It seems to me, though, that Meskin might have done better to attribute a kind of contextualism to the folk that, like my §2.2 subjectivism, rejects *Experience Independence*. He is trying to account for what he takes to be a surplus of resistance to aesthetic testimony: the portion of such resistance that is not justified by the testimony’s epistemic unreliability. We are to suppose that objectivists would not resist this portion of aesthetic testimony: that they would, justifiably, take such testimony to be correct. But if the kind of folk theory that Meskin alludes to in this example is responsible for the surplus resistance, we must suppose that the folk recognize like-minded aesthetic informants much less often than objectivists take aesthetic informants to be telling the aesthetic truth. If the folk recognize like-minded aesthetic informants at the same rate as objectivists recognize truth-telling aesthetic informants, there would be no surplus resistance. We need, then, to assume that the folk believe in far greater variation in aesthetic opinions than do objectivists, and are thus less willing to admit an informant to be like-minded than an objectivist is willing to admit she is right-minded. Perhaps this is not implausible: relativists presumably do focus on such variation to a greater extent than objectivists. In any case, Meskin’s earlier remark that the folk regard beauty claims as “little more than declarations of how objects seem to us, or how we experience them” suggests a folk view much more like my §2.2 subjectivism, which rejects *Experience Independence*.

accurately. These kinds of beliefs will also presumably come to be acquired in future by anyone who manages to rid herself of the mistaken meta-aesthetic views that she currently holds.

The problem with this view of the matter, however, is that it does not accurately track the psychological phenomenon, properly understood. If I have it on unimpeachable authority that the Grand Canyon is breathtaking, and I think Meskin is right in saying that it is perfectly legitimate for me to adopt this kind of view, then I ought to be able to do it, and to express my belief to others in the same unqualified way I might tell them that my friend's house, which I have never seen, is red. After all, Meskin thinks that the only reasonable bar to resisting aesthetic testimony is its epistemic unreliability, and this aesthetic claim, we are assuming, has as good an epistemic provenance as possible. But Meskin's suggestion does not dispel the sense that something has gone wrong if I describe myself, under such circumstances, as believing that the Grand Canyon is breathtaking; or if I tell others that this is so.

Meskin says that we sometimes acquire testimony-based aesthetic beliefs, and that sometimes these beliefs count as knowledge. I do in fact think he is right about this, as I will go on to explain. But he has not successfully dealt with the psychological phenomenon. The phenomenon at issue is that we do not describe ourselves as acquiring *unqualified* aesthetic beliefs on the basis of testimony, and we do not make unqualified testimony-based aesthetic claims, *even when* the testimony in question can be taken to be utterly reliable. So our response to aesthetic testimony is not quite as he depicts it; and I take it that it is moreover resistant to change of the kind that his suggestion concerning the influence of mistaken meta-aesthetic views would seem to mandate. We see the same problem with Mackie's error theory for moral claims: the theory essentially consigns this entire segment of our discourse to the scrapheap, but in practice moral language and moral thought are not given up so easily. This suggests that it is the theory, not the discourse, which is at fault. It is not of course impossible that a large section of our discourse could come to be seen as mistaken, and to be subject to change accordingly; this might in fact be an accurate description of what has happened to religious discourse, at least for some, in recent times. In the aesthetic case, however, I think that we can explain the practice in a way that shows why it is useful—why it makes sense. This of course would head off a diagnosis like Meskin's in terms of a commonly held mistake.

I suspect that Meskin's focus on the epistemology of aesthetic testimony makes him both underestimate the psychological phenomenon and overestimate the epistemological phenomenon. The latter overestimation is more evident in his later paper, where he barely mentions the mistaken meta-aesthetics explanation at all, and seems to intend for the epistemic unreliability of aesthetic testimony to bear more of the overall explanatory burden. If this is in fact what he intends, it makes his proposal there much the same as the one I considered and rejected in the last section. But if the epistemology of aesthetic testimony is going to bear the brunt of explaining the psychological phenomenon, it is not surprising that one would assume its epistemic status to be worse than it actually is, as the psychological phenomenon persists when the testimony's epistemic status is as good as possible. Meskin says above that the instances in which we make warranted aesthetic judgments on the basis of testimony are "perhaps quite rare." They are entirely rare, if what is meant by "making warranted aesthetic judgments on the basis of testimony" is "acquiring warranted unqualified aesthetic beliefs, or making warranted unqualified aesthetic claims, on the basis of testimony." But they are not rare at all (even if, to be sure, less common than for some other kinds of testimony) if what is meant is—in a way I

will explain more fully below—simply believing what one is told, and having one’s belief be warranted.

On the other hand, although Meskin seems inclined to exaggerate the unreliability of aesthetic testimony, he is also inclined to admit that it can’t *all* be epistemically unsound. But if the inferior epistemic status of aesthetic testimony is thought to be the principal, or perhaps the only, factor responsible for the psychological phenomenon, it is once again not surprising that one’s view of that phenomenon will be distorted: hence Meskin’s tentative suggestion above that for some of the most epistemically reliable aesthetic testimony, the phenomenon disappears. Once again, though, it is important to keep in mind what the psychological phenomenon actually involves. Meskin characterizes it as resistance to accepting aesthetic testimony, a resistance he assumes disappears for the most reliable testimony. But the genuine underlying phenomenon, on which the characterization of it as ‘resistance’ is based, is our unwillingness to admit to holding unqualified testimony-based aesthetic beliefs, or to make unqualified testimony-based aesthetic claims. *This* phenomenon does not disappear, in general, for the most reliable testimony. And, as I will go on to argue, it does not indicate at all that we are resistant to accepting aesthetic testimony.

2.4.2 Robert Hopkins and ‘Unusability’

Hopkins (2011) tries to make the best case for an aesthetic objectivist who, like Meskin—or, at least, the 2004 Meskin—accepts that there is no epistemic problem with much aesthetic testimony, but who, unlike Meskin, nevertheless thinks it usually cannot be a legitimate source of aesthetic belief. This makes such a person, in Hopkins’ terms, ‘pessimistic’ about aesthetic testimony; Meskin is an example of an ‘optimist,’ someone who thinks aesthetic testimony can be a legitimate source of aesthetic belief. Hopkins himself is officially neutral as to whether he is an optimist or a pessimist, but takes the theory he offers here to be the best option for pessimism.

Hopkins suggests—on behalf of the pessimist—that the phenomena in question can be explained by positing two different kinds of norm governing the legitimate formation of aesthetic belief: a norm of Availability and a norm of Use. The Availability norm comprises the epistemic standards that govern belief formation in general: it determines when belief counts as knowledge. According to Hopkins’ pessimist (and, it is clear, Hopkins himself), much aesthetic testimony does make knowledge ‘available,’ in the sense that it would meet the relevant epistemic standards for counting as knowledge, if the hearer would believe it on this basis. It is nevertheless illegitimate to acquire such a belief, according to the pessimist, because of some norm peculiar to the aesthetic case which makes such knowledge ‘unusable.’

Thus Hopkins’ pessimist, unlike Meskin, attributes the ‘psychological phenomenon’ to a norm of some kind, rather than to a mistake. However, the pessimist agrees with Meskin that the phenomenon in question is to be characterized as resistance to forming aesthetic beliefs on the basis of testimony. The norm in question, then—the norm of Use—is some norm or other that supposedly restricts the acquisition of aesthetic belief.

The challenge for Hopkins’ view is in showing how the two kinds of norm—Availability and Use—can possibly work together. Although he only speculates about what specifically the norm of Use might be in the aesthetic case, it is clear that it can’t be an epistemic norm—it can’t, that

is, determine whether an aesthetic belief counts as knowledge.²⁴ This is because it was posited to explain why it would be illegitimate to form aesthetic beliefs on the basis of testimony, even when such a belief, if adopted, would amount to knowledge. Given this, however, the combination of Availability and Unusability puts us in the awkward position of being offered knowledge that it would be illegitimate to take.

Hopkins is aware that this combination will strike his readers as problematic. He says:

One might wonder whether it is possible for belief to be governed by norms of the kind Unusability postulates. Belief aims at truth, and aspires to the status of knowledge. Epistemic norms determine whether it hits this target. How can belief be governed by nonepistemic norms, norms that govern something other than whether it counts as knowledge? If someone's testimony offers one knowledge, how can one fail to mold belief to fit? *Ex hypothesi* the claim being made is true and the situation is such as to put one in a position to know it. What room is there for a norm governing belief that requires one to refuse the knowledge thus offered? It seems that knowledge is not merely a goal of belief, but *the* goal. Epistemic norms are thus the only ones that should govern its formation. (145)

Nevertheless, Hopkins argues that it is perfectly coherent for belief to be governed by nonepistemic norms.

He first points out that we are not being asked to suppose that nonepistemic norms can govern the *formation* of belief: "That would be to form it in part without a view to its truth or its status as knowledge, and that may indeed be inconsistent with the goals constitutive of belief." (146) We are supposing only that nonepistemic norms may require us to *resist* belief. Belief aims at truth, and aspires to the status of knowledge, but perhaps in some cases it aspires to other things as well.

Hopkins acknowledges, though, that this still might seem unsatisfactory: "Isn't belief's role *nothing more* than to reflect the facts? If so, it can be governed *only* by epistemic norms. So one may remain tempted to think there is no room for norms of the kind Unusability views postulate." (146) His answer to this worry is to provide what he takes to be a clear example of a context in which an Unusability norm operates. If the norm clearly operates in this area, then we must allow that it could also operate in the aesthetic case.

The context Hopkins has in mind is that of expertise. Hopkins thinks that an expert, in whatever field, "is someone who ought to settle for herself questions in her domain of expertise" (146), and should resist testimony-based belief in that domain just as the pessimist claims we should resist testimony-based aesthetic belief. There will be many matters relevant to her field on which an expert will find testimony acceptable, but "for many topics, those forming the core of that field, she will not think it so." (147)

It is obviously crucial for Hopkins that his account of the norm governing the expert's belief-formation be accurate. When he says that the expert will accept testimony on many matters, but

²⁴ It may however be "epistemic in content" (149). For instance, Hopkins offers this candidate for the aesthetic Use norm: "Having the right to an aesthetic belief requires one to grasp the *aesthetic* grounds for it." This is epistemic in content, as what is required is a certain kind of knowledge. The norm is not epistemic, however, in the sense above, as the requirement does not bear on whether or not the aesthetic belief at issue counts as knowledge. Hopkins also offers another candidate norm of Use: "Having the right to an aesthetic belief requires one to have experienced for oneself the object it concerns." (150) This norm often appears in the literature, but it is always taken for granted, Hopkins says, that it is an epistemic norm. He wants to recast it as a possible norm of Use.

not on “the core” of her field, he presumably has in mind something like this: a physicist will accept testimony as to the exact publication date of Newton’s *Principia*, or even as to which naturally occurring uranium isotope is best used for fission in nuclear reactors, but will not accept testimony on fundamental theoretical principles. This seems plausible. There seem to me, though, to be two relevantly different kinds of fundamental theoretical principle. One kind, like Newton’s laws of motion for the physicist, are established principles of the field—in this case, even though they have for some time been regarded as merely good approximations of the movements of objects that aren’t very small, rather than as universally valid laws—and are taught to physicists long before they have become experts. The student physicist will, I think, accept testimony on such matters, although of course it is part of the physicist’s training to come to understand, as much as possible, whatever there is to know of the reasoning and evidence behind such theoretical beliefs. By the time the physicist is an expert, she no longer needs to acquire such knowledge, through testimony or in any other way. The other kind of theoretical principle might be the latest proposal in string theory, or on the nature of dark matter. I think it is right to say that an expert physicist won’t accept testimony as to the value of such cutting-edge theories. But I don’t think Hopkins has identified why. Such theories are controversial: the experts, even those who have considered certain theories deeply, disagree about their value. It is not surprising, then, that others in the field will not accept testimony on matters they believe to be controversial—not, at least, on those controversial matters about which equally respected experts disagree. When there is controversy of this kind, we don’t know who, if anyone, has knowledge to offer. The phenomenon is probably even starker for philosophy, the field of expertise which is surely dominating Hopkins’ thoughts, as there hardly exists a philosophical theory that is not controversial in this way.

I don’t think, then, that Hopkins has found a context in which knowledge has clearly been made available through testimony, but where an Unusability norm prevents it being accepted. There is in any case a basic structural problem with his view. There are many ways that knowledge can be made available to someone who is not in a position to take it, or who even chooses not to take it; but neither Hopkins’ expert, nor the typical recipient of aesthetic testimony, exemplify such a case. There are situations, for instance, in which knowledge is made available to a person who does not actually realize that this is the case. McDowell (2003) describes such a situation thus:

“I thought I was looking at your sweater under a kind of illumination that makes it impossible to tell what colours things are, so I thought it merely looked brown to me, but I now realize I was actually seeing that it was brown.” In saying something on these lines, one registers that one had, at the relevant past time, an entitlement that one did not then realize one had. One was in a position to acquire a bit of knowledge about the world, but because of a misapprehension about the circumstances, one did not avail oneself of the opportunity. One did not form the relevant belief, let alone get to know that that was how things were. (680-681)²⁵

A different kind of ‘misapprehension’ is in fact the basis of Meskin’s explanation of the phenomenon of aesthetic testimony: people don’t realize that others’ aesthetic testimony can be true in a way that is relevant to their own beliefs, and so they fail to avail themselves of the aesthetic knowledge that is sometimes on offer. I argued earlier, however, that many people

²⁵ Hannah Ginsborg alerted me to the existence of this passage.

can't be 'mistaken' about aesthetic testimony in this way, and that Meskin's hypothesis therefore doesn't adequately explain the phenomenon of aesthetic testimony. Sometimes we do think that others' aesthetic testimony makes knowledge available, when we take a reliable informant to be telling us what is aesthetically the case.

Hopkins remarks that this is also true for experts: he says that an expert must often be aware that a colleague's testimony makes knowledge available. For instance, he says that an expert may consider the judgment of her peers to be at least as good as her own, but that, "In such circumstances, she resists testimony, even though it offers her knowledge, and she takes it to do so. ... Since these folk [i.e. the experts] resist taking testimony even when they judge that it offers them knowledge, they cannot be responsive only to norms that determine whether their prospective belief counts as knowledge." (147) And later on the page: "The recipient [of testimony] need have no real doubt that her esteemed colleague's view amounts to knowledge, and yet feel she ought not to take her word on the matter."

There are of course many kinds of situation in which we can fail to acquire knowledge even when we know that it is available. Sometimes we simply choose not to avail ourselves of it. Maybe I can't be bothered to turn my computer on to look something up on Wikipedia; or the answer key to the exam is in the teacher's unlocked desk, but I am too fine and upstanding a student to take advantage of her momentary lapse. I may choose to forgo knowledge even when someone is trying to tell me something: "The murderer is ... !" my annoying friend tries to tell me, when I am in the middle of a whodunit, but I run out of the room with my hands over my ears, yelling, "La la la la la."

Sometimes, also, we can know that knowledge is available but are unable to avail ourselves of it, even if we want to. The internet connection is down. The answer key to the exam is locked up in the teacher's desk. Sometimes this is even the case when someone is trying to tell us something, and we want to know what it is. Maybe I ask a hotel employee in Madrid, in my proud but halting Spanish, how far the Prado is from the hotel, and she ignores my terrible accent and replies at a fast clip, "Está a unas cuatro manzanas de aquí." I am quite sure she has told me something true but I certainly do not now believe that the Prado is about four blocks away, because I haven't a clue what she actually said (in fact, I may end up with a hazy idea that the Prado must be on the Street of the Four Apples²⁶).

In a case like this, someone knows—or at least thinks—that knowledge is made available through another's testimony, because she thinks the testimony is true. She thinks the testimony is true, but nevertheless fails to believe the testimony, because she doesn't understand it. In the previous example, I believe that my friend's testimony as to the identity of the murderer in the novel is true, but fail to believe the testimony, because I can't hear it: this is simply a different way of failing to grasp it. Hopkins' expert, and the recipient of aesthetic testimony, are likewise in the position of thinking that knowledge is made available through testimony, because they think the testimony is true; but, because of the Unusability norm, they supposedly do not believe the testimony.

We have seen that this distinction, between believing that someone's testimony is true and believing the content of the testimony, is what Pettit hopes to rely on for his solution to the problem of aesthetic testimony. Pettit claims that we may think—even know—another's aesthetic testimony to be true, but cannot understand it to the extent that is necessary for

²⁶ My Spanish-speaking husband tells me that 'manzana,' which means 'apple,' also means 'city block' in Spanish.

counting as believing it. (He would presumably say, in his case, that testimony doesn't even make aesthetic knowledge available, as the testimony can't be used in any way to gain that understanding; it can only be gained from personal experience.) I argued that Pettit's attempt to establish his claim is unsuccessful, and it is in itself quite implausible. Hopkins does not explicitly make the same claim, and does not seem to have in mind at all that the recipient of either expert or aesthetic testimony fails to understand what is said to her. But in that case, it is hard to see how he can sufficiently motivate the separation between believing the testimony to be true and believing the content of the testimony. If I understand what someone tells me, and believe that what she tells me is true, then I believe what she tells me. There is, after all, simply no room for Hopkins' Unusability norm. It would be like a norm of etiquette that requires us to fail to see what is unpleasant, say. Etiquette can place restrictions on *looking*, and on talking about what we see, but I can't help seeing what I do when it is right in front of my open eyes.

Hopkins thinks there is nothing epistemically wrong with much aesthetic testimony, yet we resist it; this is what Unusability is supposed to explain. He thinks moreover that it is an advantage of his pessimist's view that it can also explain the tension between our resistance and the fact that we often accept aesthetic advice. The pessimist's explanation is that, in the context of advice-giving, the Unusability norm 'lapses' and our resistance disappears. This is supposed to be because the context is one in which the advice-taker cannot remain agnostic: I want to go see a film, but I can't see them all, and acquiring a belief about which film is most worth seeing is preferable to choosing one at random. So in this case I legitimately ignore Unusability and go ahead and acquire an aesthetic belief on the basis of a friend's or critic's testimony.

Even apart from the unworkability of the general framework Hopkins uses in this explanation, however, it relies on a misidentification of the kind of belief acquired when we take aesthetic advice. If I am told by a source I trust that *Citizen Kane*, which I've never seen, is a great film, I do not thereby take myself to believe that *Citizen Kane* is a great film, and I don't say to others, "*Citizen Kane* is a great film." We have been through this before. But the situation doesn't change in this respect if the trusted source is recommending that I actually go to see *Citizen Kane*, and I decide to take her advice: I still do not take myself to believe that *Citizen Kane* is a great film, until, perhaps, I have seen it myself; nor do I tell others that it is. The testimony's being advice rather than aimless chitchat doesn't change what kind of belief I acquire about *Citizen Kane*.

Hopkins sees an advantage in the apparent flexibility of his pessimist's view, but the pessimist's resolution of the tension between resistance and advice-taking depends on the same kind of distortion of the primary phenomenon that Meskin allows. The fact that we don't make unqualified testimony-based aesthetic claims, or admit to unqualified testimony-based aesthetic beliefs, is characterized as resistance to accepting aesthetic testimony. The Unusability norm is supposed to explain this resistance in the default case; the lapsing of the norm is supposed to explain our lack of resistance in the special case of being given aesthetic advice. But the underlying phenomenon is still present when we are taking advice, and this phenomenon is not indicative of resistance in the first place. Recognizing this allows us to avoid the problems of Hopkins' pessimist's view.

3. EVIDENTIALITY

3.1 Conveying evidentiality

Hopkins doesn't really talk about what came up towards the end of §2.3: that we have various ways to express the beliefs we do acquire from others' aesthetic claims: e.g. "Margot says the vase is beautiful", "Everyone says it's beautiful", "It's said to be beautiful", "The vase will be beautiful", "It's probably beautiful", "It must be beautiful." In fact, while such locutions are frequently brought up as examples of what it is legitimate to say instead of (e.g.) "The vase is beautiful," when we have no personal experience of its beauty, there has been little discussion of their possible significance.²⁷ It is perhaps assumed that the importance of their contrast with "It's beautiful" is merely epistemic—i.e. that their primary function is to communicate that the speaker has less confidence in the truth of "The vase is beautiful" than she would have if she had seen the vase herself and found it beautiful. This may be assumed even by those who insist that there is no general epistemic problem with aesthetic testimony: e.g. Rind, who agrees that there is no such problem, nevertheless describes Kant's position on aesthetic testimony by saying,

A description, or the testimony of others, may persuade me that a certain thing is beautiful, but I cannot legitimately express that persuasion by saying 'X is beautiful'. Rather, I must say something like 'By all accounts, X is beautiful'; or 'X must be beautiful'; or 'X is said to be beautiful'. I cannot make an *epistemically unqualified* declaration that the thing is beautiful until I have experienced ... the object for myself ... [my italics]. (2003: 66-67)

Of course, Rind's use of the italicized phrase may simply reflect the most common way to name the kinds of qualifications he has in mind (the use of epistemic modals, etc.), rather than indicating that he shares the assumption. But their status must be clarified. If such locutions are nothing more than epistemic qualifications, but there is no general epistemic problem with aesthetic testimony, and it is possible to have complete epistemic trust in another's aesthetic claim (as I think it is—remember the Grand Canyon, above), then we would expect it to be legitimate to omit any qualification in such cases. And it is not.

These locutions are not just epistemic qualifications, however. Let's focus on "X is said to be beautiful" (and the like) and "X must be beautiful." These alternatives to "X is beautiful" function to communicate, not just epistemic modality (information concerning the speaker's epistemic confidence in the claim), but, also, *evidentiality*: the speaker's evidential source for her claim.²⁸ "X is said to be beautiful" identifies the source of the speaker's claim as testimony; "X must be beautiful" identifies it as inference. "X is said to be beautiful" indicates testimony quite obviously; the evidential indication of "X must be beautiful" is less obvious, but it is nevertheless commonly acknowledged that use of epistemic *must* indicates a claim made on the basis of indirect inference.²⁹

'Evidential' and 'evidentiality' are terms from the linguistics literature. They refer to the linguistic mechanisms by which a speaker's source for her claim is conveyed: whether she

²⁷ I am not aware of any discussion of this in the aesthetics literature, where the problem with aesthetic testimony primarily lives. It has however recently become a subject of discussion in philosophy of language: see e.g. Ninan (2014).

²⁸ The possible relevance of evidentiality to the problem of aesthetic testimony was suggested to me by John MacFarlane (in an offhand comment he made during a seminar that was really about something else).

²⁹ See e.g. F. R. Palmer (1990).

asserts what she does on the basis of her own direct (sensory) experience, on the basis of hearsay, on the basis of indirect inference, etc. Some languages have a distinct grammatical category for conveying evidentiality, although these vary in how many kinds of information source must be evidentially marked. A. Y. Aikhenvald (2006) writes that some such languages carve the space into just two kinds of information—e.g. ‘firsthand’ vs. ‘nonfirsthand’ (where ‘firsthand’ may just be information acquired visually, or may include information from other senses too)—some carve the space into three, some into more (e.g. “Visual, Nonvisual sensory, Inferred, Reported,” as in Eastern Pomo from California).

Evidential information is often entangled with, and sometimes hard to distinguish from, information about the speaker’s degree of epistemic confidence. English locutions that indicate a testimony source for an aesthetic claim are noncommittal about the speaker’s degree of epistemic confidence in the claim, although context will often make it clear what the speaker thinks of the claim as a result of the testimonial evidence for it: for instance, if I say, “Margot says the painting is beautiful” and it is understood between us that I think Margot is a good judge of such things, then I will thereby convey that I have a high level of epistemic confidence in her claim. If I say, “The painting must be beautiful,” I convey that the source of my claim is an inference of some kind; I may think, for instance, that it must be beautiful because it was painted by Raphael. Aesthetic inferences may of course be based on many different kinds of reason, including what is considered to be good testimony: I may think, instead, that the painting must be beautiful because Margot says it is. And *must*, being a strong modal of epistemic necessity, indicates accordingly a high level of epistemic confidence.

3.2 ‘Must’

It has been noted, however, that use of epistemic *must* often appears to signal less confidence in a claim than utterance of the corresponding, unmodalized sentence, even though use of *must* makes a claim to epistemic necessity, and on standard treatment, *must* ϕ entails ϕ .³⁰ For instance, if, in answer to “Where’s Margot?” I reply, “She must be upstairs,” this may seem to indicate that I am less sure of Margot’s whereabouts than I would be if I had answered, “She’s upstairs.” Similarly, “The painting must be beautiful” may seem to indicate a lower degree of epistemic confidence than “The painting is beautiful.” Semanticists have often taken this appearance to show that *must* ϕ has a meaning that is weaker (i.e., commits to less) than the standard treatment predicts, and that, moreover, makes it weaker than ϕ .³¹ It seems plausible to me that this line of thought could be at least a small part of the reason why the problem with aesthetic testimony has been vaguely assumed to be epistemic: if all of the locutions that it is legitimate to use in the absence of personal experience, are thought to be epistemically weaker than “It’s beautiful”—including “It must be beautiful”—then it seems natural to assume that personal experience somehow has an epistemic leg up on any other method of acquiring aesthetic beliefs (or, at least, that the folk think it does), and that personal experience is the only kind of evidence that can provide sufficient justification for making an aesthetic claim. I therefore think it worthwhile to spend some time on *must* in order to clarify its bearing on aesthetic claims. I will conclude—with von Fintel and Gillies (2010), whose paper in defense of *must*’s strength seems to spearhead this side of the debate—that *must* is indeed strong: *must* ϕ entails ϕ . This means that someone

³⁰ See e.g. Lauri Karttunen (1972: 11-12).

³¹ See e.g. von Fintel, Kai & Anthony S. Gillies (2010).

who says “The painting must be beautiful” is (at least) as committed to the painting’s beauty as someone who says “The painting is beautiful.”

Von Fintel and Gillies offer a number of arguments in defense of *must*’s strength. They note first that *must* does not *always* appear to be weak. Taking their example: if I know the ball is in the drawer or the cupboard, see it is not in the drawer, and assert, “It must be in the cupboard,” there is nothing about my assertion that lacks confidence relative to having said “It’s in the cupboard” instead. I have already mentioned an aesthetic example of this kind: “It must be simply breathtaking,” I say, in response to my friend’s vivid description of the Grand Canyon. The use of *must* here does not signal any lack of epistemic confidence at all in the Grand Canyon’s breathtakingness. Von Fintel and Gillies suggest that it would be possible at this point to propose that *must* is sometimes strong, sometimes weak, and that its semantics should reflect this. But they offer a number of further arguments to show that, semantically, it is always strong. For instance, they say that if *must* ϕ were weaker than ϕ for some ϕ , then “*must* ϕ but perhaps not ϕ ” (e.g. “It must be beautiful but perhaps it isn’t beautiful”) would have to be a reasonable claim to make in that case, but it couldn’t be: in fact, it looks like a contradiction. They also say that there are perfectly good ways to express a relative lack of confidence, by using weaker epistemic modals—“It’s almost certainly in the cupboard”, “It’s probably in the cupboard”, “It might be in the cupboard” (“It’s probably beautiful”, “It might be beautiful”)—and say that a speaker is never indifferent to a choice between *must* and one of the weaker modals.

Von Fintel and Gillies explain the apparent weakness that *must* exhibits in some cases, in this way:

Of course, there are prejacent [the bare, unmodalized sentences] for which intuitively direct evidence is more convincing evidence than indirect inferential evidence. So, a speaker who chooses nevertheless to use the strong *must* ϕ incurs a higher degree of risk. So, we may judge that in many cases, *must* ϕ is more likely to be false than ϕ by itself would have been if there had been direct evidence for the prejacent. But a sentence being more likely to be false than another is far from an argument that it is *weaker!*” (380)

I will conclude in a moment that von Fintel and Gillies are essentially right about this. However, their discussion is incomplete. They present the problem, as they say others have done before them, simply by rehearsing specific examples of unmodalized sentences that appear to convey more confidence when asserted than their *must*-ified counterparts. They say, for instance, that when considering the following dialogue—

Where are the keys?

- a. They are in the kitchen drawer.
- b. They must be in the kitchen drawer.

— “it is natural to feel that the strongly modalized claim is less forceful.” (352) Yet there is another kind of example, which von Fintel and Gillies do not discuss, that constitutes a much greater challenge to a strong *must*. Daniel Lassiter (2014), who argues against von Fintel and Gillies in favor of the relative weakness of *must*, offers several such examples from the genealogy website Ancestry.com. Here's one:

My great grandfather John HARTMAN married a lady by the name of Margaret KESSLER. John died in Colorado 1 Oct 1896. His wife was not in his will so must have died before that time. ... Does anybody know anything about this couple? (598)

This seems to be a straightforward use of *must*, not at all jokey or exaggerated in any way. Moreover, it seems to exemplify a kind of use that is very common. I agree with Lassiter, however, that the author of the post does not seem to be “maximally committed to the truth” of the conclusion that Margaret died before John, if what this means is that the author is certain that she did. I think in this case that the author cannot be said to know that Margaret died before John, and that it would have been odd simply to have asserted this (“His wife was not in his will so she died before that time”). Yet there is nothing at all odd about the *must*-ified assertion that the author does make.

On the standard account, epistemic *must* is strong for the following reason: *must* ϕ is true just in case ϕ is true at all worlds compatible with what is known. As what is known is true, the actual world is one of these worlds; so if *must* ϕ is true, then ϕ is true. It would seem, then, that it ought to be easy for anyone justified in asserting *must* ϕ to be justified in asserting ϕ , simply by making an obvious deductive inference. Yet Lassiter’s example is one in which the author of the post seems justified in asserting *must* ϕ but not justified in asserting ϕ .

A problem similar to this is of course well known in epistemology. I take myself to know that Margot is upstairs; Margot’s being upstairs entails that she didn’t climb out the window and shin down the drainpipe; so I should straightaway be able to acquire the knowledge that Margot didn’t climb out the window and shin down the drainpipe. Yet if the subject of Margot and shinning and drainpipes comes up, I cannot take myself to rule out the possibility that that is exactly what she did, and I therefore hesitate to state that I know she didn’t. I.e.: I am prepared to assert p but not prepared to assert q , even though, if the subject comes up, I am aware that p entails q .

A common response to this latter problem is to say that the justificatory standards for knowing are not everywhere the same—they are dependent on context.³² So I take myself to know that Margot is upstairs on ordinary occasions, when I can rule out her being downstairs, and having left the house through the front door, etc. If I have not on such occasions ruled out her having climbed out the window and shinned down the drainpipe, that is no matter. If the subject comes up, however—as it might if Margot has turned to a life of crime and the cops are at the door—and I cannot in this new, worrisome context rule out her drainpipe exit, then I don’t take myself to know that she’s upstairs. In fact, simply being aware that p entails q is often to be in a context in which the subject of q has already come up. If I cannot rule out $\neg q$ in such a case, then I do not, in this case, know that p . But, ordinarily, I may know that p , and p may entail q , yet I may not have ruled out $\neg q$, because the question whether q is simply not relevant in the context at issue.

Given epistemic *must*’s connection with knowledge, we would expect it to exhibit similar variability in the justificatory standards that are applicable in different contexts of its use—and it does. On ordinary occasions, if I am asked where Margot is, I may say, “She must be upstairs,” even if I haven’t ruled out her shinning down the drainpipe. When the cops are at the door, though, the newly relevant possibilities may make me hesitate to sound so sure. However, the

³² I am steering clear here of any specific theories of this kind: contextualist, relativist, or what have you.

response to the epistemological problem that relies on contextually variable justificatory standards is not available as a response to the problem with *must*. This is because the question whether ϕ is *always* under discussion when the question of *must* ϕ is under discussion. Given the standard account of epistemic *must*, it seems that a speaker who has sufficient justification for assertion of *must* ϕ should have knowledge that rules out $\neg\phi$. But in Lassiter's example, the author of the post would presumably admit that her knowledge does not rule out Margaret's having been alive at the time of her husband's death; and, in fact, the author presumably would not flat-out assert that Margaret died beforehand. So it seems to be an example in which the possibility of $\neg\phi$ is simultaneously ruled out and not ruled out.

As I have said, von Fintel and Gillies do not discuss examples of this kind. They note that, given the indirect evidential signal that *must* conveys, asserting *must* ϕ is inappropriate when the speaker's evidence is in fact direct: if I am standing in the pouring rain, saying "It must be raining" can at best be a joke (as I might say, "Hm—slight chance of rain today, then"). But they draw exclusively on examples in which, if *must* ϕ is a felicitous assertion, because one has indirect evidence for ϕ , then assertion of ϕ is felicitous too: for instance, I may say "It must be raining" on seeing people's wet raingear when they come into the building, but I can just as appropriately say, "It's raining." In this example, the speaker's indirect evidence is in fact good enough for knowledge of the rain, so both assertions are felicitous.

Weakening *must* ϕ relative to ϕ clearly and easily fixes the problem. Von Fintel and Gillies discuss a couple of examples of theories that do just that. Angelika Kratzer (1991), for instance, does not make *must* a universal quantifier over the entire set of epistemically possible worlds. She introduces a further constraint, such that *must* ϕ is true just in case ϕ is true in all of a certain subset of the epistemically possible worlds: those worlds which "come closest to the normal course of events" (645). As things in the actual world may turn out not to be very normal, the actual world may not be in this subset; so, on this account, *must* ϕ does not entail ϕ . We can thus be justified in asserting *must* ϕ —because we have good reason to think that ϕ will be true if things turn out to be normal—but not ϕ , in those contexts in which knowing that ϕ requires ruling out abnormal possibilities.

Lassiter weakens *must* ϕ relative to ϕ in a slightly different way. He says that *must* ϕ is true just in case ϕ is highly probable (on the basis of indirect evidence: he treats the evidential component of *must* similarly to von Fintel and Gillies, as a presupposition)—where the exact degree of high probability that is required is presumably fixed by context. Again, we can be justified in asserting *must* ϕ when we have good reason to think, on the basis of indirect evidence, that ϕ is highly probable, without being in a position to claim that we know that ϕ .

Yet von Fintel and Gillies' arguments for a strong *must* are for the most part compelling—although, as I will argue, they must make room for the kinds of exceptions illustrated by Lassiter's Ancestry.com and other examples. I will focus on their argument concerning how *must* interacts with weaker modals, which seems to me to be the hardest for a weak theorist to entirely circumvent. They say that if the weak theorists are right and *must* ϕ does not entail ϕ , then *must* ϕ but perhaps not ϕ should be a perfectly acceptable thing to say, at least sometimes (e.g. "It must be raining but perhaps it isn't raining"). But, they say, it never is: it sounds like a contradiction, no matter what order the conjuncts are in. So it really does look like *must* ϕ universally quantifies over all the epistemic possibilities.

But of course this is not the last word on the matter. As Lassiter notes, “Weak theories, as a rule, deny vFG’s assumption that *perhaps* (*might*, *maybe*, *possible*, etc.) expresses existential quantification over [the epistemic possibilities]. Instead they treat *perhaps* as *must*’s dual: *perhaps* $\phi \equiv \neg$ *must* $\neg\phi$.” (602)³³ For such theories, *perhaps* is stronger than a pure existential quantifier. Theories of this kind would of course agree that *must* ϕ but *perhaps* not ϕ is contradictory.

The weak theorists’ position here seems to me to be somewhat awkward: we are to suppose, I take it, modulo further machinations, that there are contexts in which *must* ϕ , *impossible* $\neg\phi$, and $\neg\phi$ are all true. Von Fintel and Gillies acknowledge the point concerning weak theories and the duality of *must/perhaps* in a footnote, and respond: “[W]e’re not wedded to a pure existential story about *perhaps* (though that sounds pretty good) and our argument doesn’t rely on it. There are strong necessity epistemic modals. So pick one and take its dual (e.g., *there is a vanishingly small chance that*). It’ll be horrible when paired with *must*” in examples like the one above: e.g. “It must be raining but there’s a vanishingly small chance that it isn’t raining” (fn. 25).

However, Lassiter claims to find naturally occurring, felicitous examples of just such pairings; e.g. this one, from some online source:

I refuse to believe that this one game, Lost Planet 2 DX11, which was previously 100% stable remember, is crashing because my overclock is unstable It’s not impossible, granted, but IMO it is highly unlikely. There *must* be some other cause. (603)

I agree with Lassiter that this is a natural thing to say. Note, though, that the example actually uses the more common modal language (“It’s not impossible”) that is originally at issue in von Fintel and Gillies’ argument, rather than the kind of language that they advert to later, in exasperation, as a response to the weak theorists’ position (“there is a vanishingly small chance”). So we might have expected, given what Lassiter has just said about that position—i.e. that *perhaps* (*might*, *maybe*, *possible*, etc.) are *must*’s duals, so that e.g. *perhaps* $\phi \equiv \neg$ *must* $\neg\phi$ —that Lassiter and other weak theorists would have just as much work to do to explain this example as do von Fintel and Gillies. After all, Lassiter has just appeared to say (allowing for suitable rearrangement of negations) that *impossible* ϕ and *must* $\neg\phi$ are contradictory, and it is this conjunction that is at work in the Lost Planet example (“It’s not impossible” together with “There must be some other cause”). Instead of conceding this, however, Lassiter says instead that such examples “place constraints on the meanings of the other epistemic expressions [used in them] ... For example, to account for [the Lost Planet example] we must assume (quite plausibly) that *phi is impossible* requires a higher degree of confidence in $\neg\phi$ than *must* $\neg\phi$ does.” (fn. 3)³⁴ That is, Lassiter is claiming, I take it, that *impossible* ϕ and *must* $\neg\phi$ are not contradictory after all. However, I am not at all sure, in that case, whether he thinks that *possible* ϕ is not semantically equivalent to *impossible* ϕ , or whether *possible* should be removed from the list of modals that are duals to *must*, or whether it is something else he has in mind.³⁵

³³ Symbols changed for consistency with my text.

³⁴ Symbols once again changed for consistency with the text. Note that Lassiter is clearly making a point about the *semantics* of the relevant modal expressions, not just some of their possible pragmatic effects.

³⁵ There are of course well-known *pragmatic* differences between e.g. *possible* ϕ and *not impossible* ϕ . Someone who claims the latter usually implicates something along the lines of, “You may have thought ϕ impossible (but it isn’t).”

In any case, Lassiter himself mentions a more natural solution to the apparent contradiction in the Lost Planet example: he says, “It has been suggested to me that there is some kind of modal domain shifting going on in the midst of these examples.” (fn. 4) This seems to me to be exactly what’s going on: the effect of the last sentence in the example above (“There must be some other cause”) is to *dismiss* the possibility raised in the previous sentence (“It’s not impossible”). That is, the epistemic possibility that is said to obtain in the first sentence is not a member of the new set of epistemic possibilities relevant to the second sentence. The simplest explanation would be that the author has changed her mind between the two sentences—i.e., decided that the possibility raised in the first sentence is not a possibility after all. But that is not what I think is going on here. If, for instance, someone were to respond to the post by saying, “Why? You just said it *could* be an unstable overclock,” it would at least make sense for the author to repeat, “Well, yes—that’s a remote possibility. But I think it’s got to be something else”—rather than saying, “I was wrong about that—it can’t be the overclock (because …).” So we need to make sense of the author thinking, and asserting, both *possible* ϕ and *must* $\neg\phi$ —if not exactly simultaneously, then at least side by side.

At bottom, Lassiter’s Lost Planet example is simply a more explicit version of the Ancestry.com example we started with. In both examples, the authors’ modal assertions, and other features of the context, seem to indicate—assuming *must* is strong—that they both know, and do not know, the ϕ in question. In the Ancestry.com example, the author tells us what evidence she uses to draw her conclusion: Margaret must have died before John because she was not in his will. We infer from this, given our understanding that the evidence here is far from conclusive, and that the author must know this too, that the author does not know, and does not take herself to know, that Margaret died before John. Yet, in saying that Margaret must have died before John, the author seems to indicate (assuming a strong *must*) that she does know this.

In the Lost Planet example, the author seems to express two conflicting knowledge states by her explicit use of the two modal sentences. In saying that it’s not impossible that the game is crashing because of an unstable overclock, she seems to indicate that she cannot rule out this cause of the problem. In saying that the cause must be something else, however, she seems to indicate that she can rule it out.

Lassiter responds to the suggestion of modal domain shifting in the Lost Planet example by saying, “There is no way to exclude this possibility, but without detailed constraints on the inclusion and exclusion of worlds from E [the set of relevant epistemic possibilities] such a theory borders on vacuity” (fn. 4). I take the point that we can’t trot out modal domain shifting as a remedy for any apparent contradiction we feel like sweeping under the carpet—but the lack of explicitly detailed constraints of the kind Lassiter is asking for doesn’t make it any less likely that this example is subject to the phenomenon at issue. Such an explanation gains traction, moreover, by noting that the juxtaposed modals in the example (essentially “X is not impossible ... but there *must* be some other cause”) only sound natural under certain conditions: they must be juxtaposed in that order, with the right kind of pacing and the right kind of intonation. There must be a significant pause between the modals—and in this case, the shift between them is further signaled by the author’s saying that the possibility in question is “highly unlikely” (on which more below)—and the sense of evaluation given by emphasis of *must* (an emphasis which, on the page, is given by italics). (Compare: “It’s not impossible that Margaret outlived John but she must have died before him,” delivered in a straightforward, unemphasized, matter-of-fact tone.) These things help the listener, or the reader, understand that what is going on here is the

progression of a single thought, from consideration of various possibilities to a decision that results in the rejection of one of them. Reversing the modals—“It can’t be X, but X is not impossible”—starts with the decision and then instantly reneges on it, which can at best be interpreted as a change of mind mid-sentence and which, given that the about-turn is made without warning, is correspondingly infelicitous.

Note, though, that the possibility that is initially raised is not necessarily rejected *as* a possibility—i.e., as I said earlier, the author need not have changed her mind about that—but it is rejected as the *answer* to the author’s question about the cause of her game’s failure. Her use of *must* indicates a certain kind of commitment on her part to what that answer is.³⁶ The question I will finally get around to answering, now, is why, if *must* is strong, it is an appropriate vehicle for expressing that commitment, when the possibilities it rules out are in some sense still in play.

Here is one possible explanation of the Lost Planet example, which plausibly demonstrates a common exaggerated use of *must*. I am not committed to it, and will offer another below. The author sounds desperate (“I refuse to believe . . .”) and in her desperation may insist on an answer to her problem that is perhaps insufficiently grounded epistemically. In effect, she seems to be overstating her position: her frustrated “There *must* be some other cause” may in a certain way exaggerate the state of her knowledge. Von Fintel and Gillies remark thus on this phenomenon: “[T]he strength of *must* claims may be exploited in a “the lady protests too much” kind of way Just as saying *I have no doubt that he will be here very soon* may sometimes indicate the presence of at least a smidgen of doubt, saying *It must be raining* may indicate the presence of tentativeness.” (380) As they go on to say, however, “But just as it would be insane to attribute this kind of weakness to the semantics of *I have no doubt*, we have argued that no weakness should be built into the semantics of *must*.” And the Lost Planet example is equally amenable to an interpretation that leaves *must* strong in spite of the author’s possible doubt on the matter, and her acknowledgement, in a cooler moment, of another possibility. The case, if described correctly, is one in which the author, in saying *must* $\neg\phi$, *pretends* to more knowledge than she actually has. The pretense is benign, however, as context makes it clear that she is in fact “protesting too much,” and her readers will not be misled.

“Protesting too much” as an explanation does not seem feasible, however, for the Ancestry.com example, which comes across in a straightforward way free of hyperbole. Remember that the challenge (or, at least, one way of putting the challenge) is to explain why the author is prepared to assert *must* ϕ (“Margaret must have died before John”) but would not be prepared to assert ϕ (“Margaret died before John”). Von Fintel and Gillies attribute the apparent weakness of *must* to its evidential component: an assertion of *must* ϕ signals that it is made on the basis of indirect inference, which is often a less convincing kind of evidence than direct evidence or reliable testimony. Thus a speaker who asserts *must* ϕ often incurs a higher degree of risk that her assertion is false, compared to someone who asserts ϕ on the basis of direct evidence or testimony. In this particular example, it is clear that the author of the post chooses to incur a very high level of risk: for all she knows, it could easily be the case that Margaret outlived John. The question here is why she chooses to incur the high degree of risk associated with assertion of *must* ϕ , when she would not be prepared to incur the risk of asserting ϕ .

³⁶ Lassiter notes that *must* is often used to communicate the results of inference to the best explanation. Making such inferences seems to me doxastically akin to confident betting: we can make such a bet while knowing full well we are not certain to win. (We can also, of course, think something highly probable without being willing to bet on it.)

The answer is surely that *must* ϕ 's evidential component *licenses* greater risk-taking in assertion than is appropriate for assertion of ϕ . The author is presumably prepared to assert that Margaret must have died before John—even though she doesn't know that Margaret died before John, and would not be prepared to assert *that*—because such an assertion makes it clear, through its evidential component, and our knowledge of the relevant evidence, that her grounds for her assertion aren't the best grounds: the reader is prepared to expect the particular way in which the author might be wrong. The author in this case does not have good enough grounds for assertion of ϕ , but she does have good enough grounds for assertion of *must* ϕ , even though—evidentiality aside—*must* ϕ entails ϕ . She still may not assert ϕ , as doing so would mislead her readers as to the quality of her evidence for ϕ .

In effect, *must*'s evidential component, in cases like this, *lowers* the epistemic standards necessary for assertion of *must* ϕ , relative to those necessary for assertion of ϕ , even though *must* ϕ is, semantically, a stronger proposition. Note that, if it is not clear from the context in which *must* ϕ is asserted, that the indirect evidence behind the assertion is insufficient for knowledge of ϕ , then the epistemic standards for assertion of *must* ϕ in that case are not lowered in this way. To hijack (and reroute) an example from von Fintel and Gillies:

The Hollywood crew has been hard at work setting up their rain equipment for shooting the big scene. You see people coming inside with wet umbrellas, slickers, and galoshes. You are pretty sure that rain is the only explanation—filming on the big scene doesn't start until tomorrow—but there's a twinge of doubt. What do you say? Here are some options:

- a. It is raining.
- b. It must be raining.
- c. It is probably raining. (366)

Von Fintel and Gillies point out that, in this case, (b) is not a good choice; and, of course, (a) is not a good choice either: the alternative possibility at issue is strong enough to preclude you from taking yourself to know, and hence to assert, that it is raining.

So far, so good; however, the example is somewhat obscured—at least, for my purposes—by a couple of factors that I want to clear out of the way. The first is the fact that you are described as having a “twinge of doubt”: this suggests the kind of active doubt that someone who uses *must* generally does not feel (cf. my fn. 36)—as opposed to the kind of mere awareness of alternative possibilities, not necessarily ruled out, that is perfectly common with use of *must*. But we may remove this complication from the example:

... You see people coming inside with wet umbrellas, slickers, and galoshes. You are pretty sure that rain is the only explanation—filming on the big scene isn't supposed to start until tomorrow—but it's possible that filming started early. What do you say? ...

“It must be raining” still feels like an awkward choice, although it is, I think, somewhat better than before. The second complicating factor is related to the point I made earlier about the ordering of modals in the progression of a train of thought. We can make (b) for this example an even better choice by telling the story of your thought in way that leads up to *must* rather than away from it:

... You see people coming inside with wet umbrellas, slickers, and galoshes. Filming on the big scene isn't supposed to start until tomorrow, but it's possible that it started early. Nevertheless, you are pretty sure that rain is the only explanation. What do you say? ...

(b) is still somewhat awkward, however—even though, as we have seen, *must* is often used when one's indirect evidence is not considered to be conclusive. The remaining difference, I think, between this example and previous examples, is that the inconclusive nature of your evidence is not made clear by your simply saying, "It must be raining." The example stands in contrast to a slew of other rain examples in von Fintel and Gillies' paper, and we have been led at this point to see that both "It must be raining" and "It's raining" are commonplace observations on ordinary occasions of seeing people traipse inside with wet umbrellas, slickers, and galoshes. That is, in normal circumstances, wet umbrellas etc. are indirect, though generally conclusive, evidence that it is raining. What throws (b)'s game off in this particular example, then, is that the speaker does not in fact have the conclusive evidence (b) usually indicates, and, moreover, the nature of the speaker's inference concerning the unlikelihood of film-equipment-generated fake rain is not made clear either in the conversational background or the actual conversation. Altering the example appropriately in either one of these ways will make (b) clearly acceptable—even though your epistemic state remains in relevant respects the same:

1. The Hollywood crew has been hard at work setting up their rain equipment for shooting the big scene. You and your colleague have been consulting the calendar for the shoot, which says that filming on the big scene doesn't start until tomorrow—but you have both just noticed that the production manager has made a mistake in the schedule for the following week. Just then, you see people coming inside with wet umbrellas, slickers, and galoshes. Your colleague gives you a look. "Surely not," you say.³⁷ "It must be raining."
2. The Hollywood crew has been hard at work setting up their rain equipment for shooting the big scene. You have just been consulting the calendar for the shoot, and noticed that the production manager made a mistake in the schedule for the following week. Now you see people coming inside with wet umbrellas, slickers, and galoshes. "The calendar says that filming on the big scene doesn't start until tomorrow," you say. "I know scheduling errors are sometimes made, but they don't happen very often. It must be raining."³⁸

As for the Lost Planet example: whether or not "protesting too much" is the right explanation in that particular case, the author there could easily have said the same thing without the desperation involved in protesting too much ("It's not impossible, but it's highly unlikely. There must be some other cause"). Much the same pattern of assertion is made in the rain example (2) above. So we should apply to such cases the explanation I have just outlined, which says that

³⁷ Incidentally, "Surely not" clearly means "(That is) certainly not (the case)," but is often used to convey something like, "I think/hope that isn't the case (because, surprisingly, it may be)."

³⁸ Note that embedding *must* claims in *knows* clauses also prevents the lowering of epistemic standards that I have been discussing. In both (1) and (2) here, it would not have been acceptable for you to say, "I know it must be raining," instead of "It must be raining." Or compare: you bought a ticket for a lottery with a million tickets sold. They have just drawn the winning ticket, but we don't know what it is. "You must have lost," I say. "You had a one in a million shot." But I could not have said, "I know you must have lost." You would have expected me, in that case, to back that up with some such reason as, "It was rigged from the start." The purpose of embedding a *must* claim in a *knows* clause seems to be precisely to emphasize that one's indirect evidence is conclusive.

must's evidential component may lower the epistemic standards necessary for assertion of *must* ϕ , relative to those necessary for assertion of ϕ .

These are cases in which assertion of *must* ϕ immediately follows assertion of *possibly* $\neg\phi$, but where the speaker does not seem to have changed her mind in between. We may take the assertion involving *possibly* to be made according to whatever epistemic standards are appropriate in the given context for claiming knowledge of whether or not ϕ . If the speaker asserts *possibly* $\neg\phi$, then we can suppose that she does not take herself to know that ϕ . Note that in both the Lost Planet and the Rain (2) examples, the subsequent *must*-ified assertion comes after a claim that the possibility just raised is "highly unlikely" (Lost Planet) or "doesn't happen very often" (Rain (2)). This claim is what allows the subsequent assertion of *must* ϕ to be made: we are thereby given the speaker's evidential basis for her assertion of *must* ϕ , which she makes according to epistemic standards that are lower than those that were in play for her assertion of *possibly* $\neg\phi$. She asserts *must* ϕ as a result of an inference she has made concerning the low probability of an alternative that she cannot, however, entirely rule out. We are aware of her reasoning, which is why her assertion of *must* ϕ , with relatively low epistemic warrant, is nevertheless acceptable. If the speaker omits the "highly unlikely"/"doesn't happen very often", then it is possible to convey the same information tacitly by suitable pacing and tone (... and possible eye-rolling). Use of 'but' in such cases is probably indispensable ("I know scheduling errors are sometimes made but it *must* be raining!"). I assume this is easier in speech than on the page.

Towards the beginning of this discussion I mentioned two of the arguments that von Fintel and Gillies put forward for a strong *must*. It is now time to sum up their effectiveness in the face of the examples I have been discussing. Lassiter's Lost Planet example seems to be a counterexample to the first of the arguments: i.e. von Fintel and Gillies claim that the conjunction of *must* ϕ with any weaker modal claim such as *perhaps* $\neg\phi$ yields a contradiction, no matter the order of the conjuncts; but the Lost Planet example does in fact exhibit such a conjunction, and it seems entirely felicitous. I have argued, however, that the example involves a particular kind of domain shift between the two modals, which depends for its effectiveness on the evidential component of *must* together with various other factors, including the order of the conjuncts. Where no such domain shift is made evident in the conversation, the conjunction of *must* ϕ with (e.g.) *perhaps* $\neg\phi$ will come off as contradictory, as von Fintel and Gillies claim. This is also the case (as they emphasize) for the weakest possibility modal that one could dream up, such as *there is a vanishingly small chance that* $\neg\phi$. This general feature of *must* argues hard for its strength.

The other argument that I mentioned is more equivocal. It concerns the Hollywood rain example I have just been discussing. Von Fintel and Gillies use the (original) example to try to show that "when a smidgen of weakness is called for, speakers don't regularly reach for *must*" (366). As I have indicated, I think they are right that *must* is not used to *express* doubt. This is why the "twinge of doubt" in the example (among other things) makes choice (b)—"It must be raining"—seem an inappropriate choice. But, as I have been arguing, *must* is often used when a kind of epistemic openness is nevertheless *present*—sometimes strongly enough to prevent assertion of the corresponding bare prejacent. So I do not agree with a certain generalization of the example that von Fintel and Gillies perhaps hold to be true. They say, "If the content of *must* really were relatively weak, we would expect something close to indifference between the *must*-claim [(b)] and its *probably* counterpart [(c)]. But that is not what we see: the latter is clearly

preferred to the former.” (367) They are talking about this particular example, but the overall tenor of their argument tends to suggest that they think *must* ϕ and *probably* ϕ are never interchangeable, and that is not true. In the Ancestry.com example, for instance, the author could easily have substituted *probably* for *must*: “His wife was not in his will so probably died before that time.” Of course, she would have made a weaker claim in that case, but it would have made equally good sense of her epistemic position. The practical choice between them concerns what the speaker wants to convey about where, if anywhere, she is willing to put her money. In choosing to stake that money, *must*’s strength commits her to Margaret Kessler’s earlier death, and she will lose her bet if it is discovered that Margaret in fact outlived her husband. That is, use of *must* always involves a strong assertoric commitment, even when that commitment carries a high level of risk.

3.3 ‘Must,’ aesthetic claims and evidentiality

I initiated this discussion of *must* in order to clarify the status of aesthetic claims using it. *Must*, semantically, is strong: *must* ϕ entails ϕ ; and it indicates an indirect evidential source. It might have seemed, then, that anyone who felicitously asserts *must* ϕ —even for aesthetic ϕ —would have indirect, yet sufficient, evidence for ϕ . But the pattern of usage for *must* that I have just been discussing complicates the picture: *must* is often used when the speaker has *less* warrant than would be needed for assertion of ϕ . This suggests the possibility that for aesthetic claims, *must* ϕ is in fact *always* used in this way, and that indirect evidence is, in fact, never sufficient for knowledge and assertion of the bare aesthetic prejacent. Given the original puzzling datum—that simple, unmodalized aesthetic claims are never felicitously made in the absence of direct, personal experience—this possibility is a compelling one. There is, after all, this crucial difference between the aesthetic case and the non-aesthetic examples that have come up: it is never acceptable to say e.g. “It’s beautiful” instead of “It must be beautiful,” as the first requires direct experience and the second requires indirect evidence; but, in non-aesthetic cases, where there is indirect yet conclusive evidence for ϕ , assertion of *must* ϕ and assertion of ϕ are equally acceptable.

I have, however, already dismissed the epistemic solution to the problem with aesthetic testimony. Such testimony *can* provide good enough warrant for knowledge of aesthetic ϕ . So the significance of the contrast between *must* ϕ and ϕ , for aesthetic ϕ —and the significance of the contrast between the aesthetic case and non-aesthetic cases—must lie elsewhere.

Simple, unqualified assertions (“It’s raining”) are typically evidentially neutral, as we have seen. Unqualified aesthetic claims, however, as we have been trying to explain throughout, are not: if I say, “The painting is beautiful,” I convey that I make the claim on the basis of my own direct experience—I have seen the painting and found pleasure in it—and so this is not something I can felicitously say if I haven’t seen the painting. At most, I can say, “The painting must be beautiful”—as I might if someone whose judgment I trust highly has told me that it is.

I have gone through several attempts to explain this phenomenon that I believe to be unsatisfactory. It is of course not impossible for a theory that denies *Experience Independence* to try to account for the contrast between “It’s beautiful” and “It must be beautiful.” But we have seen that denying *Experience Independence* incurs semantic problems that make it overwhelmingly attractive to take aesthetic claims to have propositional content that attributes a property to an object that is independent of anyone’s actual experience. We have also seen,

however, that accepting *Experience Independence* can make the phenomena of aesthetic testimony seem inexplicable.

I want to suggest a relatively simple solution. While there are many ways to express different kinds of evidentiality in English (we have been considering *it is said that ϕ* and *must ϕ* ; there are also such locutions as *I see that ϕ* , *I hear that ϕ* , *it seems that ϕ* , *allegedly ϕ* , etc.), evidentiality does not form a distinct grammatical category in English, and, as I have said, it is not generally obligatory to include evidential information in an English assertion. My suggestion is simply that aesthetic claims (together with other evaluative claims, certain sensory claims, and some others—see below) form a class for which evidential information is, in effect, obligatory, because there is no easy, idiomatic way to make a claim in this class that does not convey it. Asserting an evidentially unmarked sentence in this class, like “The painting is beautiful,” conveys evidential information by default: it carries a direct evidential signal which in this example conveys that the source of the claim is the speaker’s direct experience of the painting. This is why we can’t assert “The painting is beautiful” without having seen the painting.^{39,40}

Moreover, because such sentences convey this evidential information, we cannot use them as the complements of *that*-clauses to ascribe aesthetic beliefs to people, when their beliefs are acquired through testimony. To say that Margot believes the painting is beautiful is, in part, to convey that Margot believes the painting is beautiful as a result of her own direct experience of the painting. The same is true when the believer is oneself: I cannot admit to thinking the painting beautiful unless I have seen it and found it beautiful myself. However, this does not mean that people fail to acquire aesthetic beliefs—beliefs about whether objects have aesthetic properties—on the basis of testimony. It just means that such beliefs have to be described differently. If Margot tells me the painting is beautiful and I am sure she is right, I can be said to believe that the painting must be beautiful.

What I am suggesting is that the primary proposition asserted by an unqualified aesthetic claim that weighs in on some aesthetic question, straightforwardly attributes an experience-independent aesthetic property to an object (or denies that it applies), in accordance with *Experience Independence*. ‘Weighing in’ on an aesthetic question is making a claim about whether or not something has some aesthetic property—so, for instance, “The painting is beautiful” and “The painting isn’t beautiful” are both unqualified claims of this kind. This primary proposition is experience-independent, and independent also of any content concerning evidential source, and we can see it, as such, in action, in some constructions that do not weigh in on an aesthetic question in this way—in, for instance, “Is the painting beautiful?” and “If the painting is beautiful, it will fetch a high price at auction,” neither of which carry evidential information. (However, someone who asks the former might expect her hearer to have seen the painting herself; otherwise, she might ask, “Is the painting thought to be beautiful?” etc.) The

³⁹ Note that von Fintel and Gillies generally contrast indirect inference on the one hand with both direct (observational) evidence and trustworthy reports (of what someone else has directly observed) on the other. For aesthetic claims, however, testimony should be lumped with inference rather than with direct evidence. Only if the speaker has observed the painting herself can she felicitously say, “The painting is beautiful”; and “The painting must be beautiful” is appropriate when the inference behind it is based on testimony.

⁴⁰ It is not unknown, incidentally, for a language that does have a grammatical category for evidentiality to exhibit a similar feature: in Kolyma Yukaghir, for instance (an almost dead language found in Siberia), use of the ‘simple non-Future’ form, which stands in contrast with the form that marks the hearsay evidential, does in fact convey that the speaker has participated in, or otherwise witnessed, the events she narrates, even though the form is not explicitly marked with an evidential. (See Maslova 2003, Ch. 6 §2.)

primary proposition is what is ‘at-issue’ when asserting an unqualified aesthetic claim in aesthetic conversation; this means, among other things, that whether or not some object has some (experience-independent) aesthetic property is what is being debated between people who disagree about an aesthetic matter. The evidential information carried by an unqualified aesthetic claim is supplementary to the claim’s main job of asserting this primary proposition.

My proposal makes sense of the data concerning aesthetic testimony. It also makes sense of another aspect of the behavior of aesthetic claims that might otherwise be puzzling, just like the testimony phenomenon: aesthetic language does not behave straightforwardly in arguments. In some ways this is not surprising; we do not generally *argue* ourselves into aesthetic opinions, except in a way that is supplemental to aesthetic experience: in looking at a picture with a friend, say, she may get me to see things in it I hadn’t noticed before, or to see things in a new way. Kant has argued that in fact we cannot validly infer an object’s beauty from the presence of any other properties it may have, and Frank Sibley has argued that we cannot infer the presence of any aesthetic property in an object from a description of the object’s non-aesthetic properties.⁴¹ It is not surprising, then, that arguments such as this one are unlikely to succeed:

- (A1) Anything rainbow-colored is beautiful.
 These shoes are rainbow-colored.
 Therefore, these shoes are beautiful.

We are unlikely (my four-year-old daughter aside) to accept any aesthetic premise like “Anything rainbow-colored is beautiful.”

However, there are many ways in which aesthetic terms can appear in valid arguments that have some chance of being accepted. For instance, here’s an apparently decent inference of a kind that has presumably been made before:

- (A2) If the vase is beautiful, it will fetch a high price at auction.
 The vase is beautiful.
 Therefore, it will fetch a high price at auction.

This argument, of course, does not have an aesthetic conclusion. We might want to say, then, that it is just (non-trivial) arguments with aesthetic conclusions that are problematic, like (A1) above, because they are bound to have a problematic premise. However, if (A2) is all right, then surely it would also be all right to run its conditional through a *modus tollens*:

- (A3) If the vase is beautiful, it will fetch a high price at auction.
 It won’t fetch a high price at auction.
 Therefore, the vase is not beautiful.

⁴¹ Someone whose Plato is better than mine has also reminded me of this in the *Phaedo*: “[I]f someone tells me that a thing is beautiful because it has a bright color or shape or any such thing, I ignore these other reasons ... but I simply ... cling to this, that nothing else makes it beautiful other than the presence of, or the sharing in, or however you may describe its relationship to that Beautiful we mentioned” (Plato 1997a: 100d)

This does give us an aesthetic conclusion. The starting point of my discussion has been, though, that we will only claim that the vase is not beautiful (etc.) if we ourselves have seen the vase. And neither of (A3)'s premises seems to require having seen the vase in order to be acceptable.

(A3) does not provide the clearest example of the phenomenon I want to illustrate, and I will come back to it. My point is that there are plenty of apparently valid arguments with premises we may very well take to be true, that have aesthetic conclusions we nevertheless balk at accepting. Here are two examples (I have chosen naturalness of expression over explicitness of logical form):

- (A4) This vase (which I see before me) is beautiful.
Physically identical objects have the same aesthetic properties.
Therefore, that (physically identical) vase (in the cupboard, which I have never seen) is beautiful.
- (A5) Margot says that *Middlemarch* (which I have never read) is a fantastic novel.
Margot is always right about books.
Therefore, *Middlemarch* is a fantastic novel.

Each of these arguments has a generalization for a premise, and it is of course possible to take issue with each of the generalizations. So, you might not think that aesthetic properties supervene on physical properties, or you might think that while some do, others don't; and you might think that while Margot is very often right about books, or even almost always right about them, it is going too far to say that she is always, without exception, right about them.

However, this kind of attempt to explain our unease with these arguments, and any others like them, finds fault in the wrong place. Even if in particular cases the generalizations are doubtful, we often rely on generalizations like these, and we cannot dismiss them all; we believe them as much as we believe other casual generalizations that inform our day-to-day casual reasoning: "Margot is always right in her mental arithmetic," etc. The problem with these arguments is simply that their conclusions convey the wrong evidential information. Conclude, instead, that the vase in the cupboard must be beautiful, that *Middlemarch* must be a fantastic novel, etc., and we will be all right.

(A3) is an awkward case because the way the argument goes muddles our expectations about the kind of evidential information that should be conveyed by its conclusion. If I know the vase won't fetch a high price at auction because I have seen it myself and on that basis know it isn't beautiful, then the argument, as a piece of reasoning I might actually go through, is redundant, but the form the conclusion takes will sound all right. If I know the vase won't fetch a high price at auction because an expert has told me so, but I haven't seen it myself, then I would have to conclude instead that it must not be beautiful.

There are of course several questions that arise about my proposal. One concerns what it is for evidence to be 'direct' or 'indirect'—at least as this is embodied in our linguistic practice. I am going to continue simply to trade on our intuitive understanding of this, although I do want to note that a speaker's conveying that she has direct grounds for a generalization (e.g. the generic "Roses are beautiful") seems to require that she have direct grounds for some sufficient number of instances (individual beautiful roses). A speaker's conveying that she has direct grounds for a

particular claim (“Margot’s rose is beautiful”) requires that she have direct grounds for that particular thing—in the case of Margot’s rose, she must have seen it and found it beautiful.

Another question concerns the primary proposition asserted by a claim that weighs in on some aesthetic question. Qualified claims like “It must be beautiful” and “It’s probably beautiful” can also be said to weigh in on an aesthetic question (on the question of an object’s beauty, in their case). I have suggested that when a speaker is more or less certain of the truth of an aesthetic claim—whatever her evidential source for it—the practical difference between choosing to assert (for instance) “It must be beautiful” rather than “It is beautiful” is simply that they differ in their evidential implications. But, on the face of it, they also assert different primary propositions. “It must be beautiful” makes a claim to epistemic necessity and “It is beautiful” does not.

Von Fintel and Gillies (§2) point out that there have been theorists who favor making the primary propositional content of *must* ϕ exactly the same as the primary propositional content of ϕ . The epistemic contribution of the modal—its claim to necessity (we are ignoring, for the moment, its evidential contribution)—could then be explained either as secondary propositional content,⁴² or as a ‘speech act modifier’: while an utterance of ϕ would be a straightforward assertion, an utterance of *must* ϕ would be a different kind of speech act. They roundly reject both ideas, however, on the grounds that they fail to account for the embedding behavior of epistemic modals. If I say to my doctor, “Must it be pneumonia?”, I am asking whether the evidence entails a diagnosis of pneumonia. If she replies, “Well, no, it doesn’t have to be pneumonia, but that’s the most likely cause of your symptoms,” she denies that it must be pneumonia.⁴³ The contribution of the modal is clearly affected by the embedding operator, as we would expect on the assumption that it contributes to primary content.⁴⁴ And presumably, when an utterance of a sentence of the form *must* ϕ is an assertion, it is the primary proposition that is asserted.

We can accept, though, that “It’s beautiful” and “It must be beautiful” have different primary propositional content, while still doing justice to the observation that, in general, the practical difference between them concerns their evidential import. Von Fintel and Gillies give the following sample conversation, adapted from the work of Mandy Simons, as an example of the kind of thing that might make you think that epistemic modals do not contribute to primary content:

Q: Why isn’t Louise coming to our meetings these days?

A: She might/must be too busy with her dissertation.

⁴² They suggest that Christopher Potts’ work on parentheticals could be adapted for epistemic modals. In an utterance, of, for instance, “Margot, a musician, is never on time,” the idea is that two propositions are expressed: *Margot is never on time* and *Margot is a musician*. The former is considered primary and the latter secondary, given their relative prominence in the sentence. Potts considers the secondary content here a kind of conventional implicature. I am not sure exactly how von Fintel and Gillies think this idea could be adapted for *must*.

⁴³ Note that, in English, *must* doesn’t embed under negation, but *has to* does. (For some more interesting contrasts between these modals, see von Fintel and Iatridou 2007.)

⁴⁴ Compare, e.g., “Is Margot, the musician, ever on time?” A “No” answer is taken to claim that she’s never on time. (If the hearer doesn’t think Margot is a musician, the question is taken to be infelicitous: “Who? Margot isn’t a musician.”)

As they point out, what is proposed as the answer to the question is that Louise is too busy with her dissertation, not that she might be, or must be, too busy with her dissertation. The contribution of the modals seems to be that they indicate the speaker's degree of confidence that she has given the right answer. Similarly, we can imagine the following sort of exchange:

- Q: Is the fresco beautiful, do you know?
A: It must be—it was painted by Raphael.

The fresco's beauty is the topic of conversation, but the person who answers the question has not seen the fresco herself and therefore answers using *must*.

Simons (2007) discusses a class of embedding verbs that she says often function in conversation as epistemic or evidential signals, when it is the embedded clause that is the 'main point' of the utterance. Her examples are constructions such as *I think that p*, *I suppose that p*, *I heard that p*, *I discovered that p*, etc. She does not discuss the role of epistemic modals, but as von Fintel and Gillies say, her observations easily transfer to them. She uses the kind of question and answer above to diagnose main point content, where the main point of an utterance is the proposition communicated by that utterance that makes it relevant to the conversation. The main point of the answer concerning Louise is thus, as already noted, that she is too busy with her dissertation; the main point of the answer about the fresco, "It must be (beautiful)", is (if we can subtract for a moment the evidential implications of putting it this way) that it is beautiful.⁴⁵ Simons' view aligns with that of von Fintel and Gillies, however, on the question of what is asserted in utterances employing the embedding verbs of interest to her. Von Fintel and Gillies argue that *must φ*, *might φ*, etc. have primary propositional content that includes their modal component, and presumably it is this propositional content that is asserted in utterances of them. Simons argues, in line with this, that even when her embedding verbs are used evidentially, and it is the embedded clause that is the main point of an utterance, the main clause proposition is still what counts as asserted. She uses facts about agreement and disagreement to support this claim, offering these sample dialogues about our friend Louise:

- A: Why isn't Louise coming to our meetings these days?
B: Henry thinks that she's left town.
C: a. But she hasn't. I saw her yesterday in the supermarket.
b. No he doesn't. He told me he saw her yesterday in the supermarket.

C's utterance in (a) is a response to B's main point, but not a denial of what B has asserted: it is a rejection of B's answer as a good one. C's utterance in (b), on the other hand, is a denial of B's

⁴⁵ I have noticed that writers on this topic often fall prey to the awkwardness resulting from there being no evidentially neutral way to express such a thing. For instance, Miles Rind (2003), whom I mentioned earlier (§2.2), says in that article, "... [T]he testimony of others ... may *persuade me that a certain thing is beautiful*, but I cannot legitimately express that persuasion by saying 'X is beautiful'" (my italics). But, just as I cannot legitimately express that belief by saying 'X is beautiful,' I cannot legitimately describe that belief by calling it the belief (or the persuasion) that X is beautiful. So Rind has misdescribed the persuasion referred to in the italicized phrase. I think he should have said, "... persuade me that a certain thing must be beautiful" (66). If he rejected this formulation, or it never occurred to him, maybe this is because he thinks the *must* construction is epistemically weak: he clearly wants to convey epistemic strength. As I have argued, however (or let von Fintel and Gillies argue for me), *must* is not epistemically weak.

claim—a claim about what Henry thinks. So, “even though B’s main point is that Louise (might have) left town, her utterance commits her to the proposition that Henry thinks this. This is the kind of commitment which accompanies assertion. Thus, what is asserted is distinct from the main point content.” (1041) And that’s fine—in these cases, and in the aesthetic case as well. Even though utterances of “It’s beautiful” and “It must be beautiful” assert different primary propositions, their main point can be—and, I claim, generally is—the same. We can, if we like, generate an analogous dialogue:

- D: Is the fresco beautiful?
E: It must be—it was painted by Raphael.
F: a. But it isn’t. I’ve seen it and it’s horrible.
b. It may not be. Most of it was done by other members of his workshop.

Of course, in this case, F’s utterance in (a) is a denial of E’s main point which, given the nature of the example, is also incompatible with E’s assertion that the fresco must be beautiful.

A more difficult question about the solution I am proposing concerns the status of the evidential content of unqualified aesthetic claims. How is it conveyed? Is it as conversational implicature, conventional implicature, presupposition, or something else? As this will be a fairly lengthy discussion, it will occupy the next chapter.

4. MULTIPLE DIMENSIONS OF CONTENT

The idea of using conversational implicature as a solution to the problem of aesthetic testimony has in fact been suggested by others. Hopkins contends, however—and Meskin echoes Hopkins' claim—that such a strategy can only account for the illegitimacy of making testimony-based aesthetic claims expressed in speech acts, and cannot account for the supposed illegitimacy of acquiring (unexpressed) aesthetic beliefs on the basis of testimony. They therefore reject it as a general solution to the problem.⁴⁶

The complaint, presumably, is that if the evidential content of an unqualified aesthetic claim gets attached to the claim only as a result of being uttered in a specific conversational context, then the sentence used to make the claim, considered in itself, is evidentially neutral; and if this is the case, then there should be no problem with believing the content of the sentence, even if the source of the belief is testimony. But it is never felicitous to attribute a testimony-based aesthetic belief of this kind to anyone. So (the complaint goes) conversational implicature can't explain the problem of unqualified aesthetic belief and testimony.

This complaint is at first glance compelling, but I will ultimately conclude that it misses the mark. I will conclude, that is, that the evidential content of unqualified aesthetic claims is in fact conveyed as a form of conversational implicature, in spite of a strong initial appearance to the contrary. In the next section (§4.1), I will set out the basic contours of conversational implicature and explain why the Hopkins/Meskin objection seems to be so strong. This will lead to the consideration in the following section (§4.2) of an explanation of aesthetic claims' evidential content in terms of alternative, semantic phenomena. I will conclude there that such phenomena are a bad fit for aesthetic claims' evidential content. In §4.3, I will return to the hypothesis that it is conversational implicature at work after all. This will be my working hypothesis throughout Chapters 5 and 6, where I will try to support the hypothesis further by showing why such implicatures are 'calculable,' given the kind of primary content that aesthetic claims carry. I will also draw comparisons between aesthetic claims and claims of many other kinds that exhibit similar behavior. However, the necessary detail of these chapters will require me to put on hold the most significant problem with an explanation of aesthetic claims' evidential content in terms of conversational implicature: that such content does not appear to be 'cancelable.' I will eventually argue that, in spite of the traditional line on cancelability—that conversational implicature content *must* be cancelable—this is not always the case, and is not the case for aesthetic claims. This has already been suggested in the linguistics literature for other kinds of claims; I will simply apply these remarks to the aesthetic domain.

4.1 Conversational implicature

A conversational implicature arises, Grice says, from a conversational context in which a speaker takes advantage of the general expectation that conversation will—in a nutshell—make sense, in order to convey a thought that her utterance does not literally express. If conversational implicature were a good fit for the aesthetic case, we would understand an unqualified aesthetic claim to be literally expressing a primary proposition, in accordance with *Experience Independence*, that attributes an aesthetic property to an object, and as conversationally implicating the secondary content that the claim is made as a result of the speaker's direct

⁴⁶ Hopkins (2000: §3) cites the source of the idea as Mothersill (1984: 85, 159-160). I certainly know of no earlier source. Meskin (2004: 76) explicitly agrees with Hopkins in rejecting the idea.

experience. The resulting explanation on offer of why we are unwilling to make unqualified aesthetic claims on the basis of testimony, would then be that such claims would suggest something the speaker would know to be false, even though that suggestion were not literally expressed by her utterance but were merely the result of conversational implicature.

Everyone's favorite Gricean example of conversational implicature concerns the philosopher who has been asked for a reference for one of his students, and writes, "Dear Sir, Mr. X's command of English is excellent, and his attendance at tutorials has been regular. Yours, etc." (1989a: 33) The clear conversational implicature is that Mr. X is no good at philosophy. This thought is not stated, or literally implied by what is stated, in the letter, but it is necessarily attributed to the writer of the reference if we are to understand him as providing the sort of information he was asked for, but did not write down. The implicature clearly arises from the specific context of giving a reference that is at work here; asserting in other contexts that Mr. X has an excellent command of English and regular tutorial attendance will happily not tend to suggest that he is no good at philosophy.

If it is conversational implicature at work in the aesthetic case, however, it would have to be generalized conversational implicature (GCI). A sentence is associated with a GCI when that sentence more or less always gives rise to the implicature no matter what its specific conversational context of utterance—but where the implicature is not, as with conventional implicature, solely due to the meaning of the uttered sentence or any of its constituent parts. Unqualified aesthetic claims would have to carry this kind of implicature because such claims more or less always convey that they are based on the speaker's direct experience, no matter what the conversational context.

As Grice noted, it can be hard to distinguish GCIs from conventional implicatures, presumably because it can be hard to distinguish the semantic effects of a form of words from any pragmatic effects that follow those words wherever they go. There is in fact ongoing debate about the extent to which GCIs are due to semantic (conventionalized) as well as pragmatic factors.⁴⁷ However, the phenomenon itself is robust, even if its theoretical explanation is somewhat up for grabs. Grice gives several examples of GCI involving the use of the indefinite article: for instance, if I say to you, "I saw a man get hit by a car yesterday" (this is my own example), you would presumably be surprised to learn that the man in question is my husband (or the car, for that matter, my car): I would normally implicate that I don't know the man who got hit. The implicature arises because if I did know the man, I would be expected to be more specific about who it was ('my husband', 'a friend of mine', or even 'John', if I know you know John too). But the implicature is not plausibly due to the meaning of the sentence—we would not want to say there is something tricky about the meaning of 'a man' here that rules out its applying to (say) my husband—and it is therefore not conventional. Standard 'scalar' implicatures are also in this class: e.g. if I say, "Some of the kindergartners can already read," I would normally be taken as implicating that not all of them can read, although what I said is certainly consistent with "All of the kindergartners can read." If I were in a position to make the claim of all of them, I would be expected to do it: the implicature is generated by my choice of *some* rather than *all*. The quantifiers (e.g. *all*, *most*, *many*, *some*) are just one of a number of 'Horn scales'⁴⁸—i.e. ranked sets of semantically related, salient alternatives—that generate GCIs in this way. (Other examples include modals—*necessarily*, *possibly*—and adverbs like *always*,

⁴⁷ See e.g. Potts (2013).

⁴⁸ After L. R. Horn.

often, sometimes.)⁴⁹ Scalar implicatures and Grice's example involving the indefinite article are standardly attributed to Grice's 'first Maxim of Quantity': i.e. "Make your [conversational] contribution as informative as is required." (1989a: 26)

Its being GCI at work in the aesthetic case—if it is conversational implicature at all—complicates the Hopkins/Meskin objection. Clearly, someone can believe without difficulty that Mr. X has an excellent command of English, that he has attended tutorials regularly and that he is good at philosophy. The fact that, in some contexts, asserting the first two beliefs will tend to suggest that the third is not true—or is not believed to be true—is obviously not a hindrance to believing all three. Thus the Hopkins/Meskin objection is a good objection to 'particularized'⁵⁰ conversational implicature: on this model, it should be unproblematic for me to believe both that X is beautiful, and that I have never seen X myself; but it is not.

A GCI, however, attaches itself to a sentence more or less whenever it is uttered, not just to particular occasions of its utterance. And it has in fact been suggested in the linguistics literature that a sentence's typical GCI is retained, as part of what is believed, when that sentence is embedded in a belief context. Chierchia, for instance, who focuses on scalar implicatures, says that "John believes that some students are waiting for him" carries the implicature that John believes that not every student is waiting for him. Chierchia admits that intuitions about such cases may not be very clear, but he points out that when someone asserts a sentence, we are generally licensed to attribute a belief to that person with that sentence as complement,⁵¹ and he says that it would be odd if the kind of implicature present in his example did *not* survive as part of the attributed belief, as "after all we [are] just reporting what [the assertion] gives us grounds for." (2004: 45)

This is not as it stands a very convincing argument, as of course it applies equally to assertions carrying particularized conversational implicatures which do not necessarily carry over to the corresponding belief attributions: someone who asserts, for instance, that John has excellent English and attends classes regularly, in a context in which this implicates that John is no good at philosophy, can be said to believe that John has excellent English and is regular in his attendance *without* the belief attribution carrying the corresponding implicature. It may or it may not carry it, depending on the context in which the belief attribution itself is made.⁵²

But I think Chierchia's point simply needs to be spelled out more fully: there is something about the nature of the connection between assertions carrying GCIs, and the belief attributions they license, that is relevant here. When an assertion goes right (so to speak), it expresses a belief held by the speaker; the corresponding belief attribution made by another reports that that person holds that belief. Of course, people believe many things that they do not assert, for one reason or another, and belief attributions frequently do not depend on people's prior assertions for their justification. However, the connection between assertion and belief attribution makes it plausible that a given belief attribution is inappropriate—at best misleading, at worst false—if the person in question would not be *willing* to make the corresponding assertion, when certain conditions are met—when, that is, she has sufficient self-knowledge, no reason to be secretive,

⁴⁹ Useful discussion in Levinson (2001: Ch. 2).

⁵⁰ Grice's term (1989a: 37).

⁵¹ Presumably allowing for suitable translation of indexicals, etc.

⁵² If I ask, "Is John any good at philosophy?" and you say, "The professor thinks that his English is excellent and that he attends class regularly," you will implicate that the professor thinks he is no good (especially if you smirk). But if we know that John is a native Spanish speaker, and I ask, "Should John be asked to translate his lecture notes for other Spanish speakers in the class?" then the same answer will not carry that implication.

and an average amount of polite helpfulness when someone asks her a direct question, say. An appropriate belief attribution thus corresponds with what a person would assert under such conditions. Moreover, an assertion of “Some students are waiting for me” will in almost all contexts implicate that the speaker thinks not all of the students are waiting for her—that is what makes the implicature a GCI, rather than a particularized implicature. But, then, a speaker will generally be unwilling to assert “Some students are waiting for me” unless she thinks that not all of the students are waiting for her.

So, if John thinks that all of the students are waiting for him, he is unlikely to assert, “Some students are waiting for me,” precisely because of the GCI that such an assertion would carry. Saying that John believes some students are waiting for him, when he believes they are all waiting for him, would likewise be misleading, as we would in this case be suggesting that it would be appropriate for John to assert that some students are waiting for him, an assertion which of course would carry the GCI that not all of the students are waiting for him.

This argument is just a rough sketch of what could be said here; as usual, trying to specify the ‘conditions’ I mention in a principled way is presumably a thankless task. In any case, I do share what I take to be Chierchia’s intuitions concerning belief attributions and scalar implicatures, which I see no reason not to take in this context as exemplary of GCIs generally. I agree that saying John believes some of the students are waiting for him will generally suggest that he believes not all of the students are waiting for him. If the phenomenon of aesthetic testimony were due to GCI, and belief attributions do carry GCIs in this way, we would thus expect it to be problematic to attribute an unqualified aesthetic belief to someone without the relevant personal experience: to say of John, for instance, that he thinks X is beautiful, when he has never seen X. The phenomenon Hopkins and Meskin both advert to as evidence against conversational implicature as a solution to the problem of aesthetic testimony—our *apparent* failure to acquire testimony-based aesthetic beliefs—thus seems to be explicable after all on the hypothesis that it is GCI at work. The general phenomenon of GCI in attitude contexts would predict that “John thinks X is beautiful” implicates that John has personal experience of X; similarly, “John thinks X must be beautiful” would implicate that John arrived at his belief concerning X’s beauty on the basis of inference.⁵³

⁵³ Chierchia takes the phenomenon of scalar implicatures that are apparently embedded in belief contexts as evidence for an account of such implicatures that relies more on grammar (i.e. semantics and syntax) than do Gricean accounts, which emphasize the pragmatic side. He claims that the latter face difficulties providing the right kind of pragmatic calculation of the implicatures arising in embedded contexts. Others disagree. E.g. Benjamin Russell (2006) argues that such implicatures are easily explained in Gricean terms. See also Potts (2013).

In Chierchia’s account, the implicature of *some students are waiting for him* is computed ‘locally’ to yield *some but not all of the students are waiting for him*, which is what is then inserted into the belief context in the rest of the grammatical processing of the sentence to give a representation of the sentence as a whole: *John believes that some but not all of the students are waiting for him*. Russell maintains that it is in fact possible under Gricean theory to yield a grammatical representation of the sentence from which the implicature is absent, and then to compute the implicature ‘globally,’ on the basis of that representation, using standard Gricean pragmatic reasoning.

I am in no position to weigh in on this debate in linguistics. However, if we take it that the evidential component of aesthetic assertions is due to GCI, and that the evidential implicatures survive in belief contexts as part of what is believed, it is clear that such implicatures cannot be computed entirely locally in the way that Chierchia computes his apparently embedded scalar implicatures. This is because the putative implicature of e.g. *the vase is beautiful* when it is inserted into a belief context (“John believes the vase is beautiful”) depends on the details of that context: what is implicated is that John has seen the vase to be beautiful. If what is claimed is that Mary believes the vase to be beautiful, what is implicated is that Mary has seen the vase to be beautiful, etc.

However, even GCIs typically exhibit all the usual hallmarks of conversational implicatures generally, in being *calculable*, *reinforceable* and *cancelable*. These features arise from the phenomenon of conversational implicature being pragmatic in nature. Conversational implicatures are ‘calculable’ in that they can be *worked out* from how the general principles governing reasonable, relevant conversational moves bear on the meaning of an utterance in its conversational context. Grice’s referee has, at the level of what is asserted, clearly not supplied the kind of information about Mr. X that the hire committee wants. The referee is not being uncooperative (he is writing the letter, after all), he must know more than he has said about Mr. X’s talents (Mr. X is his student), and he knows what kind of information the hire committee is looking for. To make sense of the referee’s conversational move, we must infer that he wants to communicate what he thinks about Mr. X but is squeamish about writing it down. This only makes sense on the supposition that he thinks Mr. X is no good at philosophy, which is what he therefore implicates. The implicature is not attributable solely to the meaning of what he says, but arises from his saying it in this particular context.

Grice also hazards a possible explanation of how we can work out the GCIs arising from use of the indefinite article. It is clear that “I saw a man get hit by a car yesterday” is subject to the same kind of conversational expectations about what kind of information, and how much, it is appropriate to offer in a given context, as the example involving the philosophy reference letter: in saying that I saw a man get hit, I have not been more specific, so it is presumed that the man in question is unknown to me. What is less clear, in this case, is just why a greater level of specificity is always expected here, if it is possible to provide it. Grice’s proffered explanation is that, as the degree of relationship between people makes a big difference to the usual consequences of any dealings they have with each other, being specific about the degree of relationship allows the hearer to shape her expectations appropriately, and preempts many further questions she may have. The presumption, then, is that the speaker will be specific about such a relationship when doing so streamlines the conversation in this way.

The calculations behind GCIs typically involve a greater degree of conventionality than those for particularized conversational implicatures. Utterance of “Some of the kindergartners can already read” will typically implicate the negation of the stronger alternative statement using *all* (“Not all of the kindergartners can read”), but it will not typically implicate the negation of any stronger statement whatsoever (e.g. the negation of “70% of the kindergartners can already read”). *Some* and *all* are in salient opposition as *some* and 70% are not, and this is for conventionalized rather than pragmatic reasons: e.g. Levinson (2001: 79) writes that “items in the same scale must be ... of the same form class, in the same dialect or register, and lexicalized to the same degree.” *Some* and 70% are neither in the same register nor lexicalized to the same degree, and thus do not form a salient opposition of a kind that can be taken advantage of for scalar GCIs.⁵⁴ However, given the salient opposition that conventionally exists between *some* and *all*, the Maxim of Quantity can be invoked to explain how these scalar implicatures arise. They are still, that is, pragmatic in nature.

Reinforceability and cancelability are two sides of the same coin, that being the fact that the proposition expressed by a sentence is consistent with the negation of any conversational implicature it happens to carry in a given context of utterance: i.e. (to put it less precisely but

⁵⁴ Levinson points out that there is no Horn scale for (*iff*, *if*) for the same reasons: “*iff* (*if and only if*) belongs to a specialized register and is not monolexemic like *if*.” (2001: 79)

more intuitively) there is a semantic distance between an utterance and its conversational implicature. Conversational implicatures, including GCIs, are thus held to be reinforceable—i.e. they can be made explicit—without an odd air of redundancy; and they are held to be cancelable, in that a speaker may felicitously disavow the implicature content. (The term ‘cancellation’ is also taken to cover cases in which a sentence’s usual conversational implicature simply doesn’t arise at all in a given context of utterance; Grice calls this ‘contextual’ cancellation.) So, I can reinforce a GCI by saying, “I saw a man get hit by a car yesterday—don’t worry, it wasn’t my husband”, or “Some—though not all—of the kindergartners can already read”; and cancel it by saying, “I saw a man get hit by a car yesterday—actually, it was my husband” or “Some of the kindergartners can already read; in fact, all of them can.”

Cancelability has received far more attention than reinforceability and has generally been taken, since Grice, to be a necessary property of conversational implicatures. It is natural to think that conversational implicatures would have to be cancelable: as their ‘calculation’ depends on reasoning about the relation between what is literally said and the context of utterance, including features of the speaker’s psychology, it seems that a speaker should always be able to block the calculation by denying that one (or more) of its premises holds. At the very least, it should be possible for the speaker to say, in order to block a conversational implicature to the effect that *p*, “but I do not mean to imply that *p*.”⁵⁵

If conversational implicatures are always cancelable, however, then it seems after all as though they cannot be invoked as the explanation for the phenomenon of aesthetic testimony. This is because the evidential implications of unqualified aesthetic claims do not seem to be cancelable. I cannot say, without oddity, “The painting is beautiful, but I’ve never seen it.” I would have to say, instead, “The painting is said to be beautiful, but I’ve never seen it,” or “The painting must be beautiful, but I’ve never seen it.” This is what, finally, gives the Hopkins/Meskin objection the most bite against the idea that the phenomenon concerning unqualified aesthetic belief and testimony could be due to GCI. I can cancel the implicature of “I saw a man get hit by a car yesterday,” and I can be said to believe that I saw a man get hit by a car yesterday, even though the man in question is my husband. But I can’t make an unqualified aesthetic claim on the basis of testimony because I can’t cancel the commitment of such a claim to the effect that it is made as a result of my own direct experience; and I can’t describe myself as holding an unqualified aesthetic belief on the basis of testimony because that description also conveys the same uncancelable commitment as part of what is believed.⁵⁶

4.2 Presupposition and conventional implicature

At this point it seems, then, that a semantic phenomenon like presupposition or conventional implicature would have to be a better candidate for the explanation at issue. Content that is conveyed as presupposition or conventional implicature is tightly tied to the meaning of the sentence uttered in a way that conversational implicatures are not, and it is thus not generally held to be calculable, reinforceable or cancelable.

How exactly to define presupposition and conventional implicature, and what kinds of constructions count as providing examples of these, are still, it seems, to some extent

⁵⁵ Grice makes this suggestion in (1989b: 44).

⁵⁶ Ninan (2014) also rejects conversational implicature as an explanation for the phenomenon of aesthetic testimony, explicitly giving as his reason that the evidential implications of the unqualified claims at issue are not cancelable.

controversial. It has even been argued that there is no such thing as conventional implicature.⁵⁷ Nevertheless, the basic idea behind both phenomena is that some constructions make semantic contributions to the sentences in which they appear, that allow those sentences to express multiple dimensions of meaning—not just a single proposition.

There are two fundamental ways in which presuppositions have been thought to differ from conventional implicatures (although, as I said, there is apparently not universal agreement here). The first concerns the logical relation between the primary proposition expressed by a sentence and the sentence's (semantically derived) secondary content. If the primary proposition does not depend logically on the secondary content, the secondary content is generally thought to be conventional implicature; if it does, the secondary content is generally thought to be a presupposition. In an utterance of "Margot's car is red," for instance, the possessive phrase is a trigger for the presupposition that Margot has a car, and if the presupposition is false, the primary proposition expressed by the utterance cannot be true (or, perhaps, the utterance fails to express a primary proposition at all). Other presupposition triggers of this kind include definite descriptions, demonstratives, and factive verbs (like 'know', 'realize', 'notice').⁵⁸

The second traditional mark of difference between conventional implicature and presupposition is the extent to which the speaker takes the secondary content of her utterance to be already in the common ground. Conventional implicature content is generally new to the conversation.⁵⁹ Presupposition triggers, on the other hand, tend to allude to something that must be taken for granted—as if it had already been mentioned—for an utterance to make sense in its conversational context. (This feature is clearly related to the first: if the primary proposition expressed by an assertion, say, cannot be true unless the utterance's presupposition is true, a hearer must take on board the truth of the presupposition in order to make sense of the speaker's assertion of the primary proposition.) In practice, however, the presupposition need not actually have been mentioned: hearers easily 'accommodate' many presuppositions that are new information for them. At a party, for instance, I might say, "My husband will be coming later" to someone who doesn't know I am married.⁶⁰

⁵⁷ See Bach (1999). Bach principally objects to Grice's original characterization of conventional implicature as being a dimension of the content of an utterance that is not part of what is *said* by the speaker who makes the utterance. Grice characterizes it in this way because he thinks that the truth value of conventional implicature content does not affect the truth value of the whole utterance. (Grice 1989a: 25-26.) Bach says that a sentence may express more than one proposition, and that when it does, the secondary proposition (the one with less importance to the conversation, when the sentence is uttered) may not be salient enough to affect people's judgments of the truth value of the sentence. Bach maintains that the secondary proposition is still, however, part of what is said.

Bach's point has not in general been taken to show that there is no such thing as conventional implicature, however. Proponents of conventional implicature (e.g. Christopher Potts) often simply adopt Bach's multiple propositions framework while nevertheless still calling certain kinds of secondary content 'conventional implicature.'

⁵⁸ See e.g. Christopher Potts (2015).

⁵⁹ It is often, however, content that is uncontroversial, that the speaker expects the hearer to take on board without comment. In saying, "X's daughter is small but mighty", I do not expect my hearer to take issue with the contrast between smallness and mightiness; I do not expect this to be news to my hearer in the sense of being new knowledge, even if it is new to the conversation. (Of course, a speaker may blunder in this, as in anything; saying e.g. "The new dean is a woman, but very capable" is not recommended.)

⁶⁰ There are some kinds of examples that nicely illustrate the common-ground difference between conventional implicature and presupposition (although there are also exceptions). For instance, parentheticals are redundant if their content has already been stated (Potts 2007b: 671): "Margot is a musician. Margot, a musician, is never on time." But this is fine: "Margot has a car now. Her car is red."

The fundamental problem with attributing the evidential implications of aesthetic claims to presupposition has to do with this logical dependence of the primary proposition of a sentence with a presupposition on its presuppositional content. This dependence is responsible for standard presuppositional ‘projection’ behavior: i.e. the fact that a sentence’s presupposition is also presupposed by the complex sentences resulting from embedding that sentence under various operators (the presupposition is said to ‘project’ over the operators). So, for instance, the presupposition of “Margot’s car is red”—that Margot has a car—is also presupposed by the following:

Margot’s car isn’t red.
Margot’s car must be red.
If Margot’s car is red, her insurance will be more expensive.
Is Margot’s car red?

In some sense, the proposition that “Margot’s car is red” tries to express doesn’t get off the ground if Margot doesn’t have a car, and when the sentence is operated on in these ways, the resulting propositions are likewise compromised if the presupposition is false.⁶¹

However, it is one of the basic intuitions supporting the acceptance of *Experience Independence* that the truth or falsity of the primary proposition expressed by an utterance of e.g. “The painting is beautiful” does not depend on the speaker’s experience of the painting. The painting is beautiful, or not, independently of *any* particular person’s having seen it to be so. Intuitively, then, the evidential implication of an utterance of “The painting is beautiful” is not presuppositional in nature, and this is borne out by the very different projection behavior of this and other aesthetic sentences. Of the operators shown above (for “Margot’s car is red”), the implication that the speaker makes the claim about the painting on the basis of her own firsthand experience projects only over negation (“The painting isn’t beautiful”). The other embedding operators either result in no evidential implications (e.g. the conditional) or different ones (e.g. “The painting must be beautiful”).

Moreover, while in simple, positive contexts, presuppositions can’t be canceled—“Margot’s car is red, but Margot doesn’t have a car”—they often can be canceled under negation, especially with the right kind of intonation: “Margot’s car isn’t red—she doesn’t have a car!” Ninan (2014: 8-9) points out, however, that aesthetic sentences do not behave in this way⁶²: I can’t say, “The painting isn’t *beautiful*—I’ve never seen it!” The infelicity of this kind of utterance presumably has to do, again, with the fact that the truth of a sentence like “The painting is beautiful” doesn’t depend on a particular person’s having seen it to be so. Yet the utterance suggests that the painting can’t be beautiful unless I’ve seen it.

The phenomenon at issue does not seem, then, to be due to presupposition. Should we therefore attribute it to conventional implicature? Constructions involving conventional implicature exhibit the right kind of logical independence of the primary content from the secondary content—an independence I take to exist for aesthetic sentences, and which puts presupposition off the table. Possibly, then, conventional implicature is a better fit.

⁶¹ This is made more precise by the view coming down from Strawson, that a sentence’s presuppositions are definedness conditions: i.e. that a sentence fails to express a proposition at all if its presuppositions are false.

⁶² Ninan presents much the same argument against presupposition as the explanation of the phenomenon at issue.

The classic (but not, I suppose, uncontroversial) example of conventional implicature is provided by ‘but.’ It is part of the meaning of ‘but’ that it makes a semantic contribution to what is conveyed that goes beyond what is conveyed by ‘and.’ If I say, “X’s daughter is small but mighty,” I indicate not only that X’s daughter is small and that she is mighty, but that there is some kind of contrast between being small and being mighty. On a conventional implicature understanding of ‘but,’ the primary proposition expressed by my utterance is that X’s daughter is small and mighty—its truth value governs people’s intuitions about the truth value of the utterance as a whole, it is probably what I primarily want to convey in uttering the sentence, etc.—and the contrast I indicate by using ‘but’ is secondary content.

Many different kinds of construction have been considered to give rise to conventional implicatures. Christopher Potts (e.g. 2007b) focuses on ‘supplements’ and ‘expressives.’ Supplements include parentheticals (“Margot, a musician, is never on time”) and certain adverbials (“Unsurprisingly, Margot is never on time”). The parenthetical and adverbial phrases, he says, contribute conventional implicature content on top of the primary proposition expressed by the sentence (that Margot is never on time). Potts’ examples of expressives include Japanese honorifics, expressive adjectives like ‘damn’ (“that damn cat”), and epithets (“that bastard cat”). While, as far as I can tell, he doesn’t focus on expressive epithets with a thicker referential component, such as racial slurs, they are clearly also on his radar, as he uses one as an introductory example in one of his papers (Potts 2007a). In the same paper he also applies his theory of expressives to the formal and familiar second-person pronouns in languages like French and German (‘vous’/‘tu’, ‘Sie’/‘du’).

It is clearly a heterogeneous list. Moreover, expressives arguably differ from others in contributing a kind of content to utterances that is nonpropositional. It is plausible that the ‘damn’ in “That damn cat has been digging up the garden again,” for instance, serves only to express the speaker’s negative attitude towards the cat, rather than adding to the propositional content of the sentence (even as a secondary proposition).⁶³ Something else again seems to be going on in the case of French and German second-person pronouns. If I address someone as ‘tu’ rather than ‘vous,’ it is not plausible that my utterance expresses a proposition to the effect that my hearer and I share a familiar relationship; nor does it sound right, though, to say that my choice of ‘tu’ expresses a feeling or attitude I have towards my hearer. I could try to put it like that, I suppose: I could clumsily say that my choice of ‘tu’ expresses my attitude of informality towards my hearer. But this would not be an attitude in the same vein as the attitude of contempt I express if I call someone a jerk. I choose ‘tu’ rather than ‘vous’ in recognition of the appropriateness of that pronoun, in light of the informal relationship I have with my hearer.

The heterogeneity of the catalogue of putative conventional implicatures may suggest that it would not be hard to add the evidential implications of aesthetic sentences to the list. However, there is even at first glance a major difference between what carries the secondary content in the aesthetic case and what carries the conventional implicature content in the kinds of examples I’ve mentioned here. The conventional implicatures I’ve listed above are associated either with individual lexical items or with grammatical forms (parentheticals, for instance). The evidential content of aesthetic sentences is clearly not associated with an entirely general grammatical form, but it also cannot be straightforwardly associated with individual lexical items—aesthetic

⁶³ Potts (2007a: 176) takes this view. The view has been proposed for referentially thicker slur terms as well. For instance, Daniel Whiting thinks that the use of a slur conventionally implicates “a non-cognitive attitude of contempt (or scorn or derision or ...)” for those to whom the referential component of the slur applies. (2013: §1)

predicates, say—because, as I noted above, the evidential content carried by a sentence employing an aesthetic predicate varies depending on the type of sentence it is (whether simple, a conditional, a modalized sentence, a question, etc.).⁶⁴ The evidential content carried by aesthetic utterances depends on the nature of what is being said in its entirety.

4.3 The ‘Direct Evidentials’ and conversational implicature

This makes the phenomenon look, after all—in spite of the cancelability issue—much more like *conversational* implicature. And this appearance is strengthened by noting that aesthetic sentences are not the only class that carry secondary evidential information in this way. A slew of claims behave similarly: normative and evaluative claims generally, including moral claims and e.g. claims about what tastes good, what is funny, what is fun to do, what is scary, or boring, or exciting; claims about people’s personality traits, like being clever or introverted; claims about perceptual appearance properties—what things look, sound, feel, taste and smell like (“It looks like a horse”, “She smells like strawberries”); and presumably many others. That is, although such claims have primary content that is plausibly held to be independent of any particular person’s response to an object on a given occasion, an unqualified claim in any of these areas conveys that it is made as a result of the speaker’s direct experience, or—as with e.g. moral claims and claims about personality—her own assessment.

This is itself a pretty heterogeneous list. There is perhaps a general sense that these claims are united by being more ‘subjective’ in nature than other kinds of claim—where I use this word to suggest claims that require either a certain kind of idiosyncratic point of view, or special discernment in judgment, or a certain kind of imagination. For instance, our aesthetic judgments reflect our aesthetic sensibilities, and our judgments about others’ personalities reflect whatever wisdom we’ve managed to acquire about other people. And many perceptual appearance claims are made as a result of idiosyncratic imagination: I might have said, for instance, many years ago, that my seventh-grade math teacher looked exactly like Richard Nixon (he did!), knowing that not everyone would have seen the resemblance (but also not intending my claim to be that he looked like Nixon *to me*—or something like that). Presumably, people disagree about these kinds of judgments more often, and more intractably, than they do about things like the location of the post office or the color of a house.

However, not every kind of claim on the list fits this mold. I may say that my seventh-grade math teacher looked like Nixon, but I may also say, “Margot’s house looks orange”; I use the same capacities, more or less, to judge directly—by seeing—that X looks orange (in this light) as I do to judge that X *is* orange, and other people with normal color vision will by and large judge the same way. Yet it is unqualified ‘looks’ claims across the board, as well as other unqualified perceptual appearance claims, that generally convey that they are based on the speaker’s direct experience, and are thus not felicitously usable at secondhand. If I have never seen Margot’s house, I can say that it is supposed to look orange under the streetlights, or that it must look

⁶⁴ Expressives and slur terms, on the contrary, demonstrate the strength of the connection between the secondary content conveyed and the mere use of the term, whatever its sentential context. Such terms typically project out of even more contexts than presuppositions, including propositional attitude contexts: in an utterance of e.g. “Margot thinks that damn cat is the most precious animal on the planet”, the negative feeling conveyed by ‘damn’ is attributed to the speaker, not to Margot. As I mentioned above, however, the secondary evidential content of aesthetic sentences in belief contexts becomes part of what is believed, as is common for both presupposition and GCI.

orange under the streetlights, but not, plainly, that it looks orange under the streetlights—even though the primary content of the latter claim is something that I may easily believe.

The list is also heterogeneous in a different way. Moral claims and personality claims—which I described above as being made on the basis of a certain kind of assessment—are interesting cases, as they demonstrate that what counts as ‘direct’ and ‘indirect’ grounds for a claim depends on the kind of domain that is at issue. The use of epistemic *must* indicates that the speaker makes a claim on indirect grounds, which is to say, in the aesthetic case, for instance, that the claim is made on the basis of inference; however, in the moral and personality cases, inference *on the basis of relevant reasons*—the putative reasons why the belief is true—is in fact a direct way to arrive at that belief. This gives the use of epistemic *must* in these cases a slightly different cast: we use it when we make claims that have been inferred from *other* kinds of reasons. I don’t call Margot clever unless I know certain things about her—that, for instance, she acs all her classes, assembles wind-powered home electricity generators from the odds and ends of the average kitchen drawer, writes beautifully impenetrable poetry—that, in my opinion, make her clever. If I am simply told that Margot is clever, without knowing why she is supposed to be, I will indicate my source as hearsay: “People say Margot is clever”, “I’ve heard Margot is clever.” If my hearsay source is impeccable, or I know something about Margot that is not itself a direct manifestation of her cleverness—her cleverness in action—but that is clearly a result of her cleverness (as, say, her winning a MacArthur Fellowship), then I can say, “Margot must be clever.” Similarly, if I really can’t see anything wrong with gay marriage, but my priest tells me that it’s wrong, and I trust my priest’s moral judgment absolutely, then I might say, “Gay marriage must be wrong (but I really can’t see why).” On the other hand, and consistently with our discussion so far, “X is wrong, but I don’t know why” has a decidedly odd ring.

It will be useful to have a name for this broad class of claims—i.e. those simple, unmarked assertions, in this variety of topics, that carry evidential information as to being made on direct grounds—so I will call them Direct Evidentials. No doubt there are other kinds of claims that behave in this way, that haven’t yet occurred to me. In the moral case, the small but overlooked linguistic fact concerning secondary evidential content, that I have been discussing, has significant bearing, I think, on two debates in ethics. The first is the exact parallel, for moral claims and beliefs, of the problem in aesthetics at issue here: it has seemed, at least to some, that there is something illegitimate or at least odd about acquiring moral beliefs purely on the basis of another’s say-so; the problem, here too, is to explain why this is the case, given that acquiring many other kinds of beliefs on this basis is entirely unproblematic.⁶⁵ The second debate I have in mind concerns the connection between moral claims and motivation: the internalist view, for instance, that believing a certain action to be right (say) necessarily involves being motivated, to some degree, to perform it. I hope to explore these matters elsewhere.⁶⁶

⁶⁵ See e.g. Sarah McGrath (2011).

⁶⁶ Dorit Bar-On and Matthew Chrisman (2009) also suggest there is a similarity between moral claims, perceptual appearance claims, and aesthetic claims, that is something like the similarity I have been discussing. They say that each kind of claim has a ‘propriety condition’ on its assertion in addition to the condition that one believe the (primary) content of what one says: a ‘looks’ claim must be made as a result of the speaker’s visual experience; an aesthetic claim (their example is “That’s ugly”) must be made as a result of the speaker’s aesthetic attitude (in the case of ugliness, of the speaker’s negative aesthetic reaction). They do not, however, explicitly diagnose the propriety condition as a linguistic one that is due to such claims’ secondary evidential content. They also have a different view of what the additional propriety condition is on utterance of a proper moral claim: their view is that in addition to believing the claim, the speaker must be in an appropriate motivational state (for instance, if she says that

5. CONTENT AND CALCULABILITY: SOME NON-AESTHETIC EXAMPLES

I have said that it appears after all to be conversational implicature at work across this class of loosely connected conversational topics, in spite of the cancelability issue. This appearance is further bolstered by consideration of calculability. We can provide a plausible GCI calculation for the fact that assertion of an evidentially unmarked, simple sentence concerning any of these topics (e.g. “The vase is beautiful”, “Her house looks orange under the streetlights”) carries the presumption that the speaker has direct grounds for her claim. It will, however, take some time. As aesthetic claims present particularly complex issues, I will save them for Chapter 6. I will discuss cancelability in Chapter 7.

5.1 Presumption of evidential source

Levinson (2001: 112) offers a version of Grice’s second Maxim of Quantity (“Do not make your contribution more informative than is required”) which he calls the “Principle of Informativeness” (or “I-Principle”), and which in rough summary he renders, “Don’t provide unnecessary information, specifically don’t say what would be obvious anyway”. He proposes that one of the effects of this principle is the enrichment of minimal utterances by the licensing of inferences to what is normal or stereotypical: what can be considered a kind of default. For instance:

[A]nyone *saying* “I scared away the birds” would have misled his audience if he did not intend to suggest (by I-inference to the stereotype) that the birds flew away. This suggestion is due of course in part to the way the world is (most kinds of birds fly); but that would only warrant a probabilistic inference. The default nature of the inference, to the surety that the birds flew away,

it would be good to help the poor, she should be motivated to help the poor). This propriety condition is, on their view, what allows us to say that moral claims ‘express’ motivational states, even though the semantic function of such claims is to express beliefs (because the sentences used to make them express—in the semantic sense of ‘express’—propositions).

The connection between moral claims and motivation is of course a matter of considerable debate. Bar-On and Chrisman’s view is very like a number of views classed as examples of ‘hybrid expressivism’ (cf. my fn. 5), which attempt to combine the best of cognitivism and noncognitivism by giving moral claims dual content: ordinary propositional content, which is supposed to account for such claims’ assertoric nature, and content that is indicative of the speaker’s motivational state, which is supposed to account for the strong connection between making a moral claim and being motivated to act in accordance with it. (I am using the term ‘dual content’ loosely—as I’m pretty sure I do throughout the paper, as well as such terms as ‘secondary content’—to include conversational implicature content as well as various kinds of semantic content.) Bar-On and Chrisman’s view differs from these in that they do not think that the way a moral claim expresses the speaker’s motivational state is a function of the claim’s semantics; nor do they think it is the effect of conversational implicature. Although they call their view ‘neo-expressivism’, their propriety condition for moral claims is not a linguistic norm, it seems.

Hybrid expressivist views are similar to the view I am propounding here, in that I am also presenting a kind of dual-content view for aesthetic (and other) claims. However, hybrid expressivists are generally keen to strip the primary content of a moral claim of any normative element, trying to explain what is normative about moral claims entirely by reference to its secondary motivational content. I do not think this normative purge is the right approach for the primary content of moral or aesthetic claims (more below on the primary content of aesthetic claims). I also think that the secondary content of moral claims is a more complex matter than is indicated by both traditional expressivism’s and hybrid expressivism’s exclusive focus on motivation. The secondary content I propose for moral claims is content concerning the ‘direct’ source of the speaker’s claim, thus conveying that it is made as a result of her own moral assessment. This is clearly different from, but, I think, complexly related to, the idea that moral claims express motivational states.

relies on the presumption that *otherwise we would have been told!* That is, given a mutual presumption about how we use stereotypes to augment the communicational load of utterances ... the effect will be to convert Bayesian probabilities or hazardous assumptions into the certainties of communicational presumptions. (2001: 53)

The bare outline of this kind of inference certainly fits our case with the Direct Evidentials: it is taken for granted that a simple unmarked assertion in this class is made on the basis of the speaker's direct grounds for it, and if the speaker does not have direct grounds, we expect to be told so. If this presumption is based on I-inference to the stereotype, however, it is not clear why this is the case, given that, in general, most assertions appear to carry no evidential implications. Why would a direct evidential source be the stereotype in the case of the Direct Evidentials?

We have already seen that the reason for the presumption in the aesthetic case is not straightforwardly epistemic: there is no general epistemic bar to acquiring aesthetic knowledge on the basis of testimony, and so there is no reason to presume on general epistemic grounds that a simple aesthetic assertion must be based on the speaker's direct experience. The same is surely true of the other Direct Evidential domains. If I say, "Margot's house looks orange under the new lights they've just installed on her street," I indicate that I have seen her house under the new streetlights myself. But the proposition that Margot's house looks orange under the new streetlights is clearly something I may come to know if someone tells me that it is so.

However, even without such a complete epistemic embargo, there is another possibility involving the epistemic status of Direct Evidential claims, suggested by Levinson's remarks on I-inference to the stereotype: maybe Direct Evidential beliefs are *typically* acquired firsthand, leading to the I-Principle presumption that a simple, unmarked Direct Evidential assertion is based on the speaker's direct grounds.

In my previous discussion of *must*, I indicated some of the complex ways in which both the type and the quality of one's evidence (what kind of evidence it is, on the one hand; and how likely, on its basis, a given claim is to be true, on the other) can be conveyed by a speaker's choice between modalized and unmodalized assertions. If I say that it's raining, I don't thereby convey anything much about the evidential source for my claim (although context may of course make this clear). If I say that it must be raining, I convey that my evidence is indirect; and, because of the context in which I say it, I also manage to convey whether or not this indirect evidence is conclusive.

There are of course other kinds of examples. Say I am watching a replay of the Wimbledon final. The time difference between me and London makes it inconvenient for me to watch the match live. I have, however, gone out of my way to avoid news reports and other tennis fans throughout the day, so as not to hear who has won. Now someone annoying wanders into the room and says, "Oh! Federer lost, you know." When the smoke stops shooting out of my ears, I ask whether it was close, and he says, "No idea. I didn't watch it. I haven't heard anything about it either. But Federer must have lost—he's such an old man now!"

Someone who says in this context that Federer lost certainly does rule out a particular evidential route for his having learned that fact: the bare statement that Federer lost is infelicitous—as it is for my annoying friend—if the speaker is simply making a prediction of the outcome based on general knowledge of the players involved. However, this kind of prediction is never actually sufficient for knowledge of the outcome of a tennis match (or any sporting event), as there is always room for an upset. The reason, then, that "Federer lost" rules out predictive inference as one's evidential source for the claim is that this kind of inference cannot

provide knowledge in such a case, and assertion without knowledge is inappropriate. My friend manages to be misleading in two ways: he asserts something he doesn't know to be the case, and, because he indicates that he does know Federer lost, he also indicates that he learned the outcome of the Wimbledon final directly, or through testimony, or possibly through some kind of indirect inference concerning what happened after the match (maybe the Swiss Embassy rolls out a Hooray Federer! parade whenever he wins a major championship).

In some cases, then, because knowledge of p can only be obtained through certain evidential routes, assertion of p will indicate those routes as the only possible sources of the speaker's knowledge. Standard lottery cases are also like this. No matter how vanishingly small (yet greater than zero) your chance of winning, I do not have sufficient grounds to say, "Your ticket lost," unless I have had epistemic contact, either directly or through testimony, with the lottery's actual outcome.

I offer these examples in contrast to the Direct Evidentials, which do not exhibit the same feature: it is just not the case that direct evidence is the only kind of evidence that provides knowledge of the truth of a Direct Evidential claim. The task now is to tell a plausible GCI story about how simple, unmarked Direct Evidential claims nevertheless convey that they are made firsthand, even though firsthand evidence isn't the only kind of evidence that gives knowledge of such claims.

I will start with a different example, outside the Direct Evidential domain, that I take to illustrate a similar story. The example is from Ninan, who uses it in discussing the issue concerning the assertibility of *must* ϕ in cases where the assertion of ϕ would be inappropriate:

Suppose A and B are friends with a couple, Carl and Diane, who have been dating for a long time and are likely to get married at some point in the future. Suppose that, prior to the following dialogue, B has not heard any recent news concerning Carl and Diane's relationship. Now consider:

- A: Carl proposed to Diane yesterday!
B: At last! She must have said "yes".
B: ?? At last! She said "yes". (2014: 16)

I agree with Ninan that B should not assert the bare "She said 'yes'." But why not? A tempting answer is that B can't know that Diane said 'yes' given the merely predictive information she has about Diane's proposal-accepting behavior. B thus uses *must* to indicate the kind of inference she is making. This answer would put Ninan's example in the same class as the tennis or lottery examples: we would be holding, I suppose, that human behavior in non-routine circumstances like these ultimately has the same kind of unpredictability as a tennis match or a lottery, precluding knowledge of outcomes based on merely predictive information.⁶⁷

However, I think it would be overly scrupulous to insist that B could not have knowledge of Diane's answer in such a case. We may imagine a slightly different setup to make this clearer, in which B has had a very recent conversation with Diane, and Diane has gone on at length about how much she loves Carl and about her plan to propose to *him*. Of course, there are doubts that

⁶⁷ Prediction of human behavior in *routine* situations is often decent grounds for simple, unmarked assertions about that behavior, and is presumably thought to provide knowledge of that behavior. If I am asked where my husband is, and have nothing else to go on but that he left the house at the usual time this morning, I may say, "He's at work." I don't have to qualify this in any way ("He must be at work" etc.).

could be raised about Diane's answer in such a case—but such doubts can be raised for any example of what we take to be ordinary knowledge that is not ironclad with certainty.

I think, then, that B can have predictive knowledge of Diane's answer,⁶⁸ but she still may not legitimately assert the bare "She said 'yes'" unless she has heard Diane say it or has been reliably informed that she said it. I want to suggest, now, that B's bare assertion of "She said 'yes'" would be illegitimate because of the evidential information it would convey: it would convey that B bases her claim on direct observation or testimony. Unlike for the tennis or lottery cases, however, this is not because these evidential sources are the only possible sources of B's knowledge. It is rather, I suggest, that these sources are *typically* what ground people's knowledge, and hence their assertions, in situations like B's, where the conversation is about what a person did in novel circumstances. It is relatively uncommon in such situations for this kind of inferential evidence to be sufficient for knowledge. Employing Levinson's framework, we may say that in B's situation, and others like it, assertion of a simple, unqualified sentence suggests by I-inference to the stereotype that it is made on the basis of direct observation or testimony.

5.2 The role of perspective: some examples with 'looks' (and one or two other things)

How, then, may we take beliefs in the Direct Evidential domains to be typically grounded by the subject's own experience, so that we may suppose I-inference to the stereotype to be behind the evidential content of Direct Evidential assertions? I said above that there is perhaps a sense in which Direct Evidential claims seem 'subjective.' What I had in mind was one element of what non-philosophers usually mean by this: that they are made from a particular perspective, a perspective that is non-standard in some way, either because it requires one to be in a special position to have it, or because it reflects a kind of discernment, imagination or judgment that we know is likely to vary from person to person in its upshots, in ways that it is often particularly difficult to adjudicate.

Perspective is relevant to the different Direct Evidential claims in different ways. I have thus far mostly avoided providing any positive account of the primary content of aesthetic (or other Direct Evidential) claims, beyond expressing allegiance to *Experience Independence*. I have put aside expressivist accounts of aesthetic utterances, and subjectivist accounts for aesthetic claims (in the philosophical sense of 'subjectivist') that would insert the speaker's aesthetic experience into the primary content of her claim. However, beyond this, I have not argued for a specific account of aesthetic primary content, and so it seems (for anything I've said) that alternative contextualist views, together with relativist and objectivist views, are all on the table.⁶⁹ My account of aesthetic claims' secondary content, and that of other Direct Evidential claims, now makes it desirable to say at least a little bit about the relationship between such claims' primary content and the secondary content that I want to attribute to a kind of GCI.

⁶⁸ Note, for what it's worth, that it would be fine for B to say to A, in discussing the matter prior to Carl's actual proposal: "Diane will definitely say 'yes' when he asks her!" Of course, making an assertion like this about the future has its own evidential implications—it is obviously not possible to have direct or testimonial evidence for what is claimed—and so it might be that such assertions have the same kind of greater risk-tolerance that affects *must*-ified assertions. On the other hand, though, B could also say, "I know Diane will say 'yes'!" (cf. fn. 38).

⁶⁹ Gibbard-inspired expressivist views for the aesthetic domain whose semantics rely on 'aesthetic standard acceptance' rather than aesthetic experience would also presumably be on the table (cf. §2.1). However, as such views are in many ways similar to relativist views, and the differences do not, as far as I can tell, have any bearing on the issues I discuss here, I will continue to put them to one side. (Cf. fn. 4.)

It is surely the case that the different kinds of Direct Evidential claims are not all best served by the same account of their primary content. Even just focusing on ‘looks’ claims—as I will do here for a moment—reveals a variety of kinds of content and hence of different ways that such claims give rise to the presumption that they are made as a result of the speaker’s own experience.

Predicates like ‘looks red’ can be used in ways that suggest some kind of implicit relativization to specific viewers and viewing conditions. Say we’re at the science museum. If I say “It looks red” when I have funny goggles on, it is presumably understood that I mean that it looks red to me, under my current conditions. You are standing next to me with different goggles on; and you are not taken to be disagreeing with me when you say, “It looks green.” In these cases, it is plausible to suppose that our claims concern how it looks to each of us in these conditions: “It looks red to me with these on”, “It looks green to me with these on.”

In this kind of context, it is simply obvious that we say what we do as a result of what we each in fact see at the time. Such examples do not, I think, clearly illustrate the Direct Evidential nature of the ‘looks’ claims themselves, as it is apparent from the context itself—regardless of any evidential information conveyed by the claim—that the claim is made on the basis of the speaker’s experience. However, context necessarily has less to say about this when a ‘looks (to me)’ claim is made to communicate a thing’s general disposition to look a certain way to the speaker, rather than how it looks to the speaker at the time she makes the claim. Now we’re at the paint store, as you want to paint your house. You are considering shades of red, as you like Margot’s red house, but you want to steer clear of orange. I say, “Actually, Margot’s house looks orange to me under the city’s new streetlights.” We both know that Margot’s house does not *currently* look orange to me under the streetlights—we are nowhere near Margot’s house, and it is daytime—and my claim is apparently a report of a disposition (to look a certain color to me under certain conditions). In general, dispositions do not have to have been manifested in order to be truly attributed to a thing: the vase need not ever have broken, say, in order for it to be true that it is fragile, or for it to be appropriate to say that it is. And I could easily come to know that Margot’s house has a disposition to look such-and-such to me without my ever having seen it at all—if, say, someone else tells me that it looks that way to her, and I know we are relevantly similar. Yet my claim clearly indicates that I have seen Margot’s house look orange under the streetlights myself; it would be odd, and misleading, for me to make the claim if I hadn’t.

Grice’s second Maxim of Quantity (“Do not make your contribution more informative than is required”) is clearly relevant here. Why bother specifying that X looks such-and-such *to me*? Presumably at the least I am hedging against the possibility that the way it looks to me is different from how it looks to others. But although it is surely sometimes possible for me to know how something is disposed to look to me, even though I haven’t seen it, and even when I think I may not be a typical viewer, it is overwhelmingly likely that the reason I know how something looks to me when I suspect I am idiosyncratic in this respect is that I *have* seen it. I have singled myself out as possibly different from others, when I could have singled out somebody else, a particular group (e.g. the people over there with funny goggles on; the red-green colorblind), or no one at all; the overwhelming implication then is that I make my claim on the basis of what I myself have seen.

However, it is not clear that a ‘looks’ claim must be relativized to a particular viewer or group of viewers, either explicitly or implicitly, although it is on the face of it at least plausible

that for color ‘looks’ claims this could be the case. If I say, “Margot’s house looks orange under the new streetlights,” context may make it clear that I mean her house generally looks orange in that circumstance, and on the contextualist construal I canvassed above for the ‘looks’ comments at the science museum, my claim here would be taken to be implicitly relativized to normal viewers. But many ‘looks’ claims are not so amenable to this treatment, even when normal human vision is in some way at issue for them. It is not the case, for instance, that “It looks like a Vermeer” will come out equivalent to “It looks like a Vermeer to normal observers.” It is not a basic capacity of the human visual system, shared by all those with normal vision, to be able to detect Vermeer lookalikes.⁷⁰

Yet both claims—“Margot’s house looks orange under the new streetlights” (to normal observers) and “The painting looks like a Vermeer”—are true or false independently of any special perspective or capacity of the speaker; both are in fact plausibly objective claims. About Vermeer-lookalike-ness, for instance, Charles Travis, whose example it is, says: “When would a painting not, in fact, look like a Vermeer? When, for example, the woman sweeping the courtyard is using a type of broom not made until the 18th century, or is wearing a hat never worn in Holland; or when the brush strokes are a bit too broad to be Vermeer’s, or feather in the wrong way, or the pigments are a bit off.” (2004: 78) Someone with normal vision and sufficient knowledge of Vermeers can often tell definitively that a painting does, or does not, look enough like one to be in the running for being one.⁷¹

‘Looks’ claims generally do suggest that they are made as a result of what the speaker has in fact seen for herself, though, even when, like these, they are plausibly objective. And accounting for the Direct Evidential nature of objective claims seems to present a special difficulty. If the content of a claim is not only independent of any particular experience the speaker has had, but is also independent of any capacity or disposition particular to the speaker, why should such a claim implicate that it is the speaker herself who has direct grounds for it? After all, as I have noted, saying that something *is* red generally conveys no evidential information. Why then is ‘*looks* red’ different, even when, however this is worked out in detail, it is a thing’s general look (under certain conditions) that is at issue? The question is particularly salient for anyone who favors an objectivist construal of any of the Direct Evidential domains—of aesthetic matters, of course, or of ethical matters, where objectivism has been an even more popular choice.

No doubt the purported ‘calculation’ in question will be different for different kinds of case. The first thing I want to note about this problem, though, is that I have already given an example of how a straightforward objective claim may carry evidential information: I argued in the last

⁷⁰ Nor is it plausible that there is an equivalence with “It looks like a Vermeer to Vermeer experts.” Either it is possible for Vermeer lookalike-ness to come apart from Vermeer expertise (even the experts are not wise to certain visible changes in Vermeer’s pigments over time, say) or it is not. But if it is not, this must be because possession of Vermeer expertise is made contingent upon the ability to detect visually whether or not a painting (really) looks like a Vermeer.

⁷¹ Note that the ease of attributing objectivity to “It looks like a Vermeer” depends on it being clear from the context that there is a highly *specific* standard at issue for Vermeer lookalikes: the standard here is that of an art dealer, say, gauging whether, visually, a painting could pass for a Vermeer. But other uses of the same sentence are not clearly associated with a specific standard: “It looks like a Vermeer!” you cry happily, about your latest effort in an amateur painting class. (Travis is clearly not thinking about this kind of use when he says what I have quoted above.) The relevant standard in this case is obviously lower than the art-dealing standard—no one, including you, wonders whether your painting could fool an expert—but it is not just lower, it is also much vaguer. Just how much visual resemblance to a Vermeer is needed to make your claim true? In some sense this seems an odd question; and that in itself suggests that there is something substantially different about these two uses of “It looks like a Vermeer.”

section that B's utterance of "At last! She said 'yes'" would have conveyed that B had direct or testimonial grounds for her claim. This is because such grounds are typical for the kind of knowledge that B expresses in making her claim.

The truth of a color 'looks' claim is independent of how things look to any particular observer when it is the class of normal observers that is relevant to it; color 'looks' claims of this type differ from (e.g.) 'is red' claims only in their reference to non-standard viewing conditions. However, the non-standard condition in the claim about how Margot's house looks under the streetlights is precisely what makes it different, evidentially speaking, from the 'is red' claim. The presumption is that it is the speaker herself who has witnessed the state of things in the special circumstances at issue: not because that is the only way she could have come to know how things are in those circumstances, but because *typically*, for these everyday sorts of observations, it is the speaker for whom those circumstances have been relevant—typically, she knows about them because they were in her life. Of course, they may not have been—but if not, she will say so. I want thus to suggest that a simple, unqualified assertion about the color something looks to be, to normal viewers but in special circumstances, conveys by I-inference to the stereotype that it was the speaker herself who saw it to look that way.

The situation for Travis's example of "It looks like a Vermeer" is different. The speaker is not talking about how the painting looks in non-standard conditions; she is talking about how it looks, in general, overall: it has the same look (to very high, art-dealing standards) as a painting by Vermeer.⁷² For this kind of claim, the calculation becomes more subtle, and actually may not materialize at all. Whether it does or not depends on a slew of details about the conversational context concerning what is known about the thing in question and the speaker's relationship with it. If I am a crooked art dealer and say to my partner, about a painting just finished by a certain highly skilled person in our employ, "It looks like a Vermeer," my partner will infer that I say so on the basis of having seen it myself—even if she had no prior knowledge that I had gone to see it (and even, we may add, if she knows that many other people have seen it, and are likely to have told me about it). But consider, now, that this highly skilled person becomes famous for her forgeries (while, alas, rotting in jail), and there is a (somewhat ironic) exhibition of her work. You ask a friend—a Vermeer fan—to go with you. As an inducement, you say, "One of the paintings looks like a Vermeer." Here there is not necessarily any suggestion that you have seen the painting in question before.

I take my use of "It looks like a Vermeer" (as crooked art dealer, to my partner) to illustrate an implicature phenomenon that actually has nothing to do with appearance claims *per se*. I have said that the claim is plausibly objective, and that its being *prima facie* objective depends on the fact (cf. fn. 71) that it is clear in this case which specific standard for Vermeer lookalike-ness is in play: the claim is true if, visually, the painting can pass for a Vermeer—if there is nothing about the painting's look that rules out its being a Vermeer. But although the standard relevant to the claim is clear, *applying* the standard in this context takes great nicety of judgment: not everyone—hardly anyone, in fact—can tell whether the painting really looks like a Vermeer

⁷² Note that the speaker need not think the painting *is* a Vermeer: she could be giving her verdict on what she knows to be a highly competent forgery. "It looks like a Vermeer" may sometimes be used to communicate that the speaker thinks, on the basis of its look, that it is a Vermeer, but the sentence need not be used in that way. These different uses correspond to what M. G. F. Martin (2010) calls the 'epistemic' use—"Going by its look, (I judge that) X is Y"—and the 'comparative' use—"X visually resembles Y." I have been assuming the comparative use in the text: i.e. that the speaker wishes to communicate only that the painting's look sufficiently resembles the look of a Vermeer.

to this degree. Where the necessity for this nicety of judgment is known, it is presumed that the speaker has judged for herself: who else's judgment would it be, given that it is not clear whose judgment can be trusted? We get the same presumption for a slightly different claim, not about appearances, in a slightly different context: my morally upstanding twin, who works for Sotheby's, comes into a meeting and says, of a painting recently found in an old lady's attic, "It's a Vermeer, all right."⁷³ The lack of any implication for "One of the paintings looks like a Vermeer," when you are talking to your friend about the exhibition of forgeries, is thus due to the fact that the painting's look is now well-known to be like that of a Vermeer.

We have now seen several examples of 'looks' claims—including plausibly objective claims—that suggest they are made on the speaker's direct grounds, and I am arguing that this is a species of implicature. I have identified several different reasons for these claims to carry the implicatures that they do: the content of the claim is relativized to the speaker, who is known or suspected to be idiosyncratic in some relevant respect; the claim requires knowledge of special circumstances, which it is assumed the speaker has herself witnessed; the claim is known to require unusual skill in judgment, which it is assumed the speaker has herself exercised. I want now to turn to a further large class of 'looks' claims that belong to the Direct Evidentials for a different reason.

I will return to the example from fn. 71, in which I say, "It looks like a Vermeer!" about my own amateurish painting. It is not clear that of two possible responses to this—"You're dreaming" and "I kind of see what you mean"—one is obviously better than the other. You might think that this is because I haven't described the situation in enough detail. But even if I embellished the story, we would still be able to imagine both responses being appropriate.

The difference between this use of "It looks like a Vermeer" and the other uses I have discussed is, as I suggested in the footnote, the vagueness surrounding whatever standard of similarity is relevant to my claim. I presumably make the claim because I am struck by some similarity between my painting and Vermeerishness; but although, if pressed, I should be able to articulate what kind of similarity this is, it seems that others may *agree* that my painting is similar to a Vermeer in this respect *without* agreeing that my original claim—"It looks like a Vermeer"—is at all compelling. Even I do not take it, for instance, that others must agree with me because my picture is recognizably of Delft, or has an air of quiet intimacy, or is painted in minutely rendered detail—even if, when challenged, I offer any one of these things as my reason for saying what I did, and even if they agree with me that my painting has these features. They may agree that my painting looks like a Vermeer to the extent that it has any one of these features, without agreeing that my original claim is compelling. (They may likewise agree that my painting looks like a Vermeer in looking rectangular, without taking this to be a compelling reason for saying "It looks like a Vermeer" in the context at issue.)⁷⁴

⁷³ These issues are subtle, though. If there *is* someone else whose judgment it would obviously be—possibly controversial though the judgment is—it might be assumed that my ground for my claim is her testimony. Say I am not an authenticator for Sotheby's but an accountant. In that case, everyone in the meeting knows I have no particular expertise in Vermeers, so they assume I have been talking to someone in the Authentication Department.

⁷⁴ Martin describes this phenomenon by saying "there is a potential for subjectivity" at the heart of comparative 'looks' claims: "... typically for those appearance statements with a comparative element we are inviting the audience to make a comparison When one talks of the look of an object inviting a comparison by the audience, one indicates that one finds the objects similar, given one's visual experience, and invites the audience to do so, too." (2010: 214)

“X looks like Y” (and just “X looks Y,” where ‘Y’ is an adjective) is similar to much-discussed sentence-forms involving gradable adjectives like ‘tall’ or ‘rich,’ in that one thing can look like something else (or look a certain way) to a greater or lesser degree. ‘Tall’ and ‘rich’ must be associated with a comparison class, provided by context though not necessarily explicitly stated, in order for sentences using them to have a determinate truth value: ‘tall for a 1st grader,’ ‘rich for an American.’ This is not true of ‘looks like Y,’ because the scale associated with this kind of expression is bounded by the way Y looks—nothing can look more like Y than Y itself—but the scales for height and wealth are potentially infinite. Even when the expression at issue is ‘looks like a Vermeer’—rather than, say, ‘looks like Vermeer’s *View of Delft*’—the relevant scale is presumably bounded, though in a less well-defined way.⁷⁵ Of course, we may nevertheless include a reference class if we want to: “You look a lot like Margot for someone who’s not related to her.”

However, ‘looks like Y’ is like ‘tall,’ ‘rich’ and gradable adjectives generally in requiring there to be some ‘cut-off point’ on the relevant scale for determining whether or not the predicate should be applied to a thing: is 4 ft. 3 in. tall (for a 1st grader)? Does having a million dollars make someone rich (for an American)? Does my painting look like a Vermeer?⁷⁶

⁷⁵ Another aspect of the adjectival expression ‘looks like a Vermeer’ is that things that are not Vermeers may be higher on the ‘looks like a Vermeer’ scale than things that are. You might think, for instance, that an excellent copy of *The Milkmaid* looks more like a Vermeer than Vermeer’s actual *View of Delft*, given that the former exhibits an intimate, indoor subject much more typical of Vermeer’s oeuvre.

⁷⁶ According to a common classification, ‘tall’ and ‘rich’ are *relative* gradable adjectives: what it takes for the adjective to apply to a thing in its positive form (“X is tall”) varies according to context. Relative gradable adjectives are usually associated with *unbounded* scales: e.g. ‘tall’ and ‘short’ map objects to a numeric scale of heights that can be infinitely large or infinitely small. A comparison class is essential to determining the truth conditions for “X is tall” as, without one, there is no clear rationale for putting the cut-off in any particular region of the scale. *Absolute* gradable adjectives are associated with bounded scales, either at one or both ends. ‘Bent’ and ‘straight,’ for instance, are associated with a scale of increasing bentness that is bounded at the lower end by complete straightness. ‘Empty’ and ‘full’ share a scale that is bounded at one end by complete emptiness and at the other by complete fullness. Absolute gradable adjectives do not have scales with contextually sensitive cut-off points (although they may be subject to contextually sensitive standards of precision) and are thus not dependent on comparison classes. Their cut-off points are fixed at one or other end of the scale. “X is straight,” for instance, is true of X only when it is completely straight (‘straight’ is a *maximum* standard absolute adjective); “X is bent” is true of X when it has any degree of bentness at all (‘bent’ is a *minimum* standard absolute adjective). ‘Empty’ and ‘full’ are both maximum standard adjectives.

Some gradable adjectives with bounded scales, however, nevertheless have patterns of usage that make them best categorized as relative. ‘Looks like Y’ has a partially closed scale like that for ‘bent’—assuming for the moment that Y is well-defined, its scale is one of increasing visual similarity to Y up to the fixed point of looks-identical-to-Y—but it has a contextually variable cut-off point. In this it is similar to ‘bald’—although ‘bald’ presumably has a fully bounded scale, from completely hairless to full hair coverage. The fact that ‘bald’ is bounded at the upper end in this way, rather than stretching off into the infinitely hairy distance—or even the infinite number-of-hairs-per-unit-area distance—presumably provides something of a clue as to why it, like ‘looks like Y,’ is an exception to the rule that expressions with closed scales should be absolute. ‘Bald’ and ‘looks like Y’ are expressions with *perceptual* scales. This is obvious for ‘looks like Y’ but not so obvious for ‘bald,’ given the apparent relevance of objective measures like number of hairs per square inch. It is clear, though, that a man with a full head of hair is not really ‘balder’ than the average sea otter—which, at a million hairs per square inch, is the most densely hairy animal on earth. A man with a full head of hair sits at the upper end of the baldness scale. The relevance to the relative/absolute distinction, of having a perceptual, rather than an objectively measured scale, presumably has to do with the kind of transitions on the scale that can be significant to us. There is no perceptual significance to the transition from minimum hairiness (complete hairlessness) to having-one-hair, nor from maximum hairiness (full coverage) to full-coverage-minus-one-hair: we care about what *looks* bald to us in various ways. Similarly, our

The cut-off point can be constrained in various ways. It is very highly constrained for my use of “It looks like a Vermeer” when I am pronouncing judgment on whether it can *pass* for a Vermeer. A height of 3 ft. is definitely not tall for a 1st grader; Donald Trump is definitely rich for an American (although possibly not as rich as he’d like us to think). But there can be a sizable gray area on a gradable term’s associated scale, in which different people may position the cut-off in different places, even in the same conversation. That is, even though a particular conversational context presumably constrains a cut-off point to a significant degree, the people whose conversation it is may still disagree about who counts as, say, rich—for no other reason than that they position the cut-off at different points of the scale (we may imagine, for instance, that they are arguing about someone whose net worth in US dollars is known to both). No one is obviously mistaken in many such disputes, which have thus come to be seen as examples of the kind of “faultless disagreement” that has been associated with matters of personal taste and other matters that have inspired relativist theories.⁷⁷

Luckily, I do not have to weigh in on these semantic debates, either here or for matters of personal taste, which, as I’ve said, also exhibit Direct Evidential behavior. We can explain why such claims give rise to Direct Evidential implicatures just by focusing on the kind of phenomena of usage that I have here sketched for gradable terms. I began the discussion of gradable terms by talking about a particular ‘looks like’ claim: “It looks like a Vermeer,” said about my amateurish painting. I hope it is not surprising at this point that it turns out many claims involving gradable terms indicate that they are made as a result of the speaker’s own experience or judgment, when the claim sails in the term’s gray area—when, that is, the context does not sufficiently constrain the position of the cut-off point in a way that makes it clear it is mutually agreed upon, and the claim cuts close to the cut-off’s possible locations. When someone makes a claim of this kind, she must take a further stand on where she thinks this cut-off is: she must constrain its position to a greater extent than context has already done. It is naturally presumed, in such a case, that she takes this stand as a result of her own judgment.

This is the case even for unidimensional, measurable adjectives like ‘tall’ and ‘rich,’ although their scales involve just one quantity (height, wealth) and this quantity is objectively measurable. Each adjective’s mapping of objects to points on the scale is thus objectively fixed, and so the ordering of objects on the scale is fixed too.⁷⁸ But while this renders comparative claims (“X is taller than Y”) clearly objective, the simple positive claims (“X is tall”) are subject to cut-off issues. Much of the time, for unidimensional, measurable adjectives such as these, the cut-off (given a mutually understood comparison class) will not be an issue for speakers, and in such cases no Direct Evidential implicature will arise: Shaquille O’Neal is tall (even for a

interest in visual likeness can take many forms, and is not limited to an interest in the difference between strict visual identity and the lack of it. (See Kennedy 2007. Kennedy mentions *bald*’s exceptional nature (his fn. 30) but these speculations about it are my own.)

⁷⁷ Note that we seldom have much of an *explicit* idea about where our cut-off points are for different adjectives. They manifest themselves on the fly, and are presumably shaped by general knowledge, personal experience, and the demands of the conversational context. Note, as well, that people generally have reasons for putting their cut-offs where they do, and that when a disagreement of this kind has any robustness, it is usually because these reasons have practical significance. Who counts as rich, for instance, has significant bearing on public policy. One person may therefore persuade another to re-position her cut-off for such-and-such reasons. Someone who takes this kind of disagreement to be faultless, however, insists that such a choice cannot be considered the correction of a mistaken belief about the rich. See e.g. Mark Richard (2004, 2008) for a relativist account of gradable adjectives.

⁷⁸ ... although this is admittedly more of an idealization for ‘rich’ than for ‘tall,’ given the uncertainties involved in assigning cash value to various non-liquid assets.

basketball player); the Queen of England is rich (for a human being on planet Earth). My reasons for thinking these things are entirely testimonial—I don't know exactly how tall O'Neal is, or how much wealth is in the Queen's coffers—but I am willing to make the assertions without qualification. In some contexts, however, the location of the cut-off relative to the object at issue will not be clear, and then it will be assumed that the speaker has exercised her own judgment in making her claim.

'Looks like Y,' however, is presumably a multidimensional expression: one thing can look like another in different ways; for 'looks like a Vermeer,' multidimensionality is presumably even more striking, as its endpoint is not well-defined. Multidimensionality of course complicates the ordering of items on a scale: different people will not necessarily agree about the ordering on a multidimensional scale (though they know perfectly well how scalarity works), as they may weight the importance of the relevant dimensions differently. When the ordering of items on the scale as well as the cut-off point are not necessarily fixed in the conversation, the likelihood of Direct Evidential implicature is even greater.

We may illustrate this phenomenon with a class of Direct Evidentials that I mentioned earlier: attributions of personality traits like intelligence or shyness. It was not immediately clear to me why these should be part of the group. Many such claims are entirely uncontroversial and seem objective enough: "Albert Einstein was intelligent," for instance. It seems to me that someone who hasn't the faintest idea what Einstein actually accomplished could make this claim without suggesting, misleadingly, that it is her own considered judgment (given the magnificence of the General Theory of Relativity) that Einstein was intelligent. But then, Einstein's intelligence is famous, and there is no danger at all, in most contexts, that the claim "Einstein was intelligent" could be sailing in *intelligent's* gray area.⁷⁹

This is generally not the case for our day-to-day evaluations of people. Human personality traits show up against a background of complex behavior, and are generally multidimensional: 'intelligent,' for instance, has this feature, as there are many different ways a person can be intelligent. Our judgments in this area are beholden to the particular glimpse we've been allowed of the behavior of the person judged, and to our more general experience of people, which affects to a significant degree both the weight we give to different aspects of the trait in question and where we are likely to place the trait's cut-off. These aspects of our judgments vary a good deal from person to person, and this makes it incredibly likely, in everyday conversations about these things, that someone making a claim of this kind about someone else's personality will have to take a stand on the relevant trait's cut-off relative to that person. In general, then, such claims are presumed to be based on the speaker's own exercise of judgment, to a much greater degree than for unidimensional properties.

What about my claim that my art-class painting looks like a Vermeer? What is interesting about this kind of claim is that a speaker takes a stand on *looks like Y's* cut-off—which could, in different contexts, really be just about anywhere on *looks like Y's* scale—as a result of a particular experience: the experience of having been struck by some kind of similarity between the thing in question and Y. We can be struck by all kinds of whimsical visual comparisons: "Your boss looks like Richard Nixon"; "Margot's new dress looks like a wedding cake." (We can be struck by all kinds of comparisons whatsoever, whimsical and visual or otherwise.) Other people will agree with us about these things when they not only see the similarity in question, but

⁷⁹ There could be some contexts in which it would: maybe one of Einstein's wives often grumbled to her friends about how her idiotic husband always left the fridge open, or whatever.

are likewise struck by the aptness of the comparison. When the ‘looks like Y’ cut-off is not obviously constrained by other factors—as it is for “It looks like a Vermeer” when I am interested in its passing for one—it is taken to be constrained by the speaker by her own experience of visual similarity. Earlier, I mentioned a similar case: Pettit’s example involving the sadness of a picture (cf. fn. 16). I said that someone who says that the picture is sad, as a result of her experience of it, in large part serves to set the standard for the picture’s sadness in the conversation, rather than merely locating the picture on the sadness scale relative to an independent, given standard.

A claim of this kind thus does not seem to be aiming for some truth about a thing that goes beyond the particular experience that gives rise to it (which is why such claims come across as *invitations* to others to share our experience, rather than as laying down the law about how things are). This is typical of claims that are not objective; yet, as we saw for “It looks like a Vermeer,” if the constraints on the relevant standard of comparison are tight enough, we get a ‘looks like’ claim that does seem objective, and if I am trying to determine whether a forgery can pass visually for a Vermeer, any particular experience I have of its look does not settle this question.

5.3 The role of perspective: personal taste

We have now seen many different reasons for Direct Evidential claims to carry the implicatures that they do. To recap: the content of a claim may be relativized to a speaker who is possibly idiosyncratic in a relevant way; a claim may require knowledge of unusual circumstances, which it is assumed the speaker has witnessed; a claim may require unusual skill in judgment, which it is assumed the speaker has herself exercised; and, as shown by the latter examples, a claim may ride close to a gradable term’s cut-off point, for terms with either unidimensional or potentially more idiosyncratic multidimensional scales, which it is assumed therefore reflects the speaker’s own positioning.

What of claims of personal taste—of what is tasty, or fun (to use a couple of the best-known examples)? Such claims (almost?) always carry Direct Evidential implicatures, as far as I can tell; they certainly carry them more often than the other examples I have discussed (the last ‘looks like a Vermeer’ kind of example aside). Both contextualists and relativists acknowledge that simple claims of personal taste are governed by the speaker’s own standards of taste: I am committed to calling a food ‘tasty’ just in case it tastes good to me, and committed to calling an activity ‘fun’ just in case it is fun for me.⁸⁰ This is a *commitment*: if (say) a food tastes good to me, it is not merely *permissible* for me to believe and hence to assert that it is tasty, but it is in some sense *obligatory* that my standard of taste dictate my belief, and hence any assertion I make, about the tastiness of the food. Something has gone wrong with me if I find marshmallows delicious but claim to be agnostic about their tastiness, for instance.⁸¹

As I said earlier, though (cf. §2.3.2), it is possible in some cases to know how a thing measures up to my standards even when I haven’t experienced it. So it is possible for a claim that something is tasty (say) to be *true* in the absence of the speaker’s having actually tasted it—

⁸⁰ Objectivism about matters of personal taste is a difficult position that attracts (almost?) no one. As e.g. MacFarlane (2014: 3) points out, if there is just one standard for these things that is ‘correct,’ the linguistic license that I have just described makes no sense, given the wide variation in what people find tasty, fun, etc.

⁸¹ John MacFarlane gives the following examples of bizarre-sounding speeches that likewise show agnosticism to be off-limits whenever the speaker’s own tastes have delivered a strong verdict: “I’m not sure whether espresso is tasty, but I hate how it tastes”; “I’ve never been able to stand the taste of durian. Might it be tasty?”; “I love orange juice and hate tomato juice. But who knows? Maybe tomato juice is tastier.” (2014: 4)

on either a contextualist's or a relativist's construal of how this truth works—as long as the speaker's standard of taste is one against which the thing measures up favorably. However, *typically*, a speaker knows how a thing measures up to her standard of taste because she has in fact tasted it. The possibility of variation in what is considered tasty from person to person, together with variation in the tastes of superficially similar foods, make it relatively uncommon for us to acquire knowledge of how things measure up to our standard of taste through testimony or inference. If I want to convey my belief about how a food measures up to my standard of taste, when I have not in fact directly measured it myself, I will not use a simple, unqualified form: “That pie must be delicious—Margot says she's never tasted better,” or “That pie must be delicious—everything Margot bakes is wonderful.”

We can see the more radical subjectivity at work for claims of personal taste, compared to most of the examples we have looked at so far, by considering the nature of the scales associated with personal taste adjectives. For instance, although ‘intelligent,’ as I said, is multidimensional, and is therefore bound to yield slightly different orderings for different people, intelligence is still arguably at least somewhat objectively *measurable*: not necessarily because of the existence of IQ tests, which certainly do not measure every kind of intelligence there is, but because we have a fairly good intersubjective grasp on the kinds of things that require intelligence and on how people measure up to these kinds of tasks.⁸² But while personal taste terms like ‘tasty’ and ‘fun’ are gradable (things can be more or less tasty, more or less fun), they exhibit more interpersonal variation in application than multidimensional terms with at least some degree of measurability.

I take it, by the way, that ‘tasty’ (say) is not in fact multidimensional, in spite of a tastiness judgment's dependence on multiple aspects of a food's flavor that are themselves dimensional. Someone can be intelligent in this or that respect: good at chess, having good business acumen, socially intelligent; judging someone to be intelligent *tout court* in the context requires the person to be intelligent ‘enough’ in some one respect or many respects together. But although a food can be tasty in parts (the pie has a tasty filling but not a tasty crust), the dimensions of a food's taste that are relevant to a judgment of its tastiness—sweetness, bitterness, saltiness and sourness, and (less traditionally) ‘savoriness’ (umami) and fattiness, together with the contributions of our sense of smell that manifest subjectively as taste, not to mention food texture and temperature—are not individually judged for tastiness and then summed in some contextually appropriate way. Nor is it the case that greater tastiness is necessarily correlated with having a greater degree of any of the qualities that are measured on these scales. It is rather

⁸² Louise McNally and Isidora Stojanovic (2017) make this point. Sibley, in contrasting ‘intelligent’ with aesthetic terms, arguably makes the point as well. He says that although there are presumably no general, finitely storable necessary and sufficient conditions for being intelligent, ‘intelligent’ is nevertheless condition-governed in a looser way: “[W]ith concepts of this sort, although decisions may have to be made and judgment exercised, it is always possible to extract and state, from cases which have *already* clearly been decided, the sets of features or conditions which were regarded as sufficient in those cases. These relevant features which I am calling conditions are, it should be noted, features which, though not sufficient *alone* and needing to be combined with other similar features, always carry some weight and can count only in one direction. Being a good chess player can count only *towards* and not *against* intelligence. ... [W]hat I want particularly to emphasize about features which function as conditions for a term is that *some* group or set of them *is* sufficient fully to ensure or warrant the application of that term. ... There are individuals possessing a number of such features of whom one cannot deny, cannot but admit, that they are intelligent.” (1959: 425-426)

that the different dimensions of a food's taste come together in taste experience to be judged, in combination—in the way that they balance each other—as tasty (or not).

The way that judgments of tastiness relate to these non-evaluative flavor scales is thus somewhat similar to the way that color judgments relate to the properties of hue, saturation, brightness, etc.—although we are of course less conscious of the effects of these color properties on our color judgments than we are of the relevance of different flavors to our tastiness judgments. Color terms are not used in a multidimensional way, in spite of the dependence of color judgments on multiple dimensional properties, and the same is true of 'tasty.'⁸³

However, 'tasty' shares with multidimensional terms like 'intelligent' a difficulty with calibrating a single scale for it, even for just one person. It can be hard to make comparative judgments about intelligence when two people are intelligent in very different ways. Similarly, comparative tastiness judgments can be hard to make unless the difference in tastiness between two foods is very large (this quite dry chocolate cake versus that incredible tempura) or involves foods of the same kind (this pretty good chocolate cake versus that even better chocolate cake that has been improved by use of bittersweet rather than semisweet baking chocolate).

For tastiness, though, even large differences will not be reliably consistent from person to person: what is tasty according to one person may be disgusting according to another. There is no objective measurability for tastiness, as there arguably is for intelligence, that makes different people's scales look roughly the same at a squint. Two people's 'intelligent' orderings may be slightly different from each other, but their 'tasty' orderings may be *radically* different.

This of course makes it even harder for people to coordinate their tastiness judgments with those of others than it is for them to coordinate their intelligence judgments. 'Tasty' doesn't just have cut-off issues (how tasty must X be, to be thought or called tasty in the context?) but relies on an idiosyncratic scale—not itself multidimensional, but dependent on multiple dimensional aspects of flavor—that is not objectively measurable, and is only *subjectively* measurable in a pretty crude way.

⁸³ E.g. we don't say a thing is "red in every way except for saturation" or that a food is "tasty in its bitterness but not its sweetness," although we can say that someone is book-smart but not street-smart. (I was alerted to the point for color by Galit W. Sassoon in what I take to be a prior draft of (2013). The reference to the grammar of color terms was taken out in the published version, but I take it the point is still a good one...)

6. THE CONTENT OF BEAUTY CLAIMS: FIRST PASS

Are aesthetic claims just like claims of personal taste—a subclass, even, of personal taste claims? Lots of philosophers have seemed to think so, perhaps without thinking about the matter too much.⁸⁴ A cursory survey of my non-philosopher friends likewise finds agreement here amongst the folk. And there are of course various observations to pile up on the “more or less alike” side of the scale. Aesthetic claims are like claims of personal taste in concerning the pleasures of experiences of various kinds: tastiness concerns the pleasure of gustatory taste, fun is about the pleasure of engaging in various activities, visual beauty (say) concerns the pleasure of looking at things. Moreover, aesthetic predicates are like personal taste predicates in showing what may be radical variation in application from person to person. And both domains inspire the thought that there is “no disputing” about them, but also inspire significant critical practices in which disputes may after all be played out. (There is food criticism as well as arts criticism.)⁸⁵ And if, as a result, aesthetic claims should be understood in the same way as personal taste claims—for which I have been assuming some kind of contextualist or relativist theory is correct—then their Direct Evidential implicatures should be understood in the same way too.

There is, however, some reason to think that they are not exactly the same. As we saw earlier, a kind of difference is argued for strongly by Kant, who insists that calling something ‘agreeable’ (a personal taste claim if ever there was one) and calling it ‘beautiful’ are fundamentally different: the latter, but not the former, being a normative claim that calls for agreement from everyone. In saying that something is beautiful, that is, I say that others ought to share my pleasure in it and thus agree that it is beautiful. A claim that something is agreeable (or tasty, etc.) does not in this way ask either for others’ pleasure or their agreement in judgment. Kant’s claim about the agreeable is thus in harmony with the norm I have just discussed, which enjoins that each person is to use her own standard of liking for making simple claims about matters of personal taste.

The nature of the standard relevant to judgments of beauty, however, is not so clear, if they are normative in the way that Kant describes. And as this may affect calculability issues for claims about beauty and other aesthetic matters, I want to spend some time here. The standards relevant to making normative claims are not straightforwardly related to the speaker’s actual dispositions to respond to things, in the way that the standard relevant to making a tastiness claim is simply read off from the speaker’s disposition to like the taste of a food (or not). I can make a normative judgment (and hence hold a normative belief, or make a sincere normative claim) while at the same time acknowledging that I do not always, or will not always, live up to

⁸⁴ E.g. Andy Egan (2010) makes almost no distinction between aesthetic matters and matters of personal taste (although he does say that some aesthetic disputes are more ‘robust’ than disputes about gustatory taste). My vague impression is that Kantians (of course) think that they are not alike, but that contemporary philosophers of language have either thought that they are, or simply have not wanted to get bogged down in whatever differences there might be. There is some new interest, however, in aesthetic language, as distinct from the language of personal taste, among linguists and linguistically-minded philosophers (e.g. Liao, McNally and Meskin 2016; Liao and Meskin 2017; McNally and Stojanovic 2017).

⁸⁵ MacFarlane (2014: 147) notes that epicureanism about food is consistent with the norm I described above, that we call foods ‘tasty’ just in case they taste good to us. A speaker may even herself acknowledge that her standards of taste are ‘deficient’ in various ways—she may know she lacks experience with certain kinds of foods, say—yet she is still bound to call foods ‘tasty’ according to this norm. The existence of food criticism is likewise consistent with the norm.

it in my actual responses. I have of course included moral claims in the class of Direct Evidentials; and it is generally just taken for granted that simple moral claims depend on the exercise of the speaker's own judgment. If aesthetic claims are normative, this will presumably not be taken as a challenge to their Direct Evidential status either: it simply seems natural for normative claims to be part of the group. However, I want to investigate in detail why this is the case, given their less than straightforward relationship to speaker response, and the viability of objectivist accounts of normative claims. I have already given some examples of objective claims that carry Direct Evidential implicatures, so it would not be impossible to explain why normative claims, if objective, carry them too; but there would presumably be different reasons for this.

My discussion will follow Kant in arguing that judgments of beauty are normative—and I will also agree with him that they are not in fact objective, though they are similar in some ways to objective judgments. The Direct Evidential status of aesthetic claims will thus ultimately not seem surprising in any way. Yet my discussion will, I hope, also show why this status is particularly strong for aesthetic claims: why, that is, the speaker's experience, which is the basis of a simple aesthetic claim, is so hard to disentangle from its content, and why the Direct Evidential implicatures of aesthetic claims do not seem to be cancelable. This is similar to how it is for personal taste claims, but may be surprising if we take aesthetic claims to be normative. I do not, for instance, think that moral claims show quite the same tenacious link between their content and the fact that the speaker generally makes the claim as a result of her own judgment: I think it is easier to find examples of 'contextual' cancelation for moral claims—i.e. cases in which the Direct Evidential implicature simply doesn't arise.

As I said above, however, anyone who disagrees about the normativity of aesthetic claims, and takes them to be similar in all relevant respects to personal taste claims, already has an acceptable calculation for the Direct Evidential implicatures of simple aesthetic claims.

6.1 The normativity of beauty

For this discussion I am going to focus on beauty, which I take to be paradigmatic of the issues concerning aesthetic judgment and normativity. I will approach an argument for the normativity of judgments of beauty by following a route that Kant mentions but does not, as far as I know, follow up on himself. This route concerns the norm I have described, which rules out agnosticism for matters of personal taste when a person's own standards are clear about the matter. It seems that agnosticism is often *not* ruled out for an aesthetic judge in similar circumstances—that we may not only suspect 'deficiency' in our own aesthetic standards (as we may likewise suspect in our gustatory standards—cf. fn. 85), but that our suspicions in the aesthetic case may actually license suspension of aesthetic belief. Note that my investigation here requires that I focus on judgments made firsthand.

Kant appears to allude to the suspension of aesthetic belief here:

The judgment of others, when it is unfavorable to our own, can of course rightly give us reservations about our own ... (*CJ* §33, 5:284.)

And for some more stalwart classical backup (though from a different camp), here is the beginning of Hume's 'Of the Standard of Taste,' which appears to allude to the same kind of thing:

The great variety of Taste, as well as of opinion, which prevails in the world, is too obvious not to have fallen under every one's observation. Men of the most confined knowledge are able to remark a difference of taste in the narrow circle of their acquaintance, even where the persons have been educated under the same government, and have early imbibed the same prejudices. But those, who can enlarge their view to contemplate distant nations and remote ages, are still more surprized at the great inconsistency and contrariety. We are apt to call *barbarous* whatever departs widely from our own taste and apprehension: But soon find the epithet of reproach retorted on us. And the highest arrogance and self-conceit is at last startled, on observing an equal assurance on all sides, and scruples, amidst such a contest of sentiment, to pronounce positively in its own favour. (Hume 1965: 275)

However, the point is perhaps a tricky one. We might think that Hume's finally enlightened blowhard, in scrupling to pronounce positively in his own favor, is agnostic not about particular aesthetic judgments he has made, but is instead merely suspicious about the quality of his aesthetic standards—the standards he used to make the judgments. This is no different from what is possible for gustatory taste, where, as I've said, we may also suspect our standards to be inferior in some way. With Kant, too, earlier in the passage from which I've quoted, he mentions the possibility of exactly this: if someone is indifferent to a building, a view, or a poem whose beauty other people praise highly, he may “begin to doubt whether he has adequately formed his taste by acquaintance with a sufficient number of objects of a certain kind.” In spite of what I think is the most natural reading of the later excerpt, quoted above, this might make us want to interpret it in the same way: perhaps we have ‘reservations’ about our capacity for forming aesthetic beliefs, rather than the resultant beliefs themselves. (‘Judgment’ is used to refer both to the capacity for judging and to what this capacity produces.)

It becomes clear, though, that Kant is not saying that here. He wants to emphasize that the testimony of others concerning something's beauty cannot force my own judgment to go the same way—and, further, can never be grounds for changing my mind—and he finishes the sentence like this:

The judgment of others, when it is unfavorable to our own, can of course rightly give us reservations about our own, but can never convince us of its incorrectness.

The last part of the sentence puts it beyond doubt that (“our own”) ‘judgment’ is being used here to refer to a particular judgment arrived at, not the general capacity for judgment itself. Our capacity for judgment can be good or bad, but not incorrect. We may have ‘reservations,’ then, about particular aesthetic judgments that we have made—at least, according to Kant in this passage, and I think he is right.

What do such reservations amount to? Do they give us the kind of agnosticism that is ruled out for personal taste? Note that the kind of agnosticism at issue here is a highly specific kind of doxastic neutrality: what is ruled out for tastiness (say) is S's suspension of belief about whether X is tasty, when S believes that X tastes good *to her* (or disgusting to her, etc.). But there are many other situations in which we may suspend belief about something's tastiness, and not just because it is entirely unfamiliar. The taste of a mouthful on a specific occasion does not all by itself make true the various beliefs about tastiness that we may be inclined to acquire on its basis. We acquire these beliefs against the background of other beliefs we have about the food in question, and the beliefs we have about our standards of taste. I may come to have reservations about previous tastiness judgments I have made if I suspect that I have misjudged the nature of

this background. This can happen in various ways, for judgments about particulars as well as for generalizations.

For instance, I may be inclined to think that my husband's tofu stir-fry is pretty missable, taste-wise, but I can have many reasons for doubting this judgment. Maybe he's become a better cook since I last tasted it: if so, my belief about "my husband's tofu stir-fry" overgeneralizes. (I would in that case retreat to the belief that my husband's tofu stir-fry used to be not very tasty.) Or maybe my tastes have simply changed, in which case I have overgeneralized about *them*. (I would retreat to the belief that I used to think my husband's tofu stir-fry was not very tasty.) Maybe someone tells me I'm supposed to eat it with sriracha; this may affect not just the belief about my husband's tofu stir-fry in general but also a belief about the stir-fry I am eating right now. Maybe in my first bite of the stir-fry I am eating right now, I get a big lump of garlic that somehow made it intact into the wok, and think, "Yuck," until I realize that the taste of my current mouthful isn't representative of the taste of the dish as a whole. Or maybe I am eating this stir-fry—and have, coincidentally, eaten all previous stir-fries—with a cold.

Generic claims about the tastiness of food and drink usually convey that they concern the thing's normal way of being eaten or drunk. (If the sriracha example is not entirely convincing, consider someone who says, "Coffee is disgusting," having only ever tasted a version made by stirring coffee grounds into cold water.) And simple claims about tastiness, whether generic or particular, generally convey that they are made on the basis of the speaker's current, but normal, standards of taste: not on the basis of what used to be normal for her, nor on the basis of what is current but abnormal (as when suffering from a cold).

It is of course possible to have reservations about aesthetic judgments that are analogous to the tastiness examples here—both aesthetic generalizations and judgments about particulars. I might say of a contemporary poet that her poems are just depressing—no beauty anywhere—without having read the more recent work that manages to find beauty after all in all the gloom. Or maybe my poetic tastes have changed since I made my original judgments about her poems. Or, concerning one of the poems in particular, maybe I did not originally read it in either the psychological or external conditions most suitable for revealing its beauty (maybe I saw it written up on a bus station bathroom wall after I had been sick).

Yet it is clear that the aesthetic 'reservations' alluded to by both Kant and Hume do not concern these kinds of issues. Both philosophers refer to the opposing judgments of other people as the source of a person's reservations about his own judgments. I may read a particular poem in a comfortable chair with a good cup of tea, and it may leave me cold. But someone whose aesthetic judgment I generally trust may tell me that it is beautiful, and this may make me wonder what I might have missed. I may come to suspect in myself any number of aesthetically relevant prejudices or deficiencies. Even though the poem left me cold—and perhaps still does—I may come to reserve judgment about its beauty, or lack of beauty. This is precisely the kind of agnosticism that is ruled out for tastiness: if, after generalizing in all of the appropriate ways, I think that X tastes bad to me, I cannot be left wondering whether X is tasty, even if someone whose tastiness judgments I generally agree with has told me that it is.

If I eat something when I have a cold, I can certainly still wonder whether it's (actually) tasty or not; I may be inclined in that moment to say that I still don't know how it tastes. So it may seem as if my suspension of judgment about the poem is analogous to my suspension of judgment about the taste of something when I have a cold. In both cases, I suspect that there is something about me that prevents me from judging accurately. But there is an important

difference. As I have said, having a cold denies me access to my *usual* gustatory tastes; but an aesthetic prejudice or deficiency denies me access to aesthetic taste that is *better* than the taste I actually have, whether or not better taste is the norm for me. I do not suspend aesthetic judgment because I think that I will return to aesthetic tastes that are normal for me, once the relevant prejudice or deficiency is removed. I suspend aesthetic judgment because I think the prejudice or deficiency prevents me from judging a thing *according to its merits*, whether or not judging well in this way is something that has ever been habitual to me. What counts as accuracy in the two cases, then, is different.

I take it that this phenomenon is in alignment with Kant's insistence that judgments of beauty are normative, but judgments of the agreeable are not—although I am not aware of him discussing the matter precisely in this way.⁸⁶ It is important, though, to see, as Kant did, the unique combination of features exhibited by judgments of beauty: on the one hand, they have an essential connection with pleasure or liking, just like judgments of the agreeable; on the other hand, their normative element, which I have just illustrated, argues for their having 'universal validity'—which is to say that in making a judgment of beauty, we make a claim to others' agreement, just as for cognitive judgments.

But because judgments of beauty, like judgments of the agreeable, are made on the basis of pleasure or liking, they are nonconceptual: they are not based just on the recognition that the objects in question satisfy any given concepts. That is, unlike for cognitive judgments, there are no properties of the objects in question—apart from their giving rise to the pleasure or liking—that dictate their possession of the properties of beautifulness or agreeableness.⁸⁷ Although Kant

⁸⁶ However, the issue may be connected with Kant's ideas about 'pure' and 'impure' judgments of taste (cf. *CJ*, §13 and §14). Henry Allison (2001:107-110), for instance, suggests that Kant thinks we make 'erroneous' judgments of taste (*CJ* §8, 5:216) when we are wrong in thinking that we have managed to set aside those factors of our liking that are either merely personal (the agreeable) or determined by reason (the good). We think, that is, that we have made a pure judgment of beauty, when in fact the pleasure on which the judgment is based is adulterated by merely personal liking, or by the kind of approval that results from thinking that the thing in question is good, either for some purpose or in itself. If this is right, then presumably, for Kant, our possible reservations about our judgments of beauty may simply be doubts about the purity of those judgments. Similarly, the prejudices and deficiencies I describe as possible hindrances to judging well aesthetically could be, for Kant, hindrances to making pure judgments.

Kant's distinction does seem significant to aesthetic life. We often do come to think that aesthetic judgments we had thought to be unassailable were 'merely' the result of personal liking or dislike, possibly conditioned by our particular cultural context. Such a thought is plausibly what occurs to Hume's formerly arrogant, self-conceited gentleman.

⁸⁷ This is of course very different from the wholly implausible doctrine that concepts aren't employed *at all* in aesthetic experience. (Such a doctrine about the experience of food would be equally implausible, for that matter.) Several features of Kant's account, besides the main talk of aesthetic judgments being 'nonconceptual,' do perhaps encourage the idea that, for him, aesthetic experience is entirely concept-free: e.g. his comments on the difference between pure and impure judgments of taste; and his explanation of pure judgments in terms of the free play of the cognitive faculties, which does not employ 'determinate concepts' and reflects merely the conditions for 'cognition in general.' But Kant does not say or ever even (really) suggest that aesthetic experience employs no concepts. Any experience of objects *must* employ concepts, as he tells us in that other book. See e.g. Christopher Janaway's (1997) discussion of this issue; in particular, about the free play of the cognitive faculties, he says:

The free play of imagination and understanding was never meant to constitute the totality of any experiential episode. *S* is perceiving *o*, perhaps in a complicated, changing environment, in which *o* must first be identified as an object ... and then fastened upon with sufficient stability for the free play of the cognitive faculties to occur and the characteristic pleasure to be felt. These features demand that *S* is

does not emphasize this, we may of course give reasons, either way, for thinking something tasty, or beautiful: “The acid tang of the citrus balances the richness so well”; “He finally allows himself a moment of specific tenderness for his wife, as unredeemable and hopeless as he is.”⁸⁸ But no such reasons can be ‘proof’ of beauty or tastiness; these qualities are inevitably underdetermined by reasons; it is only the existence or possibility of the relevant pleasure or liking that can form the basis of any judgment of either.

We have seen for tastiness what is plausibly similar for the agreeable more generally, that it is each person’s own liking, or lack of it, that is the relevant standard for her beliefs and claims about what is tasty, regardless of the level of refinement of her palate. The twist for judgments of beauty, illustrated by the phenomenon I described above concerning the nature of our possible reservations about such judgments, is that we do not take our own liking or lack of it (or our own propensities for these) to be definitive in the determination of what is beautiful. Yet at the same time it can only be liking—pleasure, or the possibility of pleasure, in the beautiful—that forms the basis of such judgments. Well then—whose liking, whose pleasure, is relevant, in those cases in which we are skeptical about our own responses? Presumably nobody’s in particular: rather, perhaps, the liking of some abstract person who judges *well*, whose response is sensitive to the object in the right kind of way—as our own, we suspect, is not. My reservations about my aesthetic judgment indicate an apparent awareness, not present for the agreeable, that an object merits a certain response, which my own perhaps falls short of. The possibility of such reservations thus makes plausible Kant’s contention that judgments of beauty, unlike judgments of the agreeable, are normative judgments that make a claim to universal validity: if a beautiful object merits a response of pleasure, then everyone ought to feel pleasure in response to it, and, thus—given the relationship between such pleasure and a judgment of beauty—to agree that the object is beautiful.

Not, of course, that everybody will. An object’s meriting a certain response does not all by itself make it the case that everyone responds to it in that way. Color judgments, which are also based on a subject’s perceptual response, of course have universal validity too, but they are based on the recognition that the visual response called for (as it were) by a red object (say) in good light is determined by the capacities of normal color vision. Ordinary objective judgments—including judgments about the colors of things—are governed by relatively clear standards. Judgments of the agreeable, which do not have universal validity, are nevertheless governed by relatively clear standards too: everyone is to take their own liking to be the standard for agreeableness. In contrast, the standards for judgments of beauty are not clear. I do not necessarily take my own liking, or lack of it, as the standard for what is beautiful; but when I do not, I also cannot specify which standard should be used instead. (Compare someone with an eye disease, who suspects her color responses are skewing yellow. She is now unsure, on the

operating with concepts in experiencing *o*. Indeed, by Kant’s own lights *o* must fall under some concepts for *S* if it is to become an object of *S*’s experience at all. (472)

Janaway argues not only that aesthetic experiences that give rise to pure judgments of taste employ concepts, but that Kant’s account leaves room for the idea that increasing one’s conceptual sophistication can enhance one’s ability to make aesthetic judgments: “Kant’s analysis of the experience required when an object is judged in a pure judgement of taste allows for the education of aesthetic responses by conceptual learning, and the absurd idea that knowing nothing could improve one’s appreciation of art is not an idea of Kant’s.” (476)

⁸⁸ ... as I might say about Macbeth’s “She should have died hereafter” speech.

basis of her own vision, whether certain things are blue or whether they're green. If she really wants to know, though, she only has to ask.)

6.2 An objectivist view for beauty?

I have been discussing this matter in the context of trying to establish the calculability of Direct Evidential implicatures for aesthetic claims. I said earlier that if aesthetic claims have the same underlying structure as personal taste claims, then we can understand their implicatures similarly. I have been arguing, however, that there is an important difference between the two kinds of claims. The question now concerns the upshot of this difference.

It is of great concern to Kant of course to explain how judgments of beauty could be possible in the first place, given their particular combination of nonconceptual response-dependence with normativity. How could anyone's pleasurable response to an object—given that it is not constrained by conceptual awareness of the object's other properties—be normatively required of everyone?⁸⁹ However, his answer to this question in terms of the 'free play' or 'harmony' of the faculties of imagination and understanding is notoriously hard to understand, and in fact I know I do not understand it and I will have nothing to say about it. I am myself going to put this question of "how possible?" to one side, to see what happens to the calculability issue when we take the normativity of aesthetic judgment, and hence its universal validity, seriously. This requires going some way to answering the question of what kind of judgments these would be, if normative, rather than how they are possible.

Kant himself says that judgments of beauty, because they are nonconceptual (in the way I have spelled out) and based on the subject's pleasure, are not objective: their "determining ground cannot be other than subjective." (*CJ*, §1, 5:202) Thus far, they are like judgments of the agreeable. Note that Kant insists for judgments of both the beautiful and the agreeable that the relevant 'subjective ground' for judgment must be the subject's actual experience of pleasure. In an earlier section (cf. §2.2), I discussed Hannah Ginsborg's interpretation of Kant's remarks on this issue; she maintains that, for Kant, the subject's experience of pleasure is part of the content of her judgment, which clearly would make Kant's a subjectivist view. I have been arguing throughout that, if this is in fact Kant's view, he has not got it quite right: the subject's experience of pleasure is not part of the primary content of her claim, but is instead presumed to be the ground of her claim through a species of implicature. Nevertheless, if we shift Kant's emphasis from the subject's actual experience of pleasure to her *disposition* to feel pleasure—what I take to be the basis of her *standard* of beauty, or agreeableness—we would, I think, still respect the non-objectivist aspects of his view by fitting it (as it stands so far, the normativity of beauty aside) to a different kind of contextualism or relativism, for which it is the subject's standards that are relevant to assessing the primary content of her claim.

While this move to a standard form of contemporary non-objectivism might be fair enough for judgments of the agreeable, trouble comes if we try to manhandle normative, universally valid judgments of beauty into a contextualist or relativist shape. The point of these models is to deny, for the kinds of claims for which they are appropriate, that a single standard governs their truth. Yet a single standard is apparently what is needed to govern the truth of judgments thought to be universally valid. It seems far easier, then, to jettison Kant's non-objectivity for

⁸⁹ "How is a judgment possible which, merely from **one's own** feeling of pleasure in an object, independent of its concept, judges this pleasure, as attached to the representation of the same object **in every other subject**, *a priori*, i.e., without having to wait for the assent of others?" (*CJ* §36, 5:288)

beauty—and hence contextualist and relativist views—and to take it instead that, however this turns out to be possible, the normativity and hence universal validity of aesthetic judgments demand that they be in some sense objective.

I mentioned earlier a view of this kind, which has been suggested by John McDowell (cf. my §2.3.2 and fn. 21). He says that taking “the phenomenology of value experience at face value” requires taking values to be “part of the world” (1998a: 129-130)—which is to say, objective in some sense.⁹⁰ This is compatible, he argues, with the plausible claim that experiences of value, including aesthetic experiences, are like experiences of the traditional secondary qualities—e.g. color experiences—in being “unintelligible except as modifications of a sensibility like ours.” (1998b: 143) (To put this in relation to Kant’s views on beauty is to emphasize the essential connection between a judgment of beauty and the feeling of pleasure on which it is based.) Yet part of the face value of aesthetic and other experience of values, he stresses, is that it is unlike secondary quality experience in being experience of objects that do not merely cause but *merit* their responses.⁹¹

I will take McDowell’s remarks to suggest the following view about the beautiful: claiming that X is beautiful attributes to X the dispositional, yet normative property of being such as to merit a certain kind of pleasure.⁹² Note the similarity with the property that we saw posited by Ginsborg and Rind in §2.2 to account for judgments about beauty that do not take the simple, positive form (i.e. judgments about beauty other than “This is beautiful,” like negative judgments and generalizations). In the paper of Ginsborg’s I discussed earlier (1998), she objects to the view suggested by McDowell’s remarks as a *general* view of beauty, because she thinks it does not take proper account of how the simple, positive judgments are made. I can know that something merits the pleasure appropriate to beauty without having experienced it myself—when, say, I am told so by a reliable informant—and so, on this McDowellian view, I can both know of something’s beauty without having experienced it, and go on to make claims about its beauty on the basis of this knowledge. Ginsborg’s objection is that she thinks we do not in fact make simple, positive beauty claims without firsthand experience, contrary to what we would expect from the McDowellian view. We saw in §2.2, however, that any force the objection has against taking the simple, positive judgment to attribute this kind of objective property to an object, can also be brought to bear against taking such a property to figure in negative judgments or generalizations, which are also made on the basis of firsthand experience. I have been arguing that the objection misses its mark in both cases: that any beauty claim that ‘weighs in’ on an object’s beauty carries a Direct Evidential implicature to the effect that it is made on the basis of the speaker’s firsthand experience, and that speakers may use alternative evidential forms to

⁹⁰ On the face of it, we also take properties like tastiness to be part of the world; yet we do not consider them to be objective. I take McDowell ultimately to be (at least implicitly) distinguishing properties like tastiness from aesthetic and ethical properties through the idea of *merit*, and the relationship between the critical practices appropriate to judging merit in objects and how these practices are reflected in our talk about tastiness, beauty, rightness, etc.

⁹¹ My exposition of McDowell here is something of an amalgamation of views he expresses in ‘Aesthetic Value, Objectivity, and the Fabric of the World’ (1998a) and ‘Values and Secondary Qualities’ (1998b). The latter is, as the title conveys, about value in general, but his examples in that paper suggest that he primarily has moral rather than aesthetic value in mind there. I think that his overall argument, however, is applicable to both (and that he himself takes it to be).

⁹² I do not quite want to claim that McDowell himself endorses this view in these papers. His remarks focus on the possibility of objectivity for values, and leave many details of the possible accounts for various values unspecified.

convey claims about beauty that are not based on such experience. The McDowellian view, then, as a general view of claims about beauty, survives Ginsborg's objection.

It is also important to note that the McDowellian view is not necessarily at odds with the nonconceptuality of judgments of beauty. We may still acknowledge that the merits of a beautiful object are, at bottom, only identifiable through feeling the kind of pleasure in the object that is felt in making the firsthand judgment. This leaves only a couple of ways to judge the merits of a potentially beautiful object when we haven't felt pleasure in it ourselves. We can be told that it is beautiful by a reliable informant, or we can infer that an object that is physically identical to an object we know to be beautiful is itself beautiful. (The further we get from physical identity, the more speculative our inferences will be. If in general I admire the beautiful pieces of a certain artist, I may think that one of her pieces that I haven't yet seen is very *likely* to be beautiful as well.)

It is perhaps not obvious why such a view would count as objectivist. The objectivity at stake here is very different from what we get with paradigmatic objective matters, including the objectivity of colors, which is underwritten, as I have said, by relatively clear, specifiable standards. There is no such concrete, specifiable standard that is available for beauty: the concept of the response merited by a beautiful object is of a sort of abstract point to be aimed at, something like a mathematical limit (though not nearly so well-defined), rather than of a specifiable standard of practical use against which beauty can be measured. What arguably makes the view nevertheless objectivist is that the idealization involved means that there is in some sense just one standard that is relevant to the assessment of everybody's claims about beauty: not a concretely specifiable standard, but rather a reflection of the fact that everyone's claims are held to critical account in the same way.

6.3 Critical standards

If we allow, with McDowell, that this can be a kind of objectivity (perhaps we should prefer a slightly weakened term, like 'quasi-objectivity'—although that prefix has already been conspicuously adopted elsewhere), the difference between it and that of paradigmatic objective domains nevertheless clearly comes out in the relatively high level of disagreement about aesthetic and other normative matters. It seems that the job of an objective standard ought to be to govern its subject matter with something of an iron hand, and the critical standard I have just described for beauty clearly does not do that. There would presumably also be wide variation in the opinions of a random sample asked to judge, or at least guess, whether a particular painting looks like it could have been painted by Vermeer; but the applicability of a stringent standard to this question is widely understood, and so the random sample will generally defer on this issue to the expertise of the few hundred people in the world actually qualified to judge. We do not see this kind of deference about beauty; and though an aesthete is someone who is particularly (perhaps excessively) devoted to beauty, we do not call an aesthete an *expert* on beauty as such, any more than we call an epicure an expert on deliciousness as such, though they may be experts on the Dutch Golden Age and French haute cuisine. The relatively high level of disagreement about aesthetic matters is not smoothed out by deference to expertise.

Part of the explanation for this difference can be provided from within the McDowellian objectivist view, by focusing on the difference between how typical objective standards and critical aesthetic standards are identified and applied. In the Vermeer case, the reasonable non-expert acknowledges that whatever standards the expert uses to make her judgment about the

look of the painting are the appropriate standards to use. The non-expert generally has some idea about what these standards are—in order to look like a plausible Vermeer, a painting must not be anachronistic in dress or furnishings, must look to have pigments that were used in Dutch art of the 17th century, etc.—but is vague about the details, and expects the expert both to know what these details are and to be adept at applying the detailed standards to the matter at hand, using appropriate methods (knowing what to look for under a microscope, for instance).

But the critical aesthetic standard is simply a reflection of the fact that criticism of aesthetic judgment is equally appropriate for everyone. This makes it something of a moving target; moreover, it is a moving target that each person constructs and develops *for herself* throughout her life of aesthetic engagement. This means that there is not only variation in aesthetic judgment, but variation in what are considered to be good aesthetic standards.

I take McDowell to be making this point in what follows. He expounds his view as part of a project to defend the ‘reality’ of values. He draws an analogy between values (say, given my discussion, ‘the beautiful’) and ‘the fearful,’ which he uses as a simple example to illustrate the basic contours of the kinds of objective, normative, yet response-dependent properties he suggests values are. He says that trying to make sense of fear

must allow for the possibility of criticism; we make sense of fear by seeing it as a response to objects that *merit* such a response, or as the intelligibly defective product of a propensity towards responses that would be intelligible in that way. For an object to merit fear just is for it to be fearful. So explanations of fear that manifest our capacity to understand ourselves in this region of our lives will simply not cohere with the claim that reality contains nothing in the way of fearfulness. Any such claim would undermine the intelligibility that the explanations confer on our responses. (1998b: 144)

He acknowledges an important disanalogy between fearfulness and values, however: i.e. the much greater variation there is in response to value than there is in our fear responses (values are “contentious”). He continues:

What the disanalogy makes especially clear is that the explanations that preclude our denying the reality of the special properties that are putatively discernible from some (broadly) evaluative point of view are themselves constructed from that point of view. (We already had this in the case of the fearful, but the point is brought home when the validation of the responses is controversial.) However, the critical dimension of the explanations that we want means that there is no question of just any actual response pulling itself up by its own bootstraps into counting as an undistorted perception of the relevant special aspect of reality. Indeed, awareness that values are contentious tells against an unreflective contentment with the current state of one’s critical outlook, and in favour of a readiness to suppose that there may be something to be learned from people with whom one’s first inclination is to disagree. The aspiration to understand oneself is an aspiration to change one’s responses, if that is necessary for them to become intelligible otherwise than as defective. But although a sensible person will never be confident that his evaluative outlook is incapable of improvement, that need not stop him supposing, of some of his evaluative responses, that their objects really do merit them. He will be able to back up this supposition with explanations that show how the responses are well-placed; the explanations will share the contentiousness of the values whose reality they certify, but that should not prevent him from accepting the explanations any more than (what nobody thinks) it should prevent him from endorsing the values. (1998b: 144-145)

The fact that not just any actual response to beauty can pull itself up by its own bootstraps, reflects our awareness of the critical dimension relevant to beauty, and is revealed in practice by (among other things) the possibility of having reservations about our beauty judgments. Crucially, though, the fact that our critical explanations of such judgments are constructed from within our own aesthetic point of view, means that each person's aspirations towards aesthetic idealization—towards improved aesthetic engagement with things—are developments of her own aesthetic standards. And because of the variety of people's aesthetic responses, and the multiplicity of different ways in which people's aesthetic engagement may be developed, even standards that have been cultivated will not necessarily converge.

Because the kind of idealization at issue develops from each individual's aesthetic point of view, we must choose for ourselves the people we consider to be good judges, in a way that is essentially different from our acknowledgement of experts in other areas. I do not believe someone to be a Vermeer expert because I have compared her judgments concerning Vermeer with my own and found hers to overlap with mine. I would need to be an expert on Vermeer myself for this to be reasonable. But for me to believe a person to be a good judge of beauty requires thinking there to be some kind of sympathy between my own judgments and hers. This does not necessarily mean I have compared some of our actual judgments; I may believe you to be a good judge of beauty simply because a mutual friend (whose judgment I already trust) has told me so. But at the end of any such testimonial chain there will be some comparison of this kind.

6.4 A problem with the objectivist view

This seems a good enough account of the way we rely on others aesthetically when we have not had the opportunity of judging for ourselves: we choose our aesthetic informants carefully. We may even account for the Direct Evidential implicatures of aesthetic claims on this basis, as the relative scarcity of this kind of aesthetic authority for a particular person may make it very likely that any aesthetic belief someone comes to hold will have been acquired firsthand. But sometimes there is overwhelming and well-known testimonial weight in favor of something's beauty—this is particularly common for natural beauty—that cuts across aesthetic subcultures. If a judgment of beauty were a judgment of merit as on the McDowellian view, we might then wonder why the Direct Evidential implicature is so robust for aesthetic claims. I mentioned earlier that claiming that Einstein was intelligent probably does not carry any particular evidential implicature in many contexts, as Einstein was so *famously* intelligent: many people presumably make the simple claim that Einstein was intelligent purely on hearsay. But this is not the case for claiming that the Taj Mahal, or the *Mona Lisa*, or Hawaii, is beautiful. These things are famously beautiful, however, and there are not many contexts in which anyone wants to argue that they lack the merits appropriate to beauty. Yet it is still presumed that these simple claims suggest the speaker has firsthand experience of their beauty.

In fact, there remains a difficulty for the McDowellian view, even after acknowledging that our judgments about the responses merited by various objects are subject to the idiosyncrasies described by McDowell. This difficulty becomes apparent when we consider the contrast between judgments about a merited response to an object and our actual disposition to respond to it. We may judge that an object must merit a certain response without being disposed ourselves to respond to it in that way. The contrast is of course starkest when we have *actually* responded to an object in a way that is different from what we think the object must merit.

Evaluative judgments about food and drink may certainly conflict with our actual dispositions to respond to them. Judgments of tastiness or deliciousness are not normative in the way that judgments of beauty appear to be, as I have discussed; so the judgment that a certain dish is impressive, say, has not been in the running as a stand-in for the judgment that the dish is delicious. We would expect, then—as long as it is possible for someone to acknowledge that a dish is impressive even when her own response to its taste is not positive—that such judgments would come apart from judgments of tastiness, deliciousness, etc. This surely is possible: this kind of value is surely what epicureanism is all about, and, as mentioned above, epicureanism does come apart from such judgments (cf. fn. 85). This is because these taste judgments are tied, not to considerations of epicurean value, but to our actual dispositions to respond to a food's taste.

For instance, as a judge in a cooking competition I may certainly think that a great chef's dish is an impressive dish, and in a way merits a delighted taste response, even though the dish contains ingredients that I do not like and is one that I do not find tasty at all. I may say, "This is an impressive dish. It must be delicious, if you like that kind of thing." In this case, of course, 'delicious' has been explicitly relativized to the group of people who, say, like bouillabaisse (as I do not). I could not simply have said, "This is an impressive dish. It must be delicious," because the standard of taste relevant to the *must*-ified deliciousness claim would be my own, and I know that I do not like bouillabaisse in any form.

Just as with deliciousness (etc.), we can make judgments about what merits a response appropriate to beauty, without ourselves being disposed to have that response. For instance, I know someone with a weird aversion to birds. Yet this person understands that her aversion is idiosyncratic (possibly just this side of a phobia), and although she gets no pleasure from looking at birds herself, she is not so narcissistic as to think that everyone who thinks them beautiful sees merit where there is none. There are of course many other examples. I often find opera singing grating rather than beautiful, though I am an admirer of classical music in general; I find Renoir lurid and sappy, rather than beautiful, though I am in general an admirer of Impressionist painting; yet in both cases there is enough testimonial weight on the side of beauty to make me think I am for whatever reason simply not responsive to their merits.⁹³

⁹³ However, I have since chanced upon this from the *New York Times*, titled 'Rage Over Renoir':

With placards reading "God Hates Renoir" and "Aesthetic Terrorism," a small band of protesters ... descended on the Museum of Fine Arts, Boston, this week. In their view, the French Impressionist Pierre-Auguste Renoir painted "treacle" and "deformed pink fuzzy women" and is overrated, overvalued and should not be taking up gallery space here or anywhere else.

The R.S.A.P. (Renoir Sucks at Painting) movement said the event would be the first of "many anti-Renoir direct actions at museums." The goal of the self-proclaimed "cultural justice seekers" is to have museums take down their bad Renoirs and replace them with actual masterpieces; failure to do so, they say, amounts to "aesthetic terrorism."

The instigator of the movement, Max Geller, a political organizer based in New York, created a following through an Instagram account that parodies and derides the Renoir oeuvre.

Among the comments is one from Genevieve Renoir, who says that she is the painter's great-great-granddaughter and that she takes exception to the movement: "When your great-great-grandfather paints anything worth \$78.1 million dollars (which is \$143.9M in today's dollars), then you can criticize. In the mean time, it is safe to say that the free market has spoken and Renoir did NOT suck at painting." (Katharine Q. Seelye, *New York Times*, 10/8/2015.)

The McDowellian view predicts that judgments of beauty will be tied to this kind of judgment of merit, rather than to our actual dispositions to respond to potentially beautiful objects—in contrast to what is the case for judgments involving evaluative taste properties. However, this simply does not seem to be true for judgments of beauty either. Just as it is odd to say, “This must be delicious—although I hate the way it tastes” (a claim that we must scramble to interpret assuming a relativization of ‘delicious’ to some other individual or group that does not include myself), it is also odd to say, e.g., “This looks lurid and sappy to me, but it must be beautiful.”

It is not that having a less-than-positive experience of an object necessarily makes the claim “It must be beautiful” inappropriate. Gauging the dispositions relevant to judgments of beauty involves various complexities: I can think I am relevantly disposed to like something without having experienced it at all, but also, in many cases, even when I have experienced it but felt something other than liking for it. I can say that the Taj Mahal must be beautiful, though I have never seen it and it is eight thousand miles away; but I can also say that it must be beautiful, though I can only see it as a speck in the distance and it certainly does not look beautiful from here, or that a poem by Baudelaire must be beautiful, though I get no pleasure from hearing you read it to me, because I don’t understand French. We seem to think that in all of these cases, I am simply denied access to the conditions that would allow my standard of beauty to demonstrate itself—that I would feel pleasure in all of these things, were the relevant conditions satisfied: were I the right distance away, or fluent in French.

But consider an example that is only slightly different from the Baudelaire example: an English-speaking student who dislikes or is indifferent to the Shakespeare that she is forced to read for a class, because she finds Elizabethan English alien and impenetrable. This kind of dislike or indifference is different in character from the blankness I feel in the face of Baudelaire: it presents itself as a judgment—at the least that Shakespeare is boring, or otherwise disagreeable, or maybe just not worth bothering with—in a way that the blankness does not. In listening to Baudelaire, without understanding it, I may be bored; I may even judge that sitting here listening to Baudelaire is boring; but I do not judge that *Baudelaire* is boring. I know that I am not in a position to make this judgment. My blankness in the face of Baudelaire is a reflection of the fact that I make few judgments about it at all, during my experience of listening to it.

The character of the boredom of the student who dislikes Shakespeare is different. It does not strike her as dissociated in this way from the properties of the Shakespeare, but actually as the effect on her of Shakespeare’s boring self: a recognition of Shakespeare’s boringness. Yet it is surely possible for such a person to come to think that greater fluency in Elizabethan English would allow her to see the beauties in Shakespeare that other people insist are there, in spite of the negative response she currently has. Greater fluency in Elizabethan English does not seem to be analogous in this case, however, to the fluency in French that would activate my standard of beauty in the face of Baudelaire: at best, the student is now *ambivalent* about Shakespeare, and use of ‘beautiful’ is as inappropriate as a reflection of ambivalent standards as use of ‘delicious’ is as a reflection of ambivalence about the taste of a dish. ‘Beautiful’ is like ‘delicious’ in being a term of wholehearted, lavish approbation. The student doesn’t say that Shakespeare must be beautiful, in spite of thinking that greater familiarity with his work would allow her to see for

herself the merits that she cannot currently perceive.⁹⁴ Neither do I say that the Renoir must be beautiful (because people whose judgment I trust say that it is), when it looks the way it does to me.

In effect, then, judgments of merit are not only much more liberal than judgments of tastiness and deliciousness, but more liberal than judgments of beauty too. Judgments of beauty, like these other judgments, seem to depend on what the subject is actually disposed to like; not straightforwardly on what the subject independently judges to be worthy of liking. There is, however, a difficulty with ‘beautiful’ that does not occur for evaluative predicates of gustatory taste. We may explicitly relativize the latter predicates, as in, “It must be delicious, if you like that kind of thing,” or assume an implicit relativization as an attempt to make sense of a claim that clearly cannot rely on the speaker’s own standards, as in the example used above. This reliance on relativization helps us, among other things, to express judgments about the responses to food we think are merited, when we do not ourselves respond in those ways.

But any similar relativization of ‘beautiful’ is at the least somewhat awkward. “It must be beautiful, if you like that kind of thing” sounds a bit odd, as does anything like, “It’s beautiful to me.” While I have certainly seen examples like the latter in print (not to mention the cheesy song lyric, “You are so beautiful / To me”⁹⁵), they are, I believe, not nearly so common as, say, “It’s fun for me.” They occur somewhat often but not as often as the latter.⁹⁶ I take it that this

⁹⁴ I find a certain analogy tempting here. If the Baudelaire example is analogous to not being able to see a thing, because it is hidden or remote from view, we might say that the Shakespeare example is analogous to a certain kind of optical illusion. The complete opacity of French to me leaves me unable to form any firsthand judgment of Baudelaire (unless it is of merely the sound of the phrases, as if it were music), but the merely partial opacity of Shakespearean English to the student misleads her into thinking she can see it perfectly well, and her mind obliges with a judgment of it: “Shakespeare is boring.” We can of course become aware that we are subject to optical illusions, and, to run the analogy further, we may say that this is what happens when the student realizes that Shakespeare deserves better than her boredom; moreover, just as awareness of an illusion does not necessarily remove our susceptibility to it—the lines in the Müller-Lyer figure look to be different lengths even though I know they are the same—the student’s awareness of the inadequacy of her response does not all of a sudden completely change that response to one of wholehearted appreciation.

If the analogy were complete, however, she would at least now be able to say that parts of Shakespeare must be beautiful (“The lines in the Müller-Lyer figure are the same length, even though they look different”), just as I can say that Baudelaire must be beautiful. But to the extent that the student’s experience of Shakespeare still presents him to her negatively, the claim that Shakespeare must be beautiful will remain an odd one for her to make. No doubt her experience of Shakespeare is now different, as a result of her dawning realization: I take it that no one who has begun to suspect her aesthetic response to a thing of missing the mark, responds to that thing in exactly the same way she did before her suspicion arose; she will be on the lookout for what she has missed, and this inevitably changes the character of her response by giving it a flavor of exploratory uncertainty. This should at least derail her previous judgment that Shakespeare is boring. But her experience is presumably not different enough for her to make the claim about beauty—not immediately, anyway.

⁹⁵ From the song ‘You Are So Beautiful,’ by Billy Preston and Bruce Fisher, 1974.

⁹⁶ For instance, ‘beautiful to me’ gets 64 hits in the Corpus of Contemporary American English (at corpus.byu.edu), some of which are ‘seems beautiful to me,’ ‘looks beautiful to me,’ etc.—but plenty of which are ‘is/are/was/were beautiful to me.’ On the other hand, ‘fun for me’ gets 161 hits, typically variations on ‘is/was fun for me.’

These considerations are very difficult to nail down definitively, however. If it comes to that, ‘tasty to me’ gets no hits at all; ‘delicious to me’ gets 2, but these are both ‘sounds delicious to me.’ (There is also nothing for ‘tasty/delicious for me,’ which to my ear sounds even less likely anyway.) My guess is that evaluative taste predicates like these do not easily relativize in this way, because of the availability of the synonymous locutions ‘tastes good’ (for ‘tasty’), ‘tastes great’ (for ‘delicious’), etc. That is, given the ability to relativize sensory appearance verbs generally, and the relevance of ‘tasty,’ ‘delicious’ etc. to just the one sensory modality (unless

linguistic ambivalence about relativizing ‘is beautiful’ stems from precisely those features of judgments of beauty that I am trying to accommodate here. That is, on the one hand, judgments of beauty are based on the subject’s dispositions for liking, and this encourages the use of relativized beauty claims. On the other hand, however, judgments of beauty have a normative element, and this encourages us to think that beauty is beauty *simpliciter*, rather than something that could be there “for” one person but not for another.⁹⁷

used, like ‘delicious’ in the corpus, metaphorically), ‘tastes good/great to me’ is simply preferred to ‘tasty/delicious to me.’ Being fun, however, is not associated with a single sense, so ‘fun for me’ is an essential locution.

⁹⁷ Note that Kant is adamant that ‘is beautiful’ cannot be relativized at all:

With regard to the **agreeable**, everyone is content that his judgment, which he grounds on a private feeling ... be restricted merely to his own person. Hence he is perfectly happy if, when he says that sparkling wine from the Canaries is agreeable, someone else should improve his expression and remind him that he should say “It is agreeable **to me**” ...

With the beautiful it is entirely different. It would be ridiculous if ... someone who prided himself on his taste thought to justify himself thus: “This object (the building we are looking at, the clothing someone is wearing, the concert that we hear, the poem that is presented for judging) is beautiful **for me**.” For he must not call it **beautiful** if it pleases merely him. ... (*CJ* §7, 5:212.)

I take it, however, that he overstates the case a bit.

7. THE CONTENT OF BEAUTY CLAIMS: A HYBRID MODEL

I have argued that we may in fact rely on testimony for beliefs about both tastiness and beauty. When we make a claim based on testimony in this way, we use an appropriate evidential form: “That pie must be delicious—Margot says she’s never tasted better.” When tastiness (or deliciousness) is *must*-ified like this, the standard relevant to the assessment of the claim is still generally the *speaker’s* standard of taste. This is what we would expect, given the close connection between tastiness claims and the speaker’s own liking for the tastes of the foods she makes her claims about. There is something not quite right about saying, e.g., “Everybody raves about this coconut cream pie. It must be delicious, even though it tastes awful to me.” The best sense we can make of this claim requires supposing ‘delicious’ to be used here in a way that implicitly relativizes it to a different standard of taste; i.e., even though the speaker does not explicitly say this, we should understand her to have asserted something like: the pie must be delicious to most people, even though it tastes awful to her.

I have also argued that the connection between a judgment of beauty and the subject’s aesthetic dispositions is somewhat looser than that between a judgment of tastiness and the subject’s gustatory dispositions: we may have reservations about our judgments of beauty that are different in kind from the reservations we may have about our tastiness judgments. We may be unsure, that is, that we have responded to an object aesthetically in the way that it merits, and may therefore be agnostic about its beauty, even when our pleasure or displeasure would otherwise seem to indicate a definite judgment one way or the other. Such agnosticism is illegitimate for tastiness: if I love the taste of peanut butter and jelly sandwiches then I am bound to think them tasty, even if I think my standard of taste to be criticizable in some way.

The comparative looseness of this connection for beauty is what prompted my consideration of the McDowellian objectivist view. Attributions of beauty seem to be shaped by normative concerns, and these seem best represented by a kind of objectivist model. However, as I have said, this view predicts not only the agnosticism I have described—a prediction borne out in aesthetic discourse—but also predicts that people will come to hold definite beliefs about beauty that do not reflect their own propensities to respond aesthetically to things. The view predicts, that is, that a person may think something must be beautiful even though she does not like it, and this prediction does not seem right.

My recent discussion appears to return us to the kind of subjectivity for beauty that I said does not sit well with its normativity: for ‘beautiful,’ as for ‘delicious’ and other such terms related to liking or pleasure, the subject’s own propensities for liking and dislike can’t be irrelevant to the claims she makes, either for claims made on the basis of her own experience or for claims made in some other way (e.g. using *must*). In particular, the positive claim that something is beautiful (or must be beautiful, etc.) must be made with the backing of the speaker’s unreserved propensity to find the relevant kind of pleasure in the thing.

In fact, though, the comparison with evaluative taste predicates gives us the resources to reconcile these two apparently conflicting features of claims about beauty. Let me review some basic facts about the relationship between our responses of pleasure or displeasure, and our judgments, for both judgments of tastiness (etc.) and judgments of beauty. The complexity of the pleasure response appropriate to beauty—including the kind of complexity contributed by the normative concerns relevant to it—gives claims about what is beautiful and not beautiful a more complex relationship with the speaker’s pleasure than we find for ‘tasty,’ ‘delicious,’ etc. If you

think something is tasty, you must like its taste, and if you like the taste of something, you are bound to think it tasty. We have a similar biconditional for a negative judgment of tastiness: i.e. you are to think a thing is not tasty if and only if you don't like its taste. In the case of beauty, though, we only get the one conditional (together with its contrapositive): if you think something is beautiful, you must like it; if you don't like it, you don't think it is (must be) beautiful. But you can like the way something looks (say) without thinking it beautiful: in such a case you may either acknowledge or suspect your liking to be of a kind not related to beauty.⁹⁸ If you do acknowledge this, you may very well think the thing is not beautiful, although you happen to like it. And you can dislike something, or be indifferent to it, without thinking it is not beautiful, just as in the Shakespeare example: you may not like the thing, but you may also think that you are missing something—that you are not responding to it in the way that it merits.⁹⁹

I have discussed some of the norms that structure how our liking and dislike for what we eat and drink on particular occasions are related to claims about tastiness, nastiness, etc.: simple generalizations indicate a speaker's liking or dislike for a food's normal way of being eaten; simple claims, whether generalized or particular, are made from the perspective of the speaker's current, normal standard of taste. My proposal for 'beautiful' involves admitting that claims about what is and is not beautiful, like claims using evaluative taste predicates, are made on the basis of one's own liking and dislike. This is contrary to the McDowellian view. The difference between deliciousness (say) and beauty, concerns the further norms that structure how liking and dislike on particular occasions, in particular contexts, are related to claims about what is beautiful and not beautiful; or, for claims that are not made on the basis of firsthand experience, how one's knowledge of one's dispositions for liking and dislike are to related to claims about what must be beautiful.

Many of the norms for beauty are similar to the corresponding norms for evaluative taste properties. We want to be sure we are talking about the right particular (the third movement, say, rather than the whole symphony) or that a generic claim has generalized to the right degree, or that an object is being considered under the right circumstances (from six feet, rather than six hundred feet, away). But, whereas the focus for tastiness etc. is towards a speaker's current, normal standards of taste, the focus for beauty is towards standards that reflect what the object in question merits, just as for the McDowellian view. According to this different proposal, however, the pull towards such standards acts as a constraint on how the responses and propensities to respond that we actually have are related to our judgments, rather than being an independent standard that governs our judgments of beauty regardless of our own responses of liking and dislike.

So—to concentrate here on firsthand judgments—just as we can't think something is tasty unless we like its taste, we can't think something is beautiful unless we like (say) its look. But the kind of pleasurable experience through which a judgment of beauty is made, is more complex than the kind of pleasurable experience relevant to a judgment of tastiness, and our feeling in the former case must be distinguished from other feelings that may not be relevant to an appreciation of the thing's beauty. A positive verdict that the thing is beautiful affirms that one's liking is in

⁹⁸ I take this point to be a good part of what Kant is getting at in distinguishing the 'disinterested' pleasure associated with judgments of beauty from pleasure in the agreeable and pleasure in the good.

⁹⁹ Note that, if part of what we are saying when we call something beautiful is that it merits pleasure, then calling it *not* beautiful is to deny (at least in part) that it merits pleasure. Indifference or displeasure may never themselves be mandated by anything, as far as this goes. More below.

fact merited, in the way that beauty requires. If I dislike the thing or am indifferent to it, I may take this as an indication that the thing is not beautiful, but I may alternatively come to think that my response is not sensitive to the thing's merits; and in that case I will reserve my judgment.

7.1 Judgments based on pleasure: tastiness

I want now to look in more detail at how we make these kinds of judgments, and what their content must be like. This will help us to see why claims about beauty carry such robust Direct Evidential implicatures.

I will continue to focus on judgments made firsthand. I have mentioned Kant's view that, like judgments of the agreeable, judgments of beauty are *based on* pleasure; although, unlike the former, they make a claim to universal validity. What is it, though, for a claim about a particular to be based on pleasure in this way?

I will start with tastiness. I take a bite of pie and find the mouthful tasty: I feel a particular kind of pleasure in what I am eating. There are many different judgments I could make as a result of experiences of this kind, judgments that are based on this kind of pleasure. I could judge that what I am currently eating is tasting good to me right now. This is, presumably, just about the least I could judge on the basis of my pleasurable taste experience: it is hard to see, for instance, given the experience, how I could be wrong in my judgment. There is even a question whether having the experience and making the judgment are distinct occurrences. We might think that they are not: that having the experience is a way of making the judgment. But we might want to say, alternatively, that making a judgment is a more explicit cognitive achievement than feeling a pleasure. If so, there would still be a range of options. We could say that feeling the pleasure and making the judgment are never the same occurrence: that the judgment always follows the pleasure. But we could say, alternatively, that sometimes having the experience is a way of making the judgment (if it rises to consciousness in the right kind of way), but that sometimes it isn't. I don't want to answer this question in any detail. I will assume, though, that such experiences of pleasure provide the subject with a kind of intentional awareness, even if this awareness falls short of judgment. If I feel pleasure right now in what I am currently eating, I am both aware of my pleasure (taken in what I am eating right now), and aware of what I am eating (its tasting good to me right now). I am, that is to say, aware of a certain relation between what I am eating and myself. I will also assume some very close relationship between the awareness this pleasure provides and the judgment, and say that the judgment is made *by means of* the pleasure.

There are of course many other judgments I could make on the basis of my pleasurable taste experience, judgments that are more substantial. I could judge that the bite of pie is tasting particularly good to me right now given the drugs I am on. I could judge that the pie itself—the pie in general—tastes good to me when I am on these drugs. I could judge that the pie is tasty (to me in my normal state).

The first of these judgments, like the judgment I started with, does not ascribe a dispositional property to the pie (these judgments concern occurrent pleasures); the latter two judgments seem to do just that. I will take it, anyway, that that is what the latter do. Simple claims like “This pie is tasty” are generally taken to express judgments of dispositional tastiness that reflect the speaker's normal state of taste receptivity, as I have discussed. We make such claims, that is, when we judge, on the basis of our own experience, that the pie is disposed to taste good to us in our normal state. (Which dispositional property we ascribe to the pie in such a case is

presumably a controversial matter between contextualists and relativists about tastiness.) The importance of the judgments about dispositions is clear: we want to be able to think and talk about relatively stable properties of different foods, not just our momentary experiences of them.

An ordinary judgment about tastiness—about a thing’s disposition to taste good (under certain circumstances)—is based on pleasure in the sense that it is made by means of a pleasurable taste experience that the subject takes to be a manifestation of the relevant disposition. If I judge that the pie is tasty (to me in my normal state) on the basis of the taste of the bite of pie currently in my mouth, I must take my current experience of the bite of pie to be pleasurable, I must take this pleasant taste to reflect my normal capacities for sensing tastes, and I must take the pie to be relatively uniform in its capacity to affect my sense of taste. (Note the difference with color judgments. To judge that the pie filling is red on the basis of a glance in which it looks red to me, I must take the conditions of my glance to be normal for viewing and my perception of redness to reflect the normal visual capacities *of humanity*.)

Of course, I do not generally have to make explicit judgments about all of these factors, before I judge that the pie is tasty. Taking the facts about the uniformity and stability of the pie, and my normally functioning sense of taste, to be the case, may simply reflect background default assumptions that I do not explicitly think of, unless unexpected features of my situation force them on my notice. Moreover, my taste experiences themselves generally indicate to me whether they are the result of normal or abnormal sensory functioning. If I have a cold, it generally doesn’t seem to me that my food is bland; rather, it seems to me that I can’t taste it properly. (Similarly, if I am looking at something red in poor light, it doesn’t generally seem to me that the thing is brown, or gray; it seems to me that I can’t properly see its color.) We can, though, be wrong about these things. (Similarly, again, the optical exhibits at science museums are designed to exploit the loopholes in our visual systems, tricking us into false visual judgments that we are not usually unlucky enough to make.)

The disposition I judge the pie to have when judging it to be tasty—being disposed to taste good to me in my normal state—is the mirror image of the disposition I take myself to have with respect to the pie: being disposed, in my normal state, to take pleasure in tasting it.¹⁰⁰ In the absence of personal experience, it is often unclear what one’s relevant dispositions are for these matters; we often simply don’t know, without personal experience, whether or not we are disposed to find a certain dish tasty. The commonness of this ignorance is partly responsible for the erroneous sense that we can *never* know whether a dish is tasty (must be tasty) unless we have tasted it ourselves. I have, however, offered many examples in which we do know these things, in spite of having no personal experience of them. In such cases, we use *must* or some other evidential form to make the relevant claims.

When one does have personal experience of the dish in question, judging one’s taste experience to be a manifestation of the particular disposition that is relevant to the matter at hand is usually fairly straightforward. It is not very hard to know that one’s sense of taste is working normally, etc. Assuming one has correctly gauged this dispositional background, one’s

¹⁰⁰ As I said, contextualists and relativists presumably differ with respect to which dispositional property they take a claim like “This pie is tasty” to ascribe to the pie. A contextualist takes the claim to ascribe a fairly full-fledged property: the property of being disposed to taste good to the speaker under normal conditions, say. A relativist presumably takes the claim to ascribe the property of being disposed to taste good, or to taste good under normal conditions—with the proviso that the claim is made when the speaker thinks the pie is disposed to taste good to *her* under normal conditions. (I will elaborate on this later.)

experience of the dish is definitive (again, promoting the sense that you can't "really know" whether something is tasty unless you have tasted it). I may think I am going to like your pudding, but when I taste it I find out that it is not tasty after all: my experience of the pudding's taste simply determines whether or not it is tasty, whatever I thought about it beforehand. (The proof of the pudding is in the eating.) As I said, though, we are occasionally mistaken in our judgments of the relevant disposition. You complain that I have tasted your pudding after brushing my teeth, and that therefore I have no idea how it tastes; when I rinse my mouth and try again, I agree with you. My first response to the pudding turns out not to have been a manifestation of the relevant disposition.

When I say that the pudding isn't tasty, as a result of making this kind of mistake, I am mistaken in what I say, but my mistake doesn't concern the pudding's objective property of tastiness (it doesn't have one). Two people can disagree about the pudding's tastiness but agree that they have each made the appropriate judgment, given their different dispositions to respond to the pudding—their different standards of taste, as we say. (This is the kind of example in which disagreement appears to be 'faultless.')

Likewise, we can *seem* to disagree about the pudding's tastiness, but discover that at least one of us has made a mistake of this kind which, when resolved, leaves us in agreement.

7.2 Judgments based on pleasure: beauty

Judgments of beauty made firsthand are based on pleasure too. I turn a corner in the Accademia and am suddenly wowed by Michelangelo's *David*: I feel a particular kind of pleasure in what I see. I have chosen a particularly beautiful object for my example, so presumably I intend this pleasure to have something to do with the sculpture's beauty. But there are of course different kinds of pleasure we can take in the way things look. Maybe the coloration of the Carrara marble used for *David* reminds me of the countertop of the kitchen island in my childhood home, and my pleasure is a nostalgic one. I would not expect this pleasure to be shared by others; nor would I think they were missing out on something important about *David*, if—as I expect—they do not feel a similar nostalgic pleasure on seeing it.

In any case, at first glance at least, we presumably have the same array of possible judgments for visual pleasures of these kinds, as we had for the pleasures of gustatory taste. For the latter, we had judgments about occurrent pleasures—and, reciprocally, about the foods found pleasurable ("This is tasting good to me right now")—and judgments about dispositions. The judgments about dispositions, when based on a pleasure in the way I am currently discussing, depend in part on the subject's awareness of the qualities of her actual pleasure response, an awareness that is of course also required for judgments specifically about the pleasure itself.

My nostalgic pleasure in *David* may simply inspire a judgment about itself, without further implications for other pleasures: "This is pleasing me right now." Kant's account of pleasure in beauty, however, shows that a judgment of beauty based on such a pleasure always has implications for other possible pleasures—even if it is a judgment specifically about *that* pleasure and not about the object's disposition to cause pleasure of a certain kind. This is because pleasure in beauty is a normative pleasure: the subject takes it to be appropriate in herself and claims it should be felt by others. Moreover, a subject's experience of pleasure in beauty must itself provide her with awareness of this special quality that it has. This is, I think, less odd than it perhaps sounds: just as my awareness of a pleasant taste is generally also an

awareness of a quality of the food or drink in my mouth,¹⁰¹ my awareness of a pleasure I take to be normative is also an awareness of a quality of the object I am responding to: an awareness of the object's *merit*—of its deserving the response I am currently having to it, at least right now, in these specific circumstances. Kant points out that it must be the pleasurable experience itself that provides this awareness, given the nonconceptual nature of the judgment of beauty. I cannot judge, independently of the kind of awareness that the pleasure provides for me, that my pleasure in an object is a normative pleasure, and hence pleasure in beauty, because there can be no conceptual 'proof' of an object's beauty, and such an independent judgment would presumably have to be a conceptual proof of this kind. This means that my pleasurable experience must be produced by the object in a way that is itself sensitive to the normative relevance of the object to that pleasure.¹⁰²

Ginsborg's interpretation of Kant, discussed earlier (§2.2), takes an ordinary judgment of an object's beauty to go no further than this: it is a judgment about an occurrent pleasure—though a normative pleasure—rather than a disposition of any kind. As an interpretation of Kant, her account is disputed,¹⁰³ but it is clear both that an ordinary judgment of beauty, based on pleasure, depends on the subject's awareness of the qualities of her pleasurable experience, and that Kant focuses a lot of energy on the nature of that experience. He paraphrases the judgment of beauty ("That's beautiful") by saying it is the judgment, made by means of one's pleasure in an object, that one's pleasure is 'universally communicable' and therefore ought to be felt by everyone—which, as I said above, may seem odd, except that, taking the pleasure to be a kind of intentional awareness of the object, this is the same thing as judging that the object is, through my pleasure, *revealing its merits to me*. Whether or not Kant intends this to comprise the entirety of an ordinary judgment of beauty—what is usually communicated by saying "X is beautiful"—as Ginsborg maintains he does, it is at least a necessary part of any ordinary judgment of beauty based on such a pleasure. My paraphrase here ("X is revealing its merits to me") also shows why such a judgment has normative implications for other possible pleasures, and thus in a way goes beyond the actual pleasure response on which the subject's judgment is based: a judgment about the object's merit is contained within it. This makes it different from the other examples I

¹⁰¹ There is presumably a kind of counterexample (hence my inclusion of 'generally' above). Say I am blindfolded and told to open my mouth while someone sprays an intensely flavored mist onto my tongue, so gently that I can't really feel it. I may in this case be aware of a pleasant taste without being aware of anything that *has* the taste.

¹⁰² This is, perhaps, an overly breezy way to sidestep the welter of controversy concerning Kant's §9 of the *Critique of Judgment*, where I take him to present the point I have just made. I guess I will just say that I take him to be saying at least this much there, whatever else he means to say as well. While I hope that my claim about this is not very controversial, I cannot hope to avoid all controversy whatsoever. For instance, Paul Guyer's interpretation of §9 (1997: Chs. 3-4) involves taking the pleasure in beauty to be 'opaque'—i.e. not a provider of intentional awareness of any kind.

Incidentally, Kant makes a comment elsewhere on this topic that I find particularly illuminating: "Agreeableness is also valid for nonrational animals; beauty is valid only for human beings, i.e., animal but also rational beings, but not merely as the latter . . . , rather as beings who are at the same time animal; the good, however, is valid for every rational being in general" (*CJ* §5, 5:210). Animals as well as people have sensory pleasures; rational beings in general (including, I suppose, gods) can make judgments based on reasons about what is good; but only *people* can make judgments of beauty. I take it that this is because only people (not necessarily human beings, but, as Kant says, *rational animals*) can have this kind of sensory sensitivity to normative features. As I have mentioned, we often give reasons for finding something beautiful, even though the reasons we give cannot entail the thing's beauty. I take this practice to reflect the *partial* transparency to us of the grounds of our judgments; much about our sensory consciousness, however, is simply inaccessible to us.

¹⁰³ E.g. by Henry Allison (2001: 113-115) and (2003).

have given of judgments based on pleasures, *about* those pleasures, that do not make claims beyond the pleasures themselves: “This is tasting good to me right now,” or “This is pleasing me right now.”

But if this is all that is claimed by an ordinary judgment of beauty—let’s say, roughly, “This object is revealing its merits to me right now through this experience”—then we cannot account for claims about beauty that are made in the absence of any experience of the object: not just claims about an object’s beauty made after the fact (maybe we could rely on *remembered* experience for that¹⁰⁴, or make some allowance for past tense, as I vaguely suggested in the earlier section), but claims made—possibly in embedded contexts, or with appropriate evidential markers—when the subject has never experienced the object at all. As we saw earlier, even negative judgments of beauty are problematic on this experiential, normative, subjectivist view. We need beauty to be a property of objects that is independent of any given experience of beauty, just as tastiness and deliciousness are properties of objects that are independent of any given experience of taste. Of course, just as for tastiness and deliciousness, it needn’t be an objective property.

Yet we have also seen that we can’t treat the judgment of beauty as straightforwardly attributing aesthetic merit to the object either, where this is an objective, or objective-ish, dispositional property that the subject becomes aware of through her pleasure. This would be the McDowellian view. (The property would be objective or objective-ish, depending on your take on whether a critical standard of the kind McDowell elucidates can count as objective.) If this view were right, we would be able to attribute such a property to objects, by calling them beautiful, when we are convinced of their merit, in spite of our own dislike.

We need, in short, something that sits between the experiential normative subjectivism possibly suggested by Kant, and the objectivism of the McDowellian view. The firsthand judgment “The pie is tasty” is based on a pleasurable taste experience that the subject takes to be a manifestation of her normally functioning sense of taste; and she thereby takes the pie to be disposed to taste good to her in her normal state. Analogously, then, we may take the firsthand judgment “*David* is beautiful” to be based on a pleasurable visual experience that the subject takes to be a manifestation of *her disposition to respond appropriately to David*; thereby taking *David* to be disposed to produce pleasure in her in her *normatively characterized* state: her state of responding as she ought. In making the judgment of beauty, she must take herself to occupy a normative—or, as Kant says, a ‘universal’—perspective.

The subjectivism of Kant’s view, motivated at least in part by his emphasis on the autonomy of taste—the (apparent) need for a judgment of beauty to be based on the subject’s actual experience of pleasure in the object—is thus not wholly swept away by this suggestion. My suggestion takes beauty to be more or less the kind of hybrid property, a blend of subjective and objective features, that Kant takes such pains to insist it is. It has a normative element—someone who judges an object to be beautiful claims in part that the object merits a certain response—and this normative claim is subject to an objective (or objective-ish) critical standard. Yet no one is entitled to claim an object is beautiful unless she is disposed to respond

¹⁰⁴ This is exactly how Ginsborg claims we should explain beliefs about beauty that persist after the beautiful object is no longer being experienced, on her interpretation of Kant. (1998: 461, fn. 8) Her idea (on Kant’s behalf) is that if I assert now that *David* is beautiful, although I last saw it many years ago, my assertion is based on my current pleasurable memory of *David*, rather than being based on my earlier experience.

appropriately to its merits by feeling pleasure in it, and this makes beauty similar to gustatory taste properties, which are not objective.

Someone who makes a judgment of beauty, then, must judge the object to be disposed to produce pleasure in her under the assumption that she occupies a normative, or universal, perspective. This judgment makes a claim about the object whose content is independent of the experience by means of which the judgment is made, but the legitimacy of the claim is *not* independent of the speaker's disposition to appreciate the object's beauty. My proposal is thus a form of (what we might still call) normative subjectivism for beauty, though it is not 'experiential' in the way that I take Kant's view to be, and, for reasons that will be made clearer in the next section, this would perhaps turn out to be a misleading label, so I will steer clear of it.

If, in making a firsthand judgment of beauty, the subject must take her perspective to represent a normative, universal perspective—a perspective that everyone should occupy—we can see the normative element of beauty to be a constraint on the kind of judgment the subject makes by means of her response to the object. This is in fact exactly what Kant means to suggest by speaking of a 'universal voice':

If one judges objects merely in accordance with concepts, then all representation of beauty is lost. Thus there can also be no rule in accordance with which someone could be compelled to acknowledge something as beautiful. Whether a garment, a house, a flower is beautiful: no one allows himself to be talked into his judgment about that by means of any grounds or fundamental principles. One wants to submit the object to his own eyes, just as if his satisfaction depended on sensation; and yet, if one then calls the object beautiful, one believes oneself to have a universal voice, and lays claim to the consent of everyone, whereas any private sensation would be decisive only for him alone and his satisfaction. (*CJ* §8, 5:216)

I have said that someone who does not feel pleasure in something that she thinks must merit such pleasure, is entitled to claim neither that the object is beautiful, nor that it is not beautiful. Her standards in such a case, I have claimed, are ambivalent. Structuring beauty's subjective and 'universal' elements in the way that I have now gives us the resources to explain the circumstances under which we are willing to make claims about what is and is not beautiful. If I feel pleasure in an object but do not take the perspective from which I like the object to represent the universal, then of course I do not call the object beautiful, but use some word from the 'agreeable' family instead. In such a case I may say the object is not beautiful, though I happen to like it; and my saying it is not beautiful is based on the judgment that it is not disposed to produce a pleasure in me in my capacity as a universal judge. I may of course make the same claim—that the object is not beautiful—when I feel no pleasure in it at all. But I may *not* claim that the object is not beautiful *just* because I feel no pleasure in it: such a claim must express the judgment that the object is not disposed to produce pleasure in me in my capacity as a universal judge, and if I have reason to think that in this matter I cannot occupy a universal perspective, then I am not in a position to make this judgment. Neither, of course, can I judge that the object *is* beautiful: I must in such a case judge that the object *is* disposed to produce pleasure in me in my capacity as a universal judge, yet I neither feel the requisite pleasure nor take myself to have the requisite universal perspective. If I take no pleasure in an object that I think must nevertheless have merit from a normative perspective that I cannot occupy, I simply reserve my judgment: I can judge neither that it is beautiful nor that it is not beautiful.

I have called the position of such a person—the example I have described most extensively is the student who dislikes Shakespeare—‘ambivalent,’ but this is not the kind of ambivalence that first comes to mind when someone is said to be ambivalent about a taste or an object’s beauty. We ordinarily say we are ambivalent about a taste when we are unsure, concerning a given taste experience, whether it is pleasant or not: it may have elements we like and elements we do not like. We ordinarily say we are ambivalent about something’s beauty when our experience of the object is similarly mixed: we may, for instance, be unsure whether ‘beautiful’ is really the right word to use to describe something we find harrowing. This kind of ambivalence makes us unwilling to issue a judgment involving properties that generally involve a much clearer kind of response. Our dispositions in these cases are not clearly dispositions to respond with liking or dislike; they are dispositions to respond with a blend of the two.

My student who dislikes Shakespeare is similarly unwilling to issue a judgment about Shakespeare’s beauty, and this unwillingness also results from a kind of unsuitability in her relevant dispositions. The kind of unsuitability at issue here is different, however. She is simply debarred from playing the beauty game in this instance, by her inability to take herself to occupy a universal perspective—to have dispositions to respond to Shakespeare that she can take to be normative. We would get a similar failure to ante up in someone who simply had no stable taste responses, and who thus could not be said to have the normal dispositions to respond to foods that are indicated by saying her sense of taste is functioning normally. Such a person would not be able to say that the pie is tasty, or that it’s not tasty, but would have to say, instead, that some bites of it taste good to her and some don’t. The Shakespeare student has her different line of retreat: “I don’t like it, but I think I must be missing something.”

7.3 The structure of beauty claims

In the last section, I made a comparison between the role that *normality* plays in ordinary tastiness claims, and the role that a kind of normative *ideality* plays in beauty claims. In this section I want to explore this comparison, and the nature of the different kinds of dispositions attributed by such claims, with an eye to being more explicit about the structure of beauty claims. I will follow what I take to be an intuitive route, will not obsessively defend the choices I make along the way, and will come to no firm conclusions. I aim simply to sketch some of the issues that are raised by these ideas and to put them in sharper relief.

The natural thought to have when normality is found to shape a discourse, in the way that it seems to do for taste properties, is that genericity is somehow involved: we often say that Ks are F (e.g. “Tigers are striped”) when Ks are generally, or typically, or normally F, and some semantic accounts of generics use a conception of normalcy in their analyses.¹⁰⁵ I have already discussed some examples of generics involving taste-related properties (e.g. “Coffee is disgusting”, cf. §6.1), and how claims that use them typically involve normalcy in two ways: they convey that they are made on the basis of the food or drink’s normal way of being eaten or drunk, and on the speaker’s normally functioning sense of taste. The first of these ways is straightforwardly typical of generics: if we take it (as I have been) that taste properties are dispositions to taste a certain way, then “Coffee is disgusting” will typically be used to say that,

¹⁰⁵ See e.g. Nickel (2016), who offers a normalcy-based account of generics. Note, however, that many apparently true generics are hard to square with a normalcy-based approach: “Mosquitoes carry West Nile virus” (a common example in this context) seems true, in spite of the fact that only a small percentage of mosquitoes carry the virus and it is presumably not normal for them to do so.

normally, instances of coffee have such a disposition, or to ascribe the disposition to normal instances of coffee.

The second way that normalcy is involved concerns these dispositions themselves. The dispositional taste properties that are ascribed to particulars by typical singular taste claims (“This cup of coffee is disgusting”)—and that are also involved in generic claims, as I just described—must be understood as suitably related to the speaker’s normal standard of taste, whether one takes this relation in a contextualist or a relativist way (cf. fn. 100). The connection with genericity here is not so straightforward, but it has been noted that ascriptions of dispositions are closely related in some way to ‘habitual’ statements, which are considered to be a special kind of generic: they concern particulars, not kinds, but can be used to describe a particular’s typical or normal behavior (e.g. “Margot takes the bus home from work”). If generic sentences concerning kinds express some type of generalized quantification over members of a kind, habitual sentences express generalized quantification over situations or times relevant to the particular at issue. (Generally, situations in which Margot goes home from work are situations in which she takes the bus; or: generally, situations in which Margot leaves work for the day are situations in which she takes the bus home.)

There is an intuitive connection between disposition ascriptions and habituals: “Margot takes the bus home from work” might not at a quick glance seem to convey anything significantly different from “Margot is disposed to take the bus home from work.” This connection shows itself plainly in taste discourse through the apparent synonymy of locutions such as “X is tasty”/“X tastes good”, “X is disgusting”/“X tastes terrible” etc.: the first of each pair appears to ascribe a dispositional property to X, while the second has the form of a habitual. Moreover, both disposition ascriptions and habituals tolerate many of the same exceptions: neither claim about Margot is falsified by Margot’s occasionally walking home instead, for instance.¹⁰⁶ The two kinds of claim are perhaps used somewhat differently, however. Wasserman (2011) suggests that habituals are typically more closely tied to actual events than disposition ascriptions. An object may of course have a disposition without ever manifesting it—the cup is fragile, yet never actually breaks—but habituals are typically used to describe *actual* behavior that conforms to a regular pattern. I may be disposed to get hives after eating shellfish—something I learn through allergy testing—but if (as a result of the testing) I steer clear of shellfish my whole life, never eating even one, it would be odd to say that I get hives after eating shellfish.¹⁰⁷

¹⁰⁶ Fara (2005) offers an account of dispositions that is based on this connection with habituals, which he uses in his analysis instead of the more traditional (but problematically exception-intolerant) counterfactual conditional.

¹⁰⁷ The issue is far from clear-cut, however. Wasserman acknowledges that habituals that are used to describe an object’s *function* are perfectly appropriate even when the object has never exhibited that function in actual behavior: this bottle-opener opens bottles, even though it’s still in the packaging and has never actually opened a bottle. Even when there is no function at issue, habituals that allude to a common and well understood (though not function-related) disposition may also seem natural in the absence of a pattern of regular behavior: “This rock dissolves in water” sounds fine, even though the rock clearly hasn’t ever dissolved in water before, and if it were to, would not be able to do it again. It is only when the object’s behavior is not obviously related to a disposition of this kind that a habitual seems to require a regular pattern of actual behavior.

If it comes to that, disposition ascriptions themselves tend to suggest such behavior has occurred, if the disposition in question is unlike solubility, fragility etc. in being commonly associated with certain types of objects. My example above, for instance—“I am disposed to get hives after eating shellfish”—would be likely to suggest that I know this because it has happened. In light of these observations, it is plausible to suppose that both habituals and disposition ascriptions are sensitive to evidential phenomena. When an object’s behavior is predictable from its

Whatever the exact relationship between disposition ascriptions and habituals, it seems clear that a conception of normality provides the background against which many claims of both kinds are made.¹⁰⁸ The quirk of taste predicates is of course that they typically refer to dispositions with a connection to what is normal for a specific person or group rather than what is normal for, say, humanity at large, as with color predicates. However, while typical simple taste claims do relate to the speaker's normal standard of taste, taste predicates are in fact highly flexible and may be used to ascribe different kinds of taste-related dispositions to objects, depending on context. If I say "This pie is delicious" at a run-of-the-mill picnic, I ascribe a disposition to the pie that must be understood in relation to my normal standard of taste. But if I am part of a group participating in an experiment in which our sense of taste is altered in various ways—and we know this—then saying "This pie is delicious" at lunch with fellow subjects of the experiment will ascribe a disposition to the pie that is understood in relation to my altered, abnormal standards.

It is important to note in any case that 'normality' is not a single way of being. One way to be 'normal' is to be statistically prevalent. This conception of normality surely underlies "Margot takes the bus home from work": the truth of this habitual presumably requires a consistent pattern of Margot actually taking the bus. But 'normal' in many contexts carries the connotations of *natural*, whether or not what is natural is statistically prevalent: 'normal' pregnancies develop according to biological design, even if environmental toxins and their harmful effects on gestation are what is statistically prevalent around here.¹⁰⁹ Biological things shaped by evolution have functions: ways they are designed to behave, or are 'supposed' to behave; and normal functioning for, say, a heart, means doing what it's supposed to do (i.e. pumping blood). Manmade things of course have functions too.

function or because it is an instance of a well understood type, it need not actually demonstrate the relevant pattern of behavior for either a habitual claim or a disposition ascription to be appropriate. In the absence of such predictability, the evidential assumption will be that the object has demonstrated such behavior. The way I have described the hives example in the main text—I mention allergy testing as the source of my knowledge—works to cancel the typical evidential assumption. In general, speaking of '*predisposition*' rather than 'disposition' has the same effect.

However, habituals and disposition ascriptions do not seem to be exactly the same with regard to these evidential phenomena. "I get hives after eating shellfish" is plausibly false if I have never actually eaten shellfish and come down in hives afterwards; in the same circumstance, an utterance of "I am disposed to get hives after eating shellfish," though possibly misleading, may surely be true. This aspect of habituals is even more pronounced when the relevant behavior is not so obviously tied to relatively unalterable features of the subject. For instance, "Margot takes the bus home from work" seems to me very clearly false if Margot has no regular pattern of taking the bus home from work.

¹⁰⁸ E.g. Mumford (1998: §4.9) emphasizes the role of 'normal' background conditions in making disposition ascriptions. He says that we must understand the conditionals traditionally used in analyses of dispositions (e.g., for solubility, 'if *x* were placed in water, *x* would dissolve') as true only given certain 'ideal' conditions, which in many cases will be normal conditions.

¹⁰⁹ Haslanger (2014) explores how the use of generics can obscure various problematic assumptions related to these concepts. "Ks are F" may suggest that Ks are statistically normally F; it may be assumed that Ks are normally F because it is their nature to be F; it may further be assumed that things ought to express their nature, and that we therefore ought to encourage Ks to be F. Use of generics can reflect all three possibilities: "cars have radios" depends on a merely statistical regularity; "birds fly" reflects the natures of ('normal') birds; "boys don't cry" may be used as a normative admonition to a crying boy, to get him to "be strong" and stop crying. As Haslanger says, however, "what's [statistically] 'normal' is not always natural, and what's natural is not always best" (365); yet, she argues, we may easily fail to notice the slide from *normal* to *natural* to *best*.

Habituals may allude to an object's statistically prevalent behavior, but they may also allude to an object's function: "This plant's sticky seed pods attach themselves to animals' fur"; "This gadget pits cherries"; "Margot fields the calls from New Guinea". They may do this even if the object has rarely or never exhibited its function in its actual behavior—the gadget is still in the packaging; this plant is growing indoors in a house with no pets; there haven't yet been any calls from New Guinea (cf. fn. 107). The behavior that the habitual describes is what the object is designed to do (for Margot and the calls from New Guinea, it is the role she has been given to play), whether or not the object actually does much of it.

Yet a habitual that alludes to an object's function will still come out false if the object consistently fails to perform its function in circumstances that call for that performance. Objects may be designed badly; or, even if they were originally designed well, the environments in which they were designed to function may have changed in ways that prevent the objects from functioning well anymore. These seed pods are supposed to stick to animals' fur (but don't, because climate change has produced animals around here that don't have any fur); this gadget is supposed to pit cherries (but doesn't, because it's a piece of crap); Margot is supposed to field the calls from New Guinea (but doesn't, maybe because she's a terrible employee, or because the calls are all mistakenly routed to Mike).

The truth of this kind of function-related habitual appears to depend, then, on what it does or would do in a certain class of special circumstances: those in which, in some sense, the object 'ought' to perform its function. This class is hard to delineate. The cherry-pitter isn't failing to do what it's supposed to do just because it doesn't pit cherries when it's lying around in a drawer: it's not designed to pit cherries all by itself in a drawer. But the seed pods *are* failing to do what they're supposed to do in their new, furless environment, even though they were not designed for that environment either. The class, as I said, is hard to delineate—hence my no doubt tendentious use of 'ought' just now.

We can, however, assimilate function-related habituals to habituals like "Margot takes the bus home from work", which clearly depend on a conception of statistical normality, in the following way. For the latter (on one reading), the class of relevant circumstances is restricted to those in which Margot goes home from work, and if some sufficient proportion of these are circumstances in which she takes the bus—where 'sufficient' in this and similar cases is presumably context-sensitive—then the habitual is deemed to be true. If we agree to shelve the difficulty concerning 'ought', then for the cherry-pitter (say), we can restrict the class of relevant circumstances to those in which it ought to pit cherries, and if it succeeds in pitting cherries in some sufficient proportion of these, then its habitual is deemed to be true. The relevant restrictors simply have different shapes.

It is not at all clear to me what proportion of these relevant circumstances are counterfactual or possible rather than actual. It seems to me that a habitual like "Margot takes the bus home from work" is simply false, rather than merely infelicitous, if Margot hasn't actually gone home from work a fair bit, and taken the bus to do it some sufficient number of these actual times. On the other hand, as I've noted, the truth of "This gadget pits cherries" does not require the gadget to have pitted any cherries at all, and is clearly dependent on various possible cherry-pitting situations. However, it is these habituals' associated dispositions—and their bearing on the dispositions that are associated with taste and beauty—rather than the habituals themselves, that are my primary concern, and the modal issues for dispositions are, if not exactly more straightforward, then at least more clearly relevant across the board. Someone asks Margot, who

is contemplating the new job she is about to start on Monday, how she will typically get home from work, and she answers, “I’m disposed to take the bus.” This is true if in various possible going-home-from-work situations she does in fact take the bus. Which possible going-home-from-work situations? Well, I’m not exactly sure, but it is tempting to say that the relevant situations are ones that are very much like the statistically normal sorts of situations that Margot will encounter on workdays, and that if she takes the bus in a sufficient proportion of these, then her disposition ascription is true. We might impose an order on the possible going-home situations, according to how normal they are, given Margot’s actual situation. Then we could draw a line around some number of the most normal, and demand that it is a sufficient proportion of these that must involve Margot taking the bus.¹¹⁰

According to this picture, Margot’s disposition ascription may be true (and remain true) even if, on Monday, it is such a nice day that, on a whim, she decides to walk home; and it may be true as well if the world takes an unusual turn on Monday and all the bus drivers go on strike. Any disposition ascription she *now* makes will be sensitive to her new abnormal situation—if she flatly says, “I’m disposed to take the bus (anyway),” we may be puzzled, and conclude she intends to steal one—but she may also allude to her continued possession of the prior disposition by saying, “I would be disposed to take the bus, if it weren’t for this strike.”

I believe this to be an intuitive sketch of how we think of dispositions, and I am going to run with it, in spite of its manifold uncertainties.¹¹¹ The example involving Margot and the bus may

¹¹⁰ Note that ‘situations’ are centered worlds: triples of world, object (in this case Margot) and time.

¹¹¹ I will mention one difficulty. Standard possible worlds semantics along the lines of Kratzer (1991) models graded possibility (p is likelier than q , p is probable, etc.) in terms of the relevant worlds’ ordering. If p is probable (say), one might intuitively expect this to mean that there are more relevant worlds—i.e. more worlds of the ‘modal base’—at which p is true than at which it is false. But if the modal base is infinite—as it generally is, in all but toy models, and uncountably infinite at that—then no obvious meaning can be attached to the idea that a majority of the base have a certain property, when that supposed majority and the remaining ‘minority’ (those worlds without the property in question) must have the same cardinality: when they are both uncountably infinite, just like the modal base that they divide. Kratzer thus models graded possibility in a different way, involving an ordering source: the worlds of the modal base are ordered according to some metric, for instance how normal or stereotypical they are. Then ‘ p is probable’ is true when there is a p -world that is closer to the ‘ideal’ determined by the ordering source than any *not- p* -world.

Now for dispositions. Intuitively, as I indicate in the main text, we want to be able to say that X is disposed to m when a ‘sufficient proportion’ of relevant X -situations are ones in which X *ms*. Dispositions are gradable: X can be more disposed to m than Y (more fragile, say), and the natural comparison with gradable adjectives suggests that, just as ‘ X is taller than Y ’ is true when X ’s degree of height is greater than Y ’s degree of height, ‘ X is more fragile than Y ’ is true when X breaks in more situations than Y . Similarly, the ‘sufficient proportion’ necessary for a positive disposition ascription (‘ X is fragile’) seems akin to the threshold on a scale required for a positive claim involving a gradable adjective (‘ X is tall’). But, as I just illustrated for gradable possibility, no obvious meaning can be attached to these relative proportions: every proportion appears to have the same cardinality, so it seems to make no sense to say that one is bigger than another or that *this* uncountably infinite proportion is ‘sufficient’ but *that* uncountably infinite proportion is not.

The obvious move in response to this problem is to try to apply Kratzer’s method: to impose an ordering on the relevant situations, and ultimately to link an object’s possession of a disposition with the object’s behavior in situations that are close to the ‘ideal’. This move is additionally attractive because some situations do seem to be more relevant to whether an object has a certain disposition than others. This is why I suggest an ordering on Margot’s going-home situations in the main text: intuitively, whether she takes the bus in humdrum, normal situations seems more relevant to her disposition than whether she takes the bus when a roc swoops down and carries her away. (Dropping a glass onto a stone floor seems more relevant to its fragility than dropping a bubble-wrapped glass onto a stone floor.) But the move doesn’t work. The problem with it is illustrated by Manley and Wasserman’s (2008) ‘Achilles’ heel’ scenario. Suppose there is a concrete block that is very sturdy in almost all

not seem like a paradigmatic disposition (like fragility, solubility, etc.) but I chose it for its clear reliance on a kind of statistical normality (although this reliance is probably clearer for the associated habitual, which depends on Margot's demonstrated behavior regarding the bus). The question now is what to say about a typical utterance of the form "X is tasty." What kind of disposition does such an utterance ascribe to X? I have said that the disposition in question is related to the subject's normal standard of taste; but is the relevant conception of normality here that of statistical prevalence, or does it have more in common with natural function?

If we opt for statistical prevalence of the kind I sketched for Margot, then we would say that a subject's standard of taste is disposed to function normally when it is in the state it is in *usually*: when its response in various possible food-tasting situations is, in a sufficient number of cases, what (given the facts about my life) it would usually be. Maybe this is the right way to go. It does seem, after all, that the kinds of things that we think of as distortions to our sense of taste—toothpaste, illness—are relatively uncommon. But I'm not sure. I don't know how to untangle the issues concerning someone, for instance, whose sense of taste permanently deviates—through illness or mishap—from its naturally designed function. The human sense of taste definitely *has* a naturally designed function—though clearly the range across individuals of what counts as functioning according to design is much greater than it is for, say, color vision.¹¹² My question here concerns whether our taste discourse is tethered to our taste faculty's 'correct' functioning or to its usual functioning. Say, as a result of illness, my sense of taste alters in a way that makes sweet things unpleasant to me. It is dessert time at a dinner party and no one knows of my idiosyncrasy. Pie is served, and, because I used to love pie, I try a small bite, hoping against hope. "The pie is awful," I sigh, disconsolately—thereby startling the host and other guests. Where on the scale from true but misleading, to broken in some more serious way, to outright false—given the semantic and pragmatic issues taken together—does my claim fall? ("The pie tastes awful to me" is, I would say, less but still somewhat misleading in this context.)

I'm not sure, but such cases might incline us to an alternative to a purely statistical approach to the requisite normality of the subject's sense of taste, involving an appeal to its natural function. We could say that a subject's sense of taste is disposed to function *right*—without

circumstances, but has an 'Achilles' heel': if dropped at exactly the right angle on a particular corner with exactly the right amount of force, it will break. This block is surely not to be counted as fragile, though it clearly has one specific weakness. The problem this scenario poses for modeling dispositions in terms of closeness to an ideal world is, as Manley and Wasserman note, that the Achilles' heel scenario provides a "recipe for generating counterexamples" (68): for *any* world that is supposed to be closest to the ideal, we can posit an object that breaks in that world but not in many others, and hence should not be counted as fragile. (And so on for any disposition.)

Our original intuition concerning dispositions, that they concern an object's behavior in a *range* of situations—though possibly ordered for relevance as well, as in Kratzer's system—thus seems hard to give up. Yet it is not at all clear how to make the intuition respectable. Manley and Wasserman speculate (79-82) that there may be a non-arbitrary measure on the relevant sets of situations that would allow us to compare their relative proportions. They liken the problem here with that of legitimizing the intuition that, on a real line between 1 and 100 meters, there are in some sense fewer points between 1 and 2 than there are between 2 and 100. The cardinal number of points in each interval is the same, but if we compare the lengths of each interval rather than the number of points in each interval, we get a respectable (if somewhat modified) version of the original intuition. Their hope is that we could find a similar substitute to cope with dispositions. With the setup I have outlined in the main text concerning Margot's disposition to take the bus home from work, one way to cope would be to find some non-arbitrary way to draw the line around the 'most normal' situations that would render this set—the set of relevant situations—finite. However, while I think that this procedure would accord with our intuitions about dispositions, I have no idea how to do it. (See also Vetter 2014.)

¹¹² I set aside the fact that our sense of taste can be further designed by our own efforts to cultivate it.

damage or interference—when its response in various possible food-tasting situations is, in a sufficient number of cases, within the range of responses that it *ought* to deliver in those cases. It is of course a complex matter to determine how one’s sense of taste ought to respond in any given case (much more complex than determining this for the cherry-pitter, or even for the sticky seed pods), and it is important to emphasize that two people can both have senses of taste that are working right—right for them, as it were—even though they deliver opposing verdicts on the tastes of some foods.¹¹³ But I take it that the basic idea here—that there is such a thing as one’s sense of taste’s function, and that it can therefore operate in line with its function or not—is uncontroversial.

A similar issue concerning the relevant kind of ‘normality’ for taste dispositions arises for the foods that we make taste claims about: are the dispositions they have to taste a certain way to be understood in terms of how they are usually eaten, or how they are supposed to be eaten? Again, I am not sure. Maybe the reason it is hard to make this call is that, as human intention is so heavily involved in determining how foods are supposed to be eaten, and human agency so generally successful at making it the case that foods are eaten as they are supposed to be, the way foods are supposed to be eaten is usually how they are eaten: the two properties coincide. Similarly, so much leeway is given to an individual’s sense of taste for counting as functioning correctly, that what is usual for anyone’s sense of taste is generally within the bounds of its correct functioning. The outliers—those whose sense of taste is irretrievably damaged—do not, I think, generate clear intuitions.

I said earlier, about firsthand tastiness judgments (“The pie is tasty”) made on the basis of a given taste experience, that the subject takes the experience to be a manifestation of her normally functioning sense of taste, thereby taking the object (the pie) to be disposed to taste good to her in normal presentations of itself to her normal sense of taste. I have now spelled out in a little more detail what such a disposition could amount to, whether we take the relevant conceptions of normality, both for the subject’s sense of taste and for the pie itself, to be statistical or functional. If my sense of taste simply does not work right (for the functional version), or clearly has no state that it is usually in (for the statistical version: it changes wildly from one moment to the next, for instance), then my utterance of “The pie is tasty” in a typical context—a context in which one or other version of normality is expected to constrain the disposition that I am attributing to the pie—will fail.¹¹⁴ I will have to use different locutions, and various explanations, in order to buck the typical contextual expectations of the simple claim.

Even if we do not ultimately opt for the functional version of taste dispositions, thinking about taste dispositions in this way puts us on a promising-looking track for thinking about beauty. I said that, analogously with tastiness judgments, firsthand judgments of beauty (“*David* is beautiful”) are made on the basis of an experience that the subject takes to be a manifestation of her disposition to respond appropriately to the object (*David*), thereby taking the object to be disposed to produce pleasure in her in her normatively characterized state: her state of

¹¹³ It may also be the case that a certain taste response that would be in the range of well-functioning for one person would not be in the range of well-functioning for another person.

¹¹⁴ Will it be false or just misleading? This is an interesting question in its own right. The reason I hesitate to opt for ‘false’ is that there are so many dispositions in the neighborhood, so to speak, that it is very easy to recalibrate on discovering that a speaker meant something different from what you as hearer expected, even if you had a right to expect it; as well as, in fact, on discovering that the speaker did not mean anything very determinate at all. If the speaker was not very mindful of which disposition she meant, she and her hearer may hash out between them which particular disposition it makes most sense to focus on.

responding as she ought. Taking the idea of natural function as a model for both aesthetic object and aesthetic viewer, we can say that the relevant disposition (the beauty attributed to *David*) is constrained by how the object is ‘supposed’ to be seen and how the viewer is ‘supposed’ to see it.

If I have never seen the Taj Mahal, I can say that it must be beautiful, because the disposition I thereby attribute to it depends on my response to it in situations that have not actually obtained: it depends on the response I would have on viewing it from an appropriate distance, in appropriate light, etc., given my actual aesthetic sensibilities, which I take to be how they ought to be with respect to the Taj Mahal’s beauty. If I don’t understand French, I can say that Baudelaire must be beautiful, for much the same reason: I take my claim to depend on the response I would have to Baudelaire if I were able to read it or hear it from the appropriate vantage of fluency in French. But the Shakespeare-hating student can’t say that Shakespeare must be beautiful, because she has come to suspect that her aesthetic dispositions are not what they need to be for a just aesthetic appraisal of Shakespeare.

However, as suggestive as the comparison between beauty and natural function is, two questions immediately arise. The easier one, which I will take first, concerns what I have said about the leeway given to the sense of taste: that two people who have opposing verdicts about the taste of a food may yet each have a sense of taste that is operating consistently with its correct function. If natural function does in fact provide a good model for aesthetic response, the latter nevertheless appears to differ from the sense of taste in this respect: two people who disagree about an object’s beauty cannot both be seeing the object as it is ‘supposed’ to be seen, with faculties that are ‘supposed’ to see it that way. This accords with a common intuition that disagreements about beauty are more ‘robust’ than disagreements about matters of gustatory taste, but the question arises: why can’t aesthetic sensibilities be allowed to have as much leeway as the sense of taste?

The answer involves each faculty’s relationship with the objects of its appraisal. Our sense of taste has a function, but whether it is exercising that function in its appraisal of a given food is not fully determined by its response to that food. That food does not demand a particular response that all by itself determines whether the subject’s sense of taste is functioning right, insofar as it delivers that response or not. But, assuming for the moment that there is a sense—even a loose or metaphorical one—in which the aesthetic faculty has, or mimics having, a function, whether or not it is functioning as it is supposed to do has *everything* to do with the particular object of its appraisal. It is the object all by itself that demands a certain response, and it makes the same demand of everyone.

I will have a little more to say about that in the next few sections, but the second, harder question I mentioned is clearly now looming large: the comparison between beauty and natural function is surely compelling, but *what* does it amount to? Can we dispense with the scare quotes I have been putting around words like ‘ought’ and ‘supposed’, and, if not, can we dispense with the words?

Kant, of course, trod similar ground well over a century before any of the machinery I have used to get to this place was available. He not only uses the language of ‘ought’ in talking about the demand made by a beautiful object on our response to it, but is so struck by the affinity between beauty and teleology in nature that the *Critique of Judgment* is about them both together. His characterization of this affinity, however, is quite obscure. By the Third Moment of the *Critique of Judgment*, where he raises the issue of function (‘purposiveness’), he has

already established that judgments of beauty are nonconceptual, and this means (among other things) that such judgments cannot rely on the cognitive recognition that an object fulfills a given determinate purpose. It can be no such purpose that makes an object beautiful. Yet it is, at the least, *as if* it has a purpose as the ground of its beauty, a purpose that is twinned with the aesthetic sensibility of the subject, so that the object fulfills its purpose in being recognized as beautiful and the subject fulfills her aesthetic purpose in so recognizing it. Kant says that a judgment of beauty “has as its ground a merely formal purposiveness, i.e., a purposiveness without an end” (*CJ* §15, 5:226), and this description seems to gibe with an ‘as if’ characterization. Yet it is still obscure what such a characterization could actually amount to, and in any case Kant is more ambitious than to leave it at ‘as if’. ‘As if’ suggests the possibility that we are making it all up, that we are simply pretending that beautiful objects make any kind of demand and that we are responding to that demand when we recognize beauty. Kant denies ‘as if’; he tries to defend the legitimacy of our judgments of beauty; his argument that pleasure in beauty depends on the free play of the cognitive faculties is supposed to explain this legitimacy. The free play is supposed to help elucidate beauty’s paradoxical ‘purposiveness without a purpose’.

Whatever the status of the free play, I agree with Kant that there must be a legitimate distinction between judgments of beauty and judgments of mere agreeableness, and that the comparison with natural function is useful. My hunch, moreover—again, with Kant—is that the comparison provides more than merely structural similarities; that there is something deeper to it that would help to explain judgments of beauty, if only we could get at it. In any case, though, explaining the legitimacy of judgments of beauty seems to me to be in much the same boat as explaining the legitimacy of moral judgments: in both cases there is the recurrent, niggling worry that in the end, it all comes down to personal preference, and that any normative demand purporting to transcend personal preference must be a fraud.

One more thing the comparison with natural function can do, though, is suggest an avenue for thinking about the nature of that normative demand for beauty, if we take it to be legitimate (as I do). It is unclear what kind of normative demand it could actually be. We can’t *obey* a demand to feel pleasure in a beautiful thing: our response is not in that way subject to the will. Beauty’s normative demand does not seem, then, to be a moral demand or any kind of demand of practical reason (“What should I do?”). Nor does it seem to be any kind of cognitive or epistemic demand (“What should I think?”). Beauty’s nonconceptuality—its connection with the subject’s feeling, rather than just her belief—precludes such an assimilation: I do not have the right to the belief that an object is (or must be) beautiful unless I am disposed to feel pleasure in it, and this disposition is not fully responsive to conceptual or epistemic reasons for belief. However, as we have seen, the language of normativity is also at home in descriptions of the functions of things: what things ought to do, because that is their function. The suggestion is just that some objects make a certain demand on one’s response to them, given one’s aesthetic role with respect to them. Withholding one’s aesthetic judgment about a particular object would then be a reflection of one’s sense that one cannot adequately fulfill this role with respect to the object. This seems to me to be exactly what we do think about ourselves, when we withhold aesthetic judgment.

I have called this section ‘The structure of beauty claims’, and while I have made some suggestions as to how we might think about the dispositions attributed by such claims, I haven’t yet talked about the structure of the claims themselves. I want now to sketch what these claims

might look like for contextualist and relativist views, given what I have said about the dispositions. The comparison with tastiness will continue to be helpful.

Contextualist renderings for both “X is tasty” (understood as attributing the disposition expected in typical contexts to a particular, as I have discussed) and “X is beautiful” presumably look very similar. S’s utterance of “X is tasty” will be true when X is disposed to taste good according to S’s normal standard of taste, whether we take ‘normal’ to be a statistical or a functional notion. S’s utterance of “X is beautiful” will be true when X is disposed to be pleasing to what we might call S’s normative aesthetic standard. That S has a normal standard of taste (in the former case), and a normative aesthetic standard (in the latter), are presumably presuppositions, under these renderings.

The relativist case is a little more complicated, and there are presumably different ways to spell it out. For a truth-value relativism such as MacFarlane’s (2014: §4.2; also 152), two people who disagree about the tastiness of X, and therefore assess the truth of “X is tasty” from the vantage of different standards of taste, are nevertheless discussing the same proposition: simply the proposition that X is tasty. Being more explicit about the fact that tastiness is a disposition should allow us to paraphrase this as the proposition that X is disposed to taste good. But which exact disposition is it that is their common topic of discussion? Is it (at least in typical contexts) X’s disposition to taste good under normal conditions? Or is a more specific identification of the disposition *not* common between them, so that not only *whose* standard of taste is relevant, but also *what kind* of standard it is (normal, toothpaste-addled, drug-induced), may change with the context of assessment? That is, is their common topic of discussion simply that X is disposed to taste good (period)?

The latter option is more flexible, and for that reason may often be more or less in tune with the hopelessly muddled conversations we often have about these things; but it isn’t a good option. We do genuinely coordinate on which disposition we mean to be talking about, when we bother to be clear about it; we don’t get the typical assertions and denials that are taken to be evidence for relativist views, when it is only a known difference concerning which kind of disposition is salient that is at odds between interlocutors:

- A: This fresh-squeezed orange juice is delicious.
B: [With the same normal orange-juice-tastes as A, but having just brushed her teeth]
?? No way! It’s really gross!

B is much more likely to have said something like, “Glad to hear it—but this toothpaste makes it taste awful.”

So, it seems to be, in a typical context, that both interlocutors take their common topic of discussion to be X’s disposition to taste good under normal conditions. “X is tasty” will then be true as assessed by someone whose normal standard of taste evaluates X favorably and false as assessed by someone whose normal standard of taste does not evaluate X favorably. Note, as with contextualism, that we must presuppose each agent to have a normal standard of taste.

Treating “X is beautiful” similarly, under a relativist construal, would require the common topic of discussion between interlocutors to be X’s disposition to be (as we might say) *normatively pleasing*. Analogously with ‘tasty’, “X is beautiful” would then be true as assessed by someone whose normative aesthetic standard evaluates X favorably and false as assessed by someone whose normative aesthetic standard does not evaluate X favorably.

In one way this is quite a nice result, and in another it is kind of strange. What is nice about it is that, on this construal, the topic under discussion is the proposition that X is normatively pleasing—that it is disposed to produce pleasure in what I have called a ‘universal’ judge. This proposition is independent of any actual judge, and does seem to be what is in dispute between people who disagree about X’s beauty. This apparent advantage of the relativist construal over the contextualist construal reflects, of course, one of the general advantages touted by proponents of relativism, who claim for the relevant domains that relativism accounts for disagreement better than contextualism does.

What is strange about it, though, is that there can be no sense in which people who disagree about X’s beauty—one person claims that X is beautiful and the other claims that X is not beautiful—can both be faultless in their claims. It is a hallmark of relativism that it aims to account for disagreement even when both parties have made no error of any kind in their assertions: I say X is tasty when it tastes good to me, and you say X is not tasty when it does not taste good to you, and those are the claims we each *should* make, given our standards of taste. But if we disagree about X’s beauty, we can’t both have the requisite normative aesthetic standard: if X makes an aesthetic normative demand, it makes the same demand of everyone. So while we can presuppose that each party to a dispute about tastiness has a normal standard of taste (i.e. normal for them), we cannot presuppose that each party to a dispute about beauty has a normative aesthetic standard. The upshot is that relativism seems wasted on beauty.

The shape for beauty claims that emerges from the attempt to fit them to relativism appears as stubbornly hybrid as the contextualist version, though in a different way. It is as if the relativist attempt reduces down to giving beauty claims an objectivist shape after all, but with a subjective presupposition: as if the primary content of “X is beautiful” is the objectivist proposition that X is disposed to be normatively pleasing, but with the essential presupposition that the subject herself has the requisite normative aesthetic standard—that X is disposed to be normatively pleasing *to her*. Such a presupposition would be quite unlike canonical examples of presupposition: in the example I gave earlier, for instance (§4.2)—“Margot’s car is red”—its primary content is logically dependent on its presuppositional content, and this is not the case with the putative primary content and presuppositional content of “X is beautiful.” It is only by making the comparison with tastiness that we have seen how the two dimensions of content for ‘beautiful’ may be so intimately connected together.

I don’t ultimately know how to resolve these tangled issues. What I hope I have established is that there are two strands of meaning in the claim that X is beautiful—even after putting its Direct Evidential implicature aside—however we ultimately decide to organize them: the claim that X is disposed to produce pleasure in a universal judge, and the claim that the subject occupies the perspective of such a judge with respect to X. Contextualism runs both claims together, and in so doing appears to privilege the second, beauty’s subjective aspect; the alternative, which I don’t quite know how to classify, privileges the objective first claim as more or less ‘at issue’ and gives an essential but subordinate role to the subjective second. I think I prefer the latter alternative but nothing I will say in the following sections depends on making a choice here.

7.4 Negative judgments of beauty

In this section and the next I will draw out some of the consequences of what I have discerned about the structure of beauty claims. This section will concern the judgment that X is not beautiful; the next will concern disagreements about beauty.

I have indicated something of what goes on when someone judges, on the basis of experience, that an object is not beautiful. On the contextualist rendering just outlined, this is the judgment that X is not disposed to be pleasing to the subject's normative aesthetic standard (not disposed to be pleasing to the subject in her capacity as a universal judge). On the alternative I presented, it is the judgment that X is not disposed to be normatively pleasing (period), still presupposing, however, that the subject has a normative aesthetic standard with respect to X. This difference in shape between these two accounts will not affect what follows.

The kind of experience that a subject bases a negative judgment of beauty upon is typically a failure to feel a pleasure in the object that she can take to be a normative pleasure. However, as I have hinted above, there is an asymmetry between positive and negative judgments of beauty, made firsthand, with respect to how they are based on the experiences that give rise to them. The relationship between the pleasurable experience that is the basis of a judgment of beauty, and the dispositions that are attributed by it, is very different from the relationship between any experience that is the basis of a judgment that a thing is not beautiful, and the relevant dispositions for the negative judgment. This asymmetry is due to the special nature of the kind of pleasurable experience that beauty gives rise to.

No such asymmetry is present for the judgments that something is tasty or that it is not tasty. I base the former on a good taste experience that I take to be a manifestation of my normal taste dispositions with respect to the food; I base the latter on an indifferent or outright bad taste experience that I take to be a manifestation of my normal taste dispositions with respect to the food. Both kinds of taste experience have the same relationship with my judgment concerning the relevant dispositions. As I said above, the experiences themselves usually indicate to me whether or not they stem from my normally functioning sense of taste. But in neither case does the experience's intrinsic quality, in this respect, *determine* whether or not it is normal for me to have experiences of that kind in eating the food: good taste experiences, or indifferent-to-bad taste experiences. In neither case does the experience's intrinsic quality make it the case that my normal dispositions are what they are. In both cases, then, I have the same materials to work with (so to speak) in *judging* my normal dispositions to be what they are: the intrinsic qualities of my good or not-good taste experiences, plus whatever other information I may have about what my normal dispositions are likely to be. (If I have just brushed my teeth, for instance, I may be wary of basing a tastiness judgment on my current taste experience, even if it *seems* normal to me.)

However, the complexity of the pleasurable response through which a judgment of beauty is made, gives this response a fundamentally different relationship with the positive judgment than the mere *lack* of such a response can have with the negative judgment. As I said above, Kant stresses that the subject's pleasure in a judgment of beauty must be produced by the object in a way that is itself sensitive to the normative relevance of the object to that pleasure. But ostensibly being aware of the object's normative relevance in this way, through one's pleasure, just is ostensibly to be aware that one is a universal judge of beauty with respect to the object: being aware of the object's merit, through one's pleasure, and being aware that one's pleasure is appropriate to the object, are two sides of the same relational awareness. Attributing the relevant

dispositions to the object as a result of this normative pleasure—the disposition to produce a normative pleasure in the right circumstances, and moreover to produce such a pleasure in the subject—does not require any further judgment on the part of the subject, on any other grounds, about whether or not she is a universal judge of beauty with respect to the object. In fact, given the nonconceptuality of judgments of beauty, there could simply be no further judgment of this kind, if the judgment of beauty is based on the normative pleasure in the way currently under discussion: no further judgment that would *supplement* the awareness provided by the normative pleasure in a way that would turn the ability to make some kind of partial judgment of beauty into the ability to make a full judgment of it.

Of course, it is possible to come to doubt a judgment of beauty one has made: maybe the object is actually quite gaudy (though you love that kind of thing) and not really beautiful at all. This is an example of the kind of reservation about our judgment with which I began my discussion of judgments of beauty. In such a case the subject comes to doubt that what seemed to be a sensitivity to normative features was really any such thing. Our judgment may certainly be influenced in this kind of case by general considerations about whether we are likely to be good judges of the matter at hand. But the judgment that such general considerations may influence is the judgment as to whether or not one's original judgment should be taken at face value. My taking myself to be generally a good judge of such things influences my *confidence* in my judgment, rather than adding more *evidence* for it.

When I judge that an object is not beautiful, on the basis of my experience of it, my experience must of course be something other than a pleasure in the object that I take to be normative. Symmetry for positive and negative judgments of beauty would seem to demand that a negative judgment be made on the basis of a *displeasure* (or indifference) that the subject takes to be normative. Nothing I have said rules out the possibility of normative displeasures—maybe judgments of the ugly, as distinct from judgments that an object is not beautiful, could be said to be based on experiences of this kind¹¹⁵—but it is clear that they cannot be necessary for making firsthand judgments that an object is not beautiful. I can after all make such a judgment on the basis of what I take to be an idiosyncratic pleasure, as well as on the basis of my indifference to or dislike of the object. What is important about these experiences for the negative judgment of beauty is simply that they are not normative pleasures. If I judge, on their basis, that an object is not beautiful, it is because I am prepared to make a certain inference: that because the object does not look beautiful to me, it must not be beautiful. Underlying this inference—what makes me prepared to make it—is my confidence that in this matter I am a universal judge. Refusing to make the inference—withholding judgment, because one believes or at least suspects that one is not a universal judge—can be a sign of aesthetic wisdom; or, alternatively, a sign of possibly neurotic aesthetic insecurity. But whatever level of confidence I have in my ability to act as a universal judge with respect to the object, must be founded in this case on general considerations about this ability. For me to consider myself such a judge, I must be prepared to agree that if the object did deserve to be called beautiful, I would feel the requisite normative pleasure in it, and there is no particular experience that I can base this latter conditional judgment upon.¹¹⁶

¹¹⁵ Henry Allison takes it as a criterion of adequacy for any interpretation of Kant that it allow for the possibility of normative displeasures. (E.g. 2001: Ch. 3) Needless to say, not everyone agrees with him: Ginsborg's (2003: 166) interpretation of Kant, for instance, allows no room for normative displeasures.

¹¹⁶ If there are normative displeasures on which judgments of ugliness can be based, symmetrically with the normative pleasures for judgments of beauty, then, given that the judgment that X is ugly would entail that X is not

Making this same conditional judgment is part of what is required for me to make a positive judgment of beauty that is not made firsthand: for judging that an object must be beautiful on the basis of good testimony, say. I must take myself, in this way, to be a universal judge of the object, and I must of course also take it that the object does deserve its pleasure, so that my capacity to be a universal judge with respect to it would be made manifest to me by means of a normative pleasure that I would take in it under the appropriate circumstances.

Testimony is involved in making this kind of judgment in a way that is clearly very different from how it is involved in the acquisition of objective beliefs. The content of an objective belief is simply transferred, through testimony, from one believer to another: whatever standard governs the truth of the matter is understood to be an objective standard, applicable to both beliefs. If I acquire a belief about beauty through testimony, however, the standard governing my acquisition of the belief is a different standard from the one that is used by my informant. She believes that the object is disposed to cause pleasure in her, in her capacity as a universal judge; I come to believe, on the basis of her testimony, that the object is disposed to cause pleasure in me, in my capacity as a universal judge. This kind of change in standards across the testimonial divide also operates, of course, for matters of personal taste, and for many of the other Direct Evidential matters that I have discussed. In the aesthetic case, though, the judgment that I must make about the object and about my own dispositions, when my judgment is not based on my own direct experience, is particularly complex.

7.5 Disagreement about beauty

I originally brought up the McDowellian objectivist view as a way to acknowledge the normativity apparently at work in discourse about beauty. Normative claims purport to have universal validity, and it therefore seems that they should be governed by a single standard, applicable to everyone. Yet the overall shape of discourse about beauty does not confirm the availability of such a standard for everyone. For typical objective matters, a single standard is understood to be universally applicable, even when many people cannot apply the standard successfully themselves, and must rely on those who can for the relevant judgments. This is how it is for the judgment about the look of the Vermeer-like painting, and whether, visually, it can pass for a Vermeer. This is how it is for all matters for which there is expertise. Even if we think that the standard for aesthetic claims must be a critical standard of the kind McDowell elucidates, rather than the more concretely specifiable kind of standard typical for ordinary object matters, we can see that a critical standard of this kind is not wholly responsible for people's claims about beauty. I can agree that something must merit the pleasure appropriate to beauty without being disposed to feel pleasure in it myself, and in such a case I do not say that the thing must be beautiful. The standards that are relevant to my claims about beauty are

beautiful, there is a sense in which the latter judgment *could* be based on a normative displeasure in a way that is symmetrical to how judgments of beauty are based on normative pleasures: the subject would judge on the basis of a normative displeasure that X is ugly, and then simply infer that X is not beautiful. My point here is that even if this is the case, not all judgments that X is not beautiful can be based on normative displeasures. There are surely countless cases in which it would be hard to find a normative demand for displeasure in an object judged not to be beautiful: this eraser on my desk, for instance, is plainly not beautiful, but I would have to view someone who told me I ought not to like it as bizarre. Three people may judge that the eraser is not beautiful on the basis, respectively, of their idiosyncratic pleasure in it, their indifference, or their dislike; the latter two need not claim any normative requirement for their responses (could there be a normative indifference? maybe), and moreover can be said to agree with each other and with the first about the object's lack of beauty.

dependent on my propensity for liking, and if I have no such propensity in a given case, then I am simply denied access to a standard of beauty that would allow me to judge, or claim, that the thing must be beautiful.

Yet, because of beauty's normative element, some aspects of a genuinely objective discourse are still visible for the beautiful, as a result of the normative constraints relevant to beauty being, after all, in a certain way applicable to everyone's experience. This is because everyone's beauty judgments are still subject to criticism in the same way, even if the scope of reasonable criticism for such judgments is narrower than for typical objective matters, in the way that I will go on to describe below.

I will once again concentrate on firsthand judgments. When someone makes a claim about tastiness that we think reveals her bad taste ("Cheez Whiz is really tasty"), we may criticize her standards ("Wow—did you grow up in a gas station mini-mart?"), but we must acknowledge that her tastiness judgment itself, given that her standards are what they are, is the appropriate judgment for her to have made. Given that she likes the taste of Cheez Whiz, she is obliged to think that Cheez Whiz is tasty; and we acknowledge this, even though we think that Cheez Whiz is awful.

But for judgments of beauty the situation is different. We may be critical not just of another's response—that she is pleased by what is not worth it, or unmoved by what should be valued—but we may also be critical of the judgment concerning beauty that she makes by means of that response. If she likes something that we think is trash, we do not admit that she is obliged to think it beautiful. She is pleased by it, but her resulting judgment that it is beautiful depends on her taking herself to be a universal judge with respect to the object, and we do not think she is right about this. Similarly, if she thinks that Macbeth's "She should have died hereafter" speech (or whatever) is just boring—certainly not beautiful—we may object to her judgment, not just because we are inclined to criticize her standards, but because of the particular way we hold her judgment to account against the backdrop of the speech's merits. She presumably cannot help her dislike—at least, not in this instant—but a better judge would have acknowledged that in this matter she is simply not a universal judge, and withheld her judgment accordingly.

This means that when I say that X is beautiful and you say that X is not beautiful, we both think the other's claim to be at fault. We disagree not just in our attitude to X—in liking or not liking it—or even just in our evaluations of each other's standards as good standards or not, but in the *claims* we make about the kind of response that is appropriate to X. Part of my claim is that a universal judge of X is bound to like it, and part of yours is that a universal judge of X is bound to no such thing. While this is not a typical objective matter with a concretely specifiable standard available for resolving the issue, it is at least subject to a critical standard that applies in the same way to both sides, in a way that can at least be partially articulated. I think that you are not a universal judge in this matter, and that your claim should therefore be retracted, even if your feeling cannot be brought to match mine in a way that would allow us both to affirm that X is beautiful; and, of course, you think the same about me: that although I seem to myself to be a universal judge, in fact I am not, and I should retract my claim in favor of recognizing that X is not beautiful, although I happen to like it.

It is worth noting, though, that the asymmetry I discussed above between positive and negative judgments of beauty has an interesting asymmetrical consequence for disagreements about beauty. It seems to me that, when we are faced with the disagreement of others, we are much more likely to reconsider and retract negative claims about beauty, made firsthand, than

firsthand positive claims. We are much more likely, that is, to suspect that the responses on which our negative judgments are based, are in some sense deficient, than we are to think that our positive responses are unmerited. The asymmetry between positive and negative judgments readily suggests an explanation of this phenomenon. The positive judgment is based on a pleasure in the object that the subject takes to be normative, and which therefore seems simply to reveal the object's beauty to her. Her pleasure seems to her to be a direct experience of the object's beauty. But the indifferent or negative experience on which the negative judgment is typically based is not similarly revelatory of the object's lack of beauty. It is, rather, one element of an inferential judgment that the subject must make, which depends on her taking herself to be a universal judge of the object, which in such a case is itself a complex judgment that is not made by means of any particular experience. It is thus much easier to doubt the negative judgment—to suppose that one is after all not a universal judge and is simply missing something important about the object—than to doubt the positive judgment. Doubting the negative judgment just requires thinking one may be blind, but doubting the positive judgment requires thinking one may have been *duped*.

It is also worth noting that in spite of the possibility of being 'at fault' that I describe here, it is generally just mean-spirited and rude to flatly impugn others' aesthetic points of view when they differ from our own. Civil aesthetic engagement with others involves exploring our different responses with generosity, rather than dismissing other people as aesthetically hopeless. But we should be clear about the structure of the judgment, which does make genuine disagreement possible (even if rude to belabor): if you think something is beautiful and I disagree, I am (in part) denying that you are a universal judge of the object in question. It is easy to be uncomfortable with this possibility: how dare I? Yet if we react to this discomfort by erasing the distinction between being beautiful and seeming beautiful—by saying, in effect, that whenever a subject takes herself to be a universal judge of an object, she just is one—then of course we thereby erase the distinctively normative content of the judgment of beauty, and 'taking oneself to be a universal judge' of an object turns out to have no more significance than merely taking oneself to judge it, by liking it, or not.

The discomfort is of course familiar from ethics. Moral relativism as a metaethical thesis grew to some prominence in the 20th century in anthropology, in reaction to a history of Western engagement with other cultures that uncritically assumed the superiority of Western cultural values. Clearly such an assumption is wrong-headed: it turns out, for instance, that Victorian corsetry is not essential to life well lived for women in sub-Saharan Africa. Even so, however—and even given more recent interest in moral relativism by philosophers—rejecting this wrong-headed assumption does not require adopting moral relativism. Relativizing the truth of moral claims to the attitudes or practices of societal groups (or individuals) makes it obscure at best how we are to accommodate the critical dimension of moral life. The claim that X is morally wrong is no longer obviously a normative claim, but seems instead to be a descriptive claim concerning the attitudes or practices of some group or individual. This difficulty with moral relativism is of course well known.

What rejecting the wrong-headed assumption, and others like it, does require, if we are to continue to make normative judgments, is a commitment to *good* criticism. This is true of aesthetic as well as ethical criticism. Such a critical practice will involve withholding judgment a lot of the time. The discomfort we may feel with normative judgment comes from the accurate sense that normative judgment invites snobbery, and the various social power plays that come

along with snobbery. This is a genuine pitfall, but the price of refusing to risk it at all is giving up on normative judgment altogether. The trick, mastered by people we tend to call ‘wise’, is to wield normative judgment as an invitation of a different kind.¹¹⁷

It is worth noting here as well that people generally don’t bother pursuing disagreements about beauty if their beliefs have not been acquired on the basis of their own experience. It is the *experience* of beauty and the ways in which different objects give rise to it that have primary interest for us in thinking and talking about beauty. The rest of the machinery I have introduced here in order to make sense of this thinking and talking clearly dances attendance on the primacy of this experience. Even when our opposing beliefs are the result of personal experience, it is not really the opposition itself (you think it’s beautiful, but I don’t), but our differing reasons for it—that is, the different ways in which we have engaged with the object—that are the substance of our discussion. In the absence of this kind of engagement with the object, dwelling on our opposition is generally not interesting and usually has no other point either. The only kind of context I can even think of that would give any oomph to a secondhand argument of this kind is that of, say, two people trying to decide which store to go to or which designer to approach in order to find something with a particularly beautiful purpose: a wedding dress, maybe. I can’t even say that such secondhand arguments generally crop up about beauty in art—for people trying to decide which movie to see or which gallery to visit—as, while people certainly do argue in this way, before the fact, about which movie or which gallery is most worth seeing, the kind of artistic merit that people are looking for in such cases comprises much more than just beauty.

Finally: I said above that although discourse about beauty has what we might think of as partial objectivity, the scope of reasonable criticism for beauty judgments is narrower than for typical objective matters. This is because the kinds of possible criticism I have described—in which we object to another’s judgment because we think she is not in fact a universal judge of the matter at hand—do not allow that person to accept a positive judgment about beauty that she could not make on the basis of her own negative experience, even if she accepts the criticism. Just as we can’t insist that someone agree that a dish must be delicious if she has no propensity to like it—this would be the wrong claim for her to make—we can’t insist that someone agree that a thing must be beautiful if she has no propensity to like it. This, too, would be the wrong claim for her to make. The normativity in play for judgments of beauty does not wholly constrain such judgments, independently of the subject’s liking. The normative element of judgments of beauty only partially constrains such judgments, in the ways that I have described.

I would say, then—although Kant would presumably object to this—that judgments of beauty cannot be universally valid to exactly the same extent that ordinary empirical judgments are. Kant in one way appears to agree: he certainly thinks that someone can legitimately agree with my claim about an object’s beauty only if she, like me, feels pleasure in the object. Of course, Kant also says that a judgment of beauty claims that everyone ought to share one’s pleasure in the object, and it is for this reason that everyone ought to agree with the judgment; and we might take this after all to argue for an exact parallel between judgments of beauty and ordinary empirical judgments. Someone can legitimately agree with my claim that the cat is on the mat only if she also believes the cat to be on the mat; and surely she ought to believe this, if she is staring right at it. There is still a difference, however. I can acknowledge that your claim

¹¹⁷ It’s hard to be wise. Maybe I shouldn’t have made that crack about Cheez Whiz. Incidentally, for an entertaining account of one man’s interrogation of what is either his good taste or his aesthetic snobbery, see Carl Wilson (2014).

that an object is beautiful must answer to the object's aesthetic merits, while finding the object disagreeable myself, and thus not being able to agree that the object must be beautiful. I cannot similarly acknowledge that your claim that the cat is on the mat answers to reality, while at the same time refusing to agree with it. You might think that I've got the parallel wrong here—that really I should be comparing my acknowledgement of the object's aesthetic merits with my acknowledgement of your cat-on-the-mat claim's *well-foundedness from your point of view*, which is perfectly compatible with my refusal to agree: my epistemic position may just be better than yours. But in acknowledging that you can appreciate aesthetic merits that I am blind to, I acknowledge that your aesthetic point of view is in fact better than mine. I acknowledge, then, something that is analogous to your experiencing the cat to be how it actually is. Yet pleasure in an object simply cannot be rationally compelled: we are rationally in charge of our aesthetic *training* but only partially in charge of our aesthetic responses at any given time (I may, for instance, try to evoke a particular response by directing my attention in a certain way, but the response itself is not entirely up to me). In matters of beauty, we can acknowledge ideals that we also acknowledge we cannot live up to. It is this possible misalignment between our own capacities and what we take to be the capacity to judge universally that is behind the fact that judgments of beauty slip partially through the net of universal validity in the way that they do, in spite of their normative element.

7.6 Calculability for aesthetic claims

I have been arguing throughout that simple aesthetic claims carry a Direct Evidential implicature to the effect that they are made firsthand, and that any aesthetic claims not made on the basis of personal experience must indicate an appropriate evidential source (“It must be beautiful,” etc.). I have been led to this view by considering what the primary content of such claims must be like, and how we can come to know about it—not just, it turns out, through firsthand experience, but in other ways as well. The phenomenon in question, I have argued, is present for a variety of different kinds of claims, and I have undertaken an extensive comparison with gustatory taste. Even for tastiness claims, however, I have not pushed a view about how exactly we should model their primary content, leaving both contextualism and relativism on the table. Claims about beauty are in some ways more complex than tastiness claims, and I likewise have no precise model of their primary content to offer. I have here just tried to outline the general contours that such a model should aim to fit. We can now see how this kind of primary content for beauty claims, sketchy though my account of it is, lends itself to a plausible calculation of their Direct Evidential implicatures.

The challenge to accounting for this calculability that may have been presented by the normativity of beauty, concerns the possibility of positing a single standard for normative claims that is taken to be independent of the speaker's own propensity to judge *for herself*. Such a standard would take the form of a certain kind of external authority. For instance, even though the normative claims of morality are subject to the same kind of critical standards that I have discussed here for aesthetic claims—and I have suggested that this is enough, in this time and place, to put moral claims in the class of Direct Evidentials—it is possible to conceive of different times and places in which the strength of an external moral authority is sufficient to preempt many moral claims from exhibiting any Direct Evidential implicature. Such times and places have presumably been actual times and places. Widespread deference to religious authority, for instance, would presumably make it commonplace (presumably has made it

commonplace) for members of the flock to make moral claims on the basis of their shepherds' testimony.

Moral claims of rightness and wrongness are not based on feeling in the way that claims of tastiness and beauty are, although of course moral feelings are an important part of moral experience. The kind of deference to authority that is at least possible in the moral case is simply not possible for beauty, because of the way a speaker's own dispositions to respond aesthetically are involved in the understanding of the aesthetic claims she makes. We may in the way that I have described accept others' aesthetic testimony when we think that we agree in standards, but we may not defer to another just because we think her aesthetic judgment must be better than ours, when we suspect we would not share the same response. Any attempt to defer to others about beauty in this way would come across, if discovered, as a kind of deceit or evasion, even if an explicit evidential marker is used. We may think that someone else must be a suitable universal judge of an object's beauty without being able to adopt her belief.

Claims about beauty present a perfect storm of conditions for Direct Evidential calculability. The contentious nature of the critical standards that are relevant to them would be enough to explain their implicatures, in the same way as for moral claims (in the absence of widespread deference to moral authority). But the standards relevant to claims about beauty are also tied to a speaker's own dispositions to respond to things with liking and dislike, as they are for claims of personal taste. As I have said, a speaker typically knows how a thing measures up to her standard of tastiness (say) because she has in fact tasted it. Simple, positive tastiness judgments must locate a taste on an idiosyncratic scale that is subjectively measurable only for very large differences or for foods of the same kind; and such judgments must, moreover, take a stand on the food in question's being 'tasty enough' in the context to call tasty.

Although there is wide variation in people's resultant tastiness scales, we shouldn't overestimate it: we are willing to take on the tastiness beliefs of others quite a lot of the time, even when we don't know them that well (or at all—online crowd-sourced review sites provide at least some level of information, frustrating though they can be)—although of course our acceptance rates go up when we have particular trust in our informants and can rule out the presence of ingredients we know we dislike. The situation is somewhat different if we compare 'tasty' with 'delicious.' 'Delicious,' as I have mentioned, is a term of *lavish* approbation: you really have to be disposed to like something a lot to call it delicious. It is a term used for special occasions: occasions of exceptionally striking tastiness. We *may* be convinced through testimony that a food really is exceptional in this way: we may say that it must be delicious. But given the relatively small portion of the tastiness scale that delicious things occupy—I assume 'delicious' is more or less synonymous with 'very tasty' or 'exceptionally tasty'—it is that much harder to say on this basis, with the level of certainty required for simple, epistemically unqualified assertion, that a particular food is such as to be exceptionally striking in this way. For the most part, we wait to be *struck* by a food that is strikingly tasty enough to call 'delicious,' and so deliciousness claims are overwhelmingly—though not exclusively—used to express a speaker's striking experience of deliciousness. It is largely in this way—through experience of the delicious—that people are willing to take a stand on deliciousness.

We saw a couple of similar examples earlier. One was Pettit's example of the sadness of a picture, and another concerned my amateurish painting, about which I said, "It looks like a Vermeer." I pointed out that *is sad* and *looks like a Vermeer* are gradable expressions, and said that, in making claims of this kind, I must take a stand on the relevant cut-off points: on how

much the thing in question must be sad (in the way that a picture is sad), or look like a Vermeer, in order for the claim to be appropriate. I said that in these contexts, what gives me reason to make the claim is simply my experience of having been struck in a particular way by the object in question. I speculated that, for the sadness of a picture, I may be struck by the picture's similarity in various ways to a person's expression of sadness; for the supposed Vermeer lookalike, I make the claim as a result of being struck by some visual similarity between my painting and a Vermeer. In both cases, the position of the cut-off point seems to be constrained by the experience itself: the claims do not seem to laying down the law about how each picture looks, but come across instead as invitations to others to share my experience. It seems natural to say for these examples that my claims express my experience, given the close connection between the experience and my reason for making the claim.

'Beautiful' is, like 'delicious,' a term of lavish approbation. Judgments of beauty are similar to judgments of tastiness (and deliciousness) in being dependent on a combination of multiple factors, though 'beautiful' (like 'tasty' and 'delicious') is not itself a genuinely multidimensional term (cf. §5.3). The different features of an object that is judged to be beautiful are not, in that judgment, evaluated independently for beauty, and the individual evaluations then summed in a contextually appropriate way.¹¹⁸ Beauty is, like gustatory taste, a matter of combination and balance.

Unlike 'delicious,' however, 'beautiful' is not associated with the 'exceptional' portion of any particular evaluative scale. 'Beautiful' is often used to describe what is exceptionally good-looking or exceptionally pretty, but it is certainly not limited to this, and, at least in my experience, the most striking examples of beauty are substantially misdescribed by those adjectives, even when given the exceptional adverb. Many different kinds of properties can be relevant to a judgment of beauty, whereas only flavors, and some other sensory properties such as perceived texture and temperature, can be relevant to judgments of deliciousness. Beauty's promiscuity in this respect is of course also present for thin evaluative terms like 'excellent,' but 'beautiful' is surely not so thin, at least not in its primary usage.

We do speak of one thing being more beautiful than another, so beauty must be gradable to some extent, or must at least be usable in a gradable way. But the difficulty I mentioned earlier for 'tasty,' concerning the calibration of a single tastiness scale—a difficulty shared with

¹¹⁸ McNally and Stojanovic (2017: §3) say that 'beautiful' is in fact multidimensional, on the grounds that something can be beautiful in some/every way or beautiful except for It is true that 'beautiful' is different from 'tasty' and 'delicious' in this respect. But 'beautiful' is not additive like a truly multidimensional term such as 'intelligent': i.e. although we can consider individual features of an object (color, shape) and judge them independently for beauty, a beautiful object is not in general beautiful because of the independent beauty of its individual features. In fact, an object's beauty may depend on some of its features being ugly: it may be less beautiful, or not beautiful at all, if (say) some part of it that is an ugly brown is changed to a beautiful blue. Cf. Plato's *Republic* Book IV—brought up by that Plato-aware person previously mentioned: "Suppose, then, that someone came up to us while we were painting a statue and objected that, because we had painted the eyes (which are the most beautiful part) black rather than purple, we had not applied the most beautiful colors to the most beautiful parts of the statue. We'd think it reasonable to offer the following defense: "You mustn't expect us to paint the eyes so beautifully that they no longer appear to be eyes at all, and the same with the other parts. Rather you must look to see whether by dealing with each part appropriately, we are making the whole statue beautiful."" (Plato 1997b: 420c-d)

This contrast with 'intelligent' is what Sibley highlights in saying that the features relevant to a judgment of intelligence "always carry some weight and can count only in one direction. Being a good chess player can count only *towards* and not *against* intelligence" (cf. fn. 82).

genuinely multidimensional terms—is magnified for ‘beautiful,’ given the open-ended multiplicity of ways that things can be beautiful. Comparative judgments of tastiness are easiest for foods of the same kind because the different dimensions of flavor that are relevant to such judgments will be similar. The comparative judgments for beauty that we are actually inclined to make are generally even more constrained in this way (e.g. “This place is even more beautiful than it was last year”).

I have described how personal experience of a food is assumed by default to be a speaker’s source of knowledge about its tastiness, and even more so for its deliciousness. But the nature of the tastiness scale makes it possible to have other sources of tastiness knowledge, quite a lot of the time. I know where many things fall on my tastiness scale, including the deliciousness part of it. If I can roughly locate foods I’ve never tasted on the scale, relative to the foods whose positions I already know, then I can make tastiness judgments about those foods, as long as their positions on the scale are far enough from the relevant cut-off in the context to be unproblematic in that respect. I can do this if there are fairly robust correlations between the foods’ other properties (which I can know about without tasting them) and their tastiness. There often are such correlations: between tastiness and the praise of a trusted judge, between tastiness and consistent manufacture, etc. It is harder for such correlations to be definitive for deliciousness than for mere tastiness, as I have described, but it does happen.

There is only one kind of beauty that seems comparable to deliciousness in its level of such predictability, and that is natural beauty. Many have noted that judgments of natural beauty are in general far less contentious than judgments about the beauty of art and artifacts. Presumably this is because the many dimensions relevant to a judgment of natural beauty are similar from case to case, and attunement to these dimensions does not require any specific training or uncommon experience (it is ‘natural’ for people to appreciate natural beauty). Certainly the most straightforward examples of the acquisition of beauty beliefs on the basis of testimony concern natural beauty: Hawaii must be beautiful, I think (though I have never been there).

For other kinds of beauty, though, we run up against a usually insuperable level of unpredictability. Small changes to a beautiful object’s physical properties may wreak havoc on its beauty. That is perhaps not surprising. What may be more surprising is our resistance to correlating beauty with the testimony of trusted judges. I have spent a long time arguing that we are not *fully* resistant to this, and I do not mean to go back on that. But it is undeniable that the testimony even of trusted judges about manmade beauty overwhelmingly tends to produce a “wait and see” attitude in its recipients, who thereby acquire—until they have witnessed the object themselves—various degrees of belief about the object’s *likelihood* of being beautiful, without crossing over to certainty.

We frequently believe that a trusted judge is genuinely sensitive to the merits of an object whose beauty she is praising. What nevertheless usually prevents us from taking her judgment wholly on board is our inability to correlate even such trusted praise as this with our own disposition to respond to the object in the way required for thinking it beautiful. We just do not know, on the basis of her testimony alone, whether the object is disposed to strike *us* with the force that is necessary for calling it beautiful.

I have described two kinds of ambivalence that it is possible to have about a thing’s beauty, even when you have experienced it yourself. Sometimes your response to the object is itself ambivalent—you have mixed feelings—and you therefore aren’t sure whether *beauty* is quite the right property to attribute to the object. Sometimes you do not take yourself to be a universal

judge of the object, and thus do not take your negative response to indicate its lack of beauty. In both cases, you refrain from the judgment that the object is beautiful and from the judgment that the object is not beautiful. There is another kind of reason for refraining from judgment, even given firsthand experience of the object: having a response that is simply not strong enough to make you want to take a stand on the object's beauty. If you tell me a thing is beautiful and I trust your judgment, there is still significant room for me to wonder whether the thing will strike *me* sufficiently for me to call it beautiful.

The *Mona Lisa* is a famously beautiful object. Yet I was frankly more struck, as experiences go, by standing in front of it and then turning around to see the hordes of tourists *videoing* it, than I was by the sight of the lady herself. —Unlike my experience of turning the corner in the Accademia and catching sight of *David*, which lives in my memory as a moment of singular clarity. I am not willing to say, on the basis of my own experience, that the *Mona Lisa* is beautiful. This is not because I quarrel with anyone who says it is—I certainly don't think it's *not* beautiful. I believe that others must be sensitive to the *Mona Lisa*'s merits, so I'm just not sure, in this case, and the many others like it, whether it's me or circumstances that have failed: whether I fail in being a universal judge of the *Mona Lisa*, or whether the appropriate conditions for making manifest its disposition to produce (enough) pleasure in me, in my capacity as a universal judge, just never quite materialized. I'm not sure, and so I refrain from judgment.

I generally cannot tell, in the absence of experience—at least with art and artifacts—what I am disposed to find striking enough, enough of a stand-out against the background of normality, to count as beautiful. Before seeing the *Mona Lisa*, I may have thought it quite likely that I would find it beautiful; in the event, I was not particularly struck by it. We are overwhelmingly agnostic, in this way, about other people's beauty claims. We may, however, combine this agnosticism with a kind of acceptance of what they say. Given the nature of claims about beauty, this acceptance goes further than our acceptance of other people's deliciousness claims. I may be agnostic about a food's deliciousness (if I haven't yet tasted it) while yet acknowledging that you have the right to call it delicious, given your striking experience of it. I may be agnostic about an object's beauty while yet acknowledging not only that you were, in fact, struck by it, but that you must be a universal judge with respect to it.

Beauty claims are thus highly idiosyncratic—certainly more so than deliciousness claims—in spite of their claim to universal validity. This may have a slightly paradoxical air. Natural beauty is of course something of an exception here: people who appreciate natural beauty are alike in being attuned to the dimensions relevant to judgments of natural beauty, and do not need any special training to be attuned in this way. But the kind of perspective from which one can appreciate the beauties of art and artifacts may be highly specialized. This is perhaps most obvious for language arts: it is a minimal requirement for appreciating the beauties of literature that you understand the language in which it is written. But it is also the case for visual arts, architecture and music, as well as more everyday objects, for which it is clear that some styles have more mass appeal than others in part because they require less specialized knowledge and experience, or fewer quirks, in order to appreciate them (hence, for instance, all the Impressionist prints in dentists' waiting rooms). It can be easy *not* to have the highly specialized perspective that is required for resonating with a particular object's unique, multidimensional beauty. But in fact this in no way detracts from the claim to universality made by a person who knows, through the right kind of experience of an object, that she occupies the perhaps highly specialized perspective from which one can see that the object merits one's pleasure, and who therefore

judges the object to be beautiful. The claim can even legitimately be made by someone, like Kant's young poet, who insists on the thing's beauty against the indifference or condemnation of everyone else in the world. Kant's poet ends up changing his mind, but sometimes the lone wolf is simply ahead of his time.

7.7 Wrapping up aesthetic content and calculability

It is in this way that we take beauty claims to *express* a speaker's experience of a beautiful thing. They are similar in this respect to deliciousness claims, but the connection with experience is even more robust for beauty. It is not that the primary content of such a claim itself depends on the experience: it does not. Nevertheless, there is a tighter relationship between beauty claims and the experience of beauty than there is in many other cases of Direct Evidential implicature, simply because of the more limited use that is made of beauty claims overall. This relationship may seem distorted simply by saying, as in these other cases, that the default epistemic route to the belief behind the claim runs through the experience. This is true, but it ignores the fact that our interest in the *property* of beauty does not go very much further than our interest in the *experience* of beauty, an interest that is reflected in the nature of beauty discourse. This is unlike our typical interest in objective properties. It generally doesn't matter to us how we acquire our knowledge of objective properties: through our own experience or through the experience of others. This is so even for plausibly dispositional, experiential properties like colors.

Calling something beautiful ascribes to it a subjective yet normative dispositional property that is idiosyncratically recognized and not predictably correlated with other properties. Recognizing an object's beauty thus has fairly limited further practical use. The importance of knowing about the property itself, apart from having experiences of the property, is just to arrange the world in such a way as to bring about more experiences of that kind, if one likes. This may have practical value: if you have a good eye, you may want to run a boutique or be an interior designer; your eye is 'good' in such a case because the things you find beautiful, other people find beautiful too. More cynically, if you don't care about beauty yourself, you may still find a way to make money from other people's taste. But, either way, such practical value simply depends on other people's interest in beauty, which, like one's own—if one has it—is fundamentally an interest in being able to experience it oneself.

We don't care about having the experiences, then, because we want to know about the property; it is rather that talking about the property allows us to rely on a sort of minimal ontological structure for organizing our thoughts about our experiences and the objects that give rise to them. This is what the primary content of beauty claims appears designed to let us do.

I have not talked much, lately, about other aesthetic properties besides beauty. They clearly comprise a very mixed bag; but at this point I think we should have covered the many slightly different reasons why claims that weigh in on them are made as a result of the speaker's own experience or judgment. Maybe some relate to feeling in the way that claims of personal taste do; maybe some that relate to feeling also have a normative element, like beauty. Evaluative aesthetic claims about art, like the claim with which I began, all those pages ago—"Middlemarch is a great novel"—relate to feeling in a way that is not at all straightforward: their explicit concern is artistic merit, which has some kind of complex relationship with artistic response that is beyond my ability to sort out right now. They are subject to critical standards like both moral claims and claims about beauty, but I am not sure which they more closely resemble in their

relationship to feeling. Maybe they sit somewhere between. I may come to believe that some behavior is morally permissible—must be permissible, I might say, when I am still grappling with the issue—while still being squeamish in my feelings about it: in such a case, my feelings haven't yet caught up with my belief, but moral belief does not require such agreement in feeling. It is not so clear to me that I can likewise say, "Mondrian must be a great artist," baffled as I am by pieces that seem to me to have all the appeal of Lego—even though I do believe, given what artists, critics and museum executives themselves believe, that Mondrian's work must have artistic merit. *That* seems all right to say, but the claim "Mondrian must be great" seems dubious. Possibly the issue here is not so much to do with aesthetic judgment, though, as it is with the use of positive evaluative adjectives, which would explain why the former is fine but the latter questionable, in spite of my attempt to mean much the same thing by both.

8. CANCELABILITY

It is now time to address the issue I left hanging in Chapter 4: the fact that the evidential implicatures of unqualified aesthetic claims do not seem to be cancelable. As I said earlier, it is natural to think that conversational implicatures would have to be cancelable, as, by definition, an utterance's primary content is consistent with the negation of any content it may implicate. This is certainly Grice's view:

[A] putative conversational implicature that *p* is explicitly cancelable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add *but not p*, or *I do not mean to imply that p*, and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. Now I think that all conversational implicatures are cancelable ... (1989b: 44)¹¹⁹

After all, if in *every* context the utterance of the form of words carries the secondary content in question, and in *none* of these contexts is the explicit disavowal of this content admissible, then surely the content in question is attached to the form of words semantically rather than pragmatically. Even generalized conversational implicatures (GCIs), which are attached to a form of words more tightly than is the case for particularized implicatures, and which therefore arise in many more contexts of utterance, are typically explicitly cancelable, and are often 'contextually' cancelable as well.¹²⁰ I have given some examples of explicit cancelation of a GCI, such as "Some of the kindergartners can already read; in fact, all of them can"; and there are also plenty of examples of GCI contextual cancelation. For instance, Levinson argues that the numerals form a Horn scale, so that saying e.g. "I have nine dollars" generally implicates "I don't have ten dollars"; but there are contexts in which the implicature fails to arise:

[Outside a movie theater]

A: It's going to cost us ten dollars to get in and I didn't bring a cent.

B: Don't worry, I've got ten dollars. (2001: 52)

Levinson's explanation, which is surely correct, is that what is overwhelmingly relevant in this conversation is having *at least* ten dollars, so that A and B can go see the movie. B's utterance therefore does not implicate that she has no more than that, as her total amount of money is in this case not remotely at issue.

Earlier (cf. §4.1), I discussed Chierchia's claim that a sentence's typical scalar implicature is retained, as part of what is believed, when that sentence is embedded in a belief context. He uses the example "John believes that some students are waiting for him" to illustrate this, saying (to my mind, very plausibly) that this belief attribution carries the implicature that John believes that

¹¹⁹ Grice goes on to say that passing a cancelability test is nevertheless not sufficient for diagnosing conversational implicature (though he regards it as necessary). One reason he offers concerns 'loose' ways of speaking: in saying "Macbeth sees Banquo, but he's hallucinating," for instance, the speaker does not cancel an *implicature* (generated by 'sees') to the effect that Banquo exists. That the seen object exists is part of the core meaning of 'see,' Grice says, even though there are loose uses of the term that suggest otherwise.

¹²⁰ As many others have noted, contextual 'cancelation' is a bit of a misnomer for this phenomenon. It is not that the implicature arises and is then canceled, as for explicit cancelations, but rather that it doesn't arise in the first place. I will nevertheless continue to use the term, given the position it now has in the literature since introduced by Grice.

not every student is waiting for him. Nevertheless, it is possible to attribute to John both the belief that some students are waiting for him, *and* the belief that all of the students are waiting for him, by saying (e.g.) “John believes that some—indeed, all—of the students are waiting for him.” In this case, the implicature has been explicitly canceled. It is also possible to make a belief attribution that would typically carry a GCI in this way, in a context in which the implicature simply doesn’t arise. We could, for instance, alter Levinson’s dialogue thus:

[Outside a movie theater]

A: It’s going to cost us ten dollars to get in and I didn’t bring a cent.

B: Don’t worry—C thinks she’s got ten dollars.

There are two related difficulties, then, for the aesthetic case. One difficulty is that, as aesthetic claims do not seem to be either explicitly or contextually cancelable, we are hard pressed to attribute the evidential phenomena of such claims to GCI in the first place. The related difficulty is that, without a GCI-based explanation of the evidential implications of aesthetic claims that allows for cancellation, we are hard pressed to understand the nature of aesthetic beliefs in a way that accounts for their primary content as I have been trying to do. Just as we want to be able to attribute to John the belief that some of the students are waiting for him, even when he also believes that they are all waiting for him, we want to be able to attribute to someone belief in the *primary content* of what is expressed by saying that Hawaii is beautiful, even though she has never been there. In the former case, we can ascribe the belief to John in a way that cancels its typical implicature. In the latter case, though, it seems we cannot: we seem unable to ascribe a simple aesthetic belief to someone in a way that separates its primary content from its supposedly implicated evidential content. We just cannot felicitously say, “She thinks Hawaii is beautiful, but she’s never been there.” Both difficulties—the one with aesthetic claims, the other with aesthetic beliefs—together present two facets of the same problem: they serve to cast doubt on my claim that primary aesthetic content can be separated from the evidential content at all. This is the problem that I left hanging at the end of §4.1.

In this chapter I hope to put these worries to rest. I will argue that conversational implicatures are not necessarily either contextually or explicitly cancelable, and will discuss some extremely robust non-aesthetic examples that at the least come very close to not being cancelable, and some that may not be cancelable at all. I will then go on to describe how the implicatures of Direct Evidential claims compare to these examples, and explain why aesthetic claims in particular approach uncancelability. Finally, I will apply my discussion of the implicatures of aesthetic claims to the account of aesthetic belief. I will conclude that the very same pressures on simple aesthetic claims that make their evidential implicatures virtually uncancelable *also* act on attributions of aesthetic belief. Someone who’s never been to Hawaii can at most claim that Hawaii must be beautiful. If we insist on asking whether or not she *believes* Hawaii to be beautiful, we can at most say that she believes it must be.

8.1 Extremely robust conversational implicatures

It is important to note that the cancelability of conversational implicatures does not actually follow from their calculability, although the two properties usually go hand in hand. An utterance’s conversational implicature is ‘calculable’ in that it is possible to work it out from the combination of the utterance’s conventional meaning together with the context in which it is

uttered. Usually the speaker has deliberately chosen her utterance to convey the implicated meaning in this way, although it seems reasonable to allow that speakers may unintentionally convey conversational implicatures, just as they may unintentionally convey conventional meanings. (Purists would presumably want to say that this is not conversational implicature, but merely the appearance of it: i.e. that speaker intention is necessary. Nothing I have to say will turn on this.)

First, then, concerning contextual cancelation: the basic dynamics of calculability in no way guarantee that there will be some contexts in which a form of words fails to give rise to its usual implicature. Such a situation would comprise GCI taken to the limit. While we would not expect this situation to be the norm, given both the independence of primary and implicated content, and the multiple complexities of utterance context, there are no general principles that rule it out. Even if we are inclined to judge by fiat that the supposedly implicated secondary content in such a case can at most be conventional and not conversational implicature—given, that is, our inability to separate the secondary content from the form of words in question wherever they appear—it must still be allowed that this secondary content differs markedly from conventional implicature content in being calculable rather than conventional. I will set this possible fiat aside and take it that secondary content that looks like conversational implicature in every way except for the possibility of cancelation is, in fact, conversational implicature.

Sven Lauer (2014) argues that the theoretical possibility here canvassed—i.e. the possibility of conversational implicatures that are not contextually cancelable—is in fact instantiated in a certain very common case. (He thinks the implicatures at issue are not explicitly cancelable either; I will get to explicit cancelation below.) He argues that utterances of unembedded disjunctions give rise to mandatory conversational implicatures to the effect that the speaker either does not know, or does not want to communicate, which of the disjuncts is true.¹²¹ This is of course a kind of ‘Quantity’ implicature that often occurs: if a speaker does not use a salient, stronger alternative (*all* instead of *some*, say), she generally implicates that she is not in a position to do so. The similarity for disjunction is brought out by these examples from Lauer (22):

Ad: Where is John? I need to track him down.
Sp: He is in Europe.
→ *Sp* does not know where in Europe John is.

Ad: Where is John? I need to track him down.
Sp: He is in Paris or London.
→ *Sp* does not know that John is in London.
→ *Sp* does not know that John is in Paris.

As Lauer notes, it is not hard to find contexts in which the implicature in the first of these cases fails to arise, just as in Levinson’s movie theater dialogue:

[Somewhere in San Francisco, CA, *Ad* and *Sp* are planning a dinner party, talking about who they should invite.]

¹²¹ Note, relative to the possible ‘fiat’ I mention above, that the disjunction implicatures Lauer discusses are not remotely plausible as conventional rather than conversational.

Ad: Is John in town?
Sp: No, he is in Europe.

The implicature fails to arise because in this conversation it doesn't matter where exactly John is; what is relevant is just whether or not he can come to the party. However, Lauer goes on to point out that the implicature for the disjunction is unaffected by similar relevance considerations:

[Dinner party context as above]
Ad: Is John in town?
Sp: No, he is in Paris or in London.
→ *Sp* does not know that John is in London.
→ *Sp* does not know that John is in Paris.

The implicature still arises, even though John's exact location is likewise irrelevant to the concerns of the dinner party planning.

Lauer's explanation of this difference for disjunction depends on several special features of the disjunction case: using a disjunction itself makes each disjunct relevant, even if they were not particularly relevant before; using either disjunct would be more informative than using the disjunction; and using the disjunction requires a longer, more complex expression than using either disjunct singly. The implicature for disjunction is thus so robust because the speaker has deliberately chosen a more complex expression to communicate less information, even though the information that is withheld is relevant in the context—it *becomes* relevant through use of the disjunction, even if it was not relevant before.¹²²

Lauer does in fact mention a kind of context in which these disjunction implicatures either fail to arise or can be explicitly canceled, so these implicatures are not quite mandatory, in spite of what appears to be Lauer's official line. He says that "[t]he implicature(s) can be absent if the disjunction was mentioned in the near-by context" and gives the following (explicitly canceled) example:

Ad: If John only were in London or in Paris!
Sp: But John is in London or in Paris. He landed in Heathrow this morning.

Lauer explains this case by saying that, "as with other instances of such 'echoic' utterance choices ... this kind of use depends on (near-)identity with a previously uttered expression." (23, fn. 5)¹²³ No doubt this is because, in trying to construct a context for the example that makes it

¹²² Lauer cites Eckardt (2007) for this explanation: "In using a disjunction, the speaker necessarily has to mention two properties which are usually more specific. These properties are presented as salient and relevant. The simpler sentences are salient alternative utterances in context. The hearer hence will look for a reason why the speaker chose a more complex expression in order to give less information." Lauer says that Eckardt is mainly concerned here with embedded disjunctions, but that, as she herself notes, her remarks apply to unembedded disjunctions as well.

¹²³ Lauer says that this kind of exception "proves the rule" that he has been espousing—but I am not sure what exactly he means by this. He has been arguing that these implicatures are mandatory—neither contextually nor explicitly cancelable—but if we accept that his example is a genuine exception, it just seems to be a straightforward counterexample to his absolute claim. –I always took the phrase "the exception that proves the rule" to be used to indicate a strange outlier that highlights what is *normally* (though clearly not always) the case.

sound natural, *Sp*'s response is heard (at least to my ear) as a somewhat wry metalinguistic comment on *Ad*'s strange disjunctive lament. It is probably not the best example: interpreting *Sp*'s response like this is the only way I can make it even approach naturalness; usually in such a context a speaker will not repeat the disjunction in this way, as Lauer's main argument predicts. (*Sp* may have said, for instance, "But he is. He landed in Heathrow this morning.") We can, however, construct a better example by applying an idea Lauer uses elsewhere. Say an advisor is looking over the rules for graduation, and says to a student, "Well, you *have* taken either 125A or 125B, so you're fine" The advisor in this case does not implicate that she doesn't know which of the classes the student has taken; she is making clear that the rule involves a particular disjunction, and that the student has satisfied the rule. The advisor's remark involves a repetition of the rule—but note that it need not be the exact form of words *per se* that is repeated. In such a case the disjunction itself (however expressed) is what is overwhelmingly relevant and this effectively swamps the introduced relevance of each separate disjunct.¹²⁴

The disjunction implicatures that Lauer describes are nevertheless very, very robust. Moreover, he gives several more examples which could likewise be construed as extremely robust conversational implicatures, if the default assumption of easy cancelability is suspended. One kind of example involves common-ground presuppositions: an utterance of e.g. "John interviewed a father of the victim" (rather than "John interviewed the father of the victim") would, because of the use of the indefinite article, be expected to convey that the victim has more than one father, an implication that is generally contradicted by common knowledge. (Lauer does not mention the complication arising from today's less homogeneous family structures.) Thus "John interviewed a father of the victim," and other similar utterances involving the indefinite article, are generally infelicitous (ignoring the complication, that is).¹²⁵ Similarly, "Mary broke all her arms" (rather than "Mary broke both her arms") is also infelicitous, because it appears to suggest that Mary has more than two arms. (26) Lauer also mentions the standard implications of 'high-negation polar questions': if someone says "Doesn't John drink?", for instance, she implies that she thought John drinks or at least expected that John drinks. Putting a question in this way always carries such an implication, as if a speaker does not want to convey this expectation, she should simply put the question differently ("Does John drink?"). (25)

Lauer takes such examples to show that an utterance can be rendered infelicitous when its conversational implicature is known to be false. He notes that this goes against the standard view, which predicts instead that whenever a putative conversational implicature is known to be

¹²⁴ Lauer suggests at several points that the disjunction implicature is simply immune to considerations of relevance, unlike standard Quantity implicatures. It seems to me, however, that the dynamic of relevance is simply different for disjunction implicatures, rather than entirely absent.

¹²⁵ I modeled my advisor example above on an example Lauer says he was given by an anonymous reviewer, which shows the implication arising from the indefinite article to be contextually cancelable in some cases as well:

[Context: The constitution of Phantasia stipulates: "A son of the king has to be present at the opening of the parliament". It is known to everyone that the current king has exactly one son, who is in attendance. Running through the regulations during the opening ceremony, the speaker of parliament declares:]
As the constitution demands, all MPs have pledged their allegiance to the king, a son of the king is present, ...
(27)

Incidentally, it seems that Phantasia's constitution has not been updated either. The rules of succession to the British throne, for instance, were amended in 2013 to make gender no longer a determining factor.

false, it is always contextually canceled. Lauer's explanation for such cases is that the effect of an utterance will be infelicity, rather than felicitous cancelation of its implicature, when the speaker has no reason to choose that utterance over an utterance that does not carry the false implicature. A speaker has no reason to choose a high-negation polar question if she does not intend to convey her contrary expectations. Similarly, she has no reason to say "Mary broke all her arms" when Mary has the normal number of arms, and "Mary broke both her arms" would do. *Usually* a speaker has no reason to utter a disjunction if she knows which disjunct is true and does not wish to conceal her knowledge; occasionally, however, the disjunction itself is so overwhelmingly salient that she does have a reason, and in such a case the implicature is contextually canceled, as the standard view predicts.

Lauer's explanation of this kind of infelicity is also relevant to the issue of explicit cancelation, although he does not specifically discuss explicit cancelation in any detail. At first glance, it is natural to think that any conversational implicature would have to be explicitly cancelable, as an utterance's primary content is logically independent from its implicature content. Why not simply disavow the implicature content whenever one likes? But an explicit cancelation of an utterance's implicature will only be felicitous when there was a reason for the speaker to choose that utterance in the first place, rather than an alternative that would not have carried the implicature. Otherwise, the explicit cancelation can at best be interpreted as a *correction* of what was said immediately beforehand. This is the case even for standard Quantity GCIs that are trotted out as prime examples of cancelability. To return to Levinson's numeric Horn scale outside the movie theater, we can compare:

- (1) A: It's going to cost us ten dollars to get in and I didn't bring a cent.
B: Don't worry, I've got ten dollars—in fact, I've got twenty, so C can come too.

with:

- (2) A: How much money do you have?
B: I've got ten dollars—in fact, I've got twenty.

The dialogue in (2) doesn't sound strange if we imagine B rooting around in her purse while she speaks, and finding another ten dollars in the middle of her sentence. But in that case "I've got twenty" is really a correction rather than a cancelation. A's question in (2) makes B's total amount of money at issue, so there is no reason for B to say she has ten dollars unless she thinks that's all she has; in particular, if she knows she has more than ten, there is no reason for her to say she has ten and then go on to cancel the implicature to the effect that she has no more. In contrast, A's remark in (1) sets the topic at issue to be whether they have enough money to go to the movie. B thus has a reason to say she has ten dollars even though she knows she has more than that. She effectively changes the topic at issue when she goes on to say that, in fact, she has twenty dollars.¹²⁶

To sum up: the dynamics governing how conversational implicatures arise do not necessitate that all implicatures are cancelable, and there are several real-world examples that are at least extremely robust. (As far as I can tell, if we allow that high-negation polar questions carry conversational implicatures as Lauer suggests, these implicatures are never cancelable.) A kind

¹²⁶ Laia Mayol and Elena Castroviejo (2013) investigate this dynamic empirically.

of utterance that generally carries a certain implicature will only lend itself to felicitous contextual cancelation of that implicature in a given case if there is some conversational reason to choose that utterance instead of an alternative that does not carry the implicature. Similarly, explicit denial of an implicature will be heard as implicature cancelation rather than utterance correction only when there is some point to choosing the original utterance in spite of its false implicature. Typical Quantity GCIs can be canceled in one way or the other when the relevance of the weaker claim swamps the salience of the stronger alternative. This overwhelming relevance gives the conversational point to choosing the weaker claim, in spite of the fact that the speaker is in a position to make the stronger claim as well.

8.2 Aesthetic and other Direct Evidential implicatures

There are three factors that make aesthetic GCIs virtually uncanceled. One of these is the counterpart to the salient oppositions that give rise to many Quantity GCIs (e.g. the opposition between *some* and *all*): that is, having direct evidence for a claim is opposed to having indirect evidence for it. While this opposition is somewhat imbalanced, in that the simple, unmarked form (“X is beautiful”) is associated with the direct evidential implicature, and a marked form (e.g. “X must be beautiful”) is used to indicate indirect evidence, the marked form need not be too long or complex. It is relatively easy to use.

Another factor distinguishes aesthetic GCIs from most other GCIs we have seen. This is that the salient opposition in question—the choice between conveying that one’s claim is made directly and conveying that it is made indirectly—doesn’t affect the claim’s main point content. Earlier, in my discussion of *must* (cf. §3.3), I noted that (e.g.) “X is beautiful” and “X must be beautiful” differ in primary content—the latter but not the former makes a claim to epistemic necessity—but that the main point of either utterance is generally the same: they are both used when the point at issue is simply whether X has a certain aesthetic property. Given that the choice the speaker makes in the aesthetic case doesn’t affect main point content, then, this means that aesthetic GCIs aren’t subject to the relevance considerations that affect other GCIs. “X is beautiful” and “X must be beautiful” are equally relevant in all contexts in which X’s beauty is at issue: the choice between them concerns the subsidiary evidential matter.

Lauer claims that it is this lack of susceptibility to relevance considerations that makes the disjunction implicatures he discusses ‘mandatory’. As I’ve mentioned, I don’t think he’s quite right about that: such implicatures are insulated from many of the relevance considerations that affect other Quantity implicatures, but are not immune to them entirely (cf. fn 124). However, the implicatures (if that’s what they are) of high-negation polar questions do appear to be unaffected by relevance. These questions are chosen over their unmarked alternatives just to indicate the speaker’s expectations; what the speaker is asking to be told is the same in both cases. This is perhaps the better comparison, then—at least in this respect—for aesthetic GCIs.

The overall idea behind the uncanceledness of aesthetic GCIs is therefore this: the opposition that is set up, between conveying that one has direct evidence for an aesthetic claim and conveying that one has indirect evidence for it, requires a speaker to choose a sentence form that is appropriate to the evidential status of her claim. As this evidential status is not related to the main point of the claim, the speaker’s choice is not affected by relevance considerations. Contextual cancelation of the direct evidential implicature of the simple, unmarked form does not arise, and explicit cancelation is not felicitous, because a speaker never has reason to choose the unmarked form when her evidence for her claim is indirect. If she chooses that form when it

is known that her evidence is indirect, her claim is taken to be infelicitous, rather than its implicature heard to be contextually canceled. The claim is likewise infelicitous if she chooses the simple, unmarked form and then tries to cancel its implicature explicitly: she simply has no reason to choose that form in the first place if her own experience is not the source of her claim.

However, I said that there are three factors relevant to the uncanceledness of aesthetic GCIs, and so far I have only mentioned two. The third factor I have in mind is a kind of prologue to the other two. After all, the opposition between conveying that one has direct evidence for one's claim and conveying that one has indirect evidence for it, is different from that between *some* and *all*, or the other oppositions that are typical examples of those that give rise to Quantity GCIs. The opposition between *some* and *all* arises relatively straightforwardly as a result of the words' meaning, which directly affects the main point content of claims that involve them. In contrast, given that most simple, unmarked assertions carry no evidential information at all, the opposition that arises for Direct Evidential claims—between conveying direct evidence and conveying indirect evidence—is a result of the more complex interplay between the primary content of such claims and their epistemology that I described in the last two chapters. It is a dynamic that is specific to them.¹²⁷

The details of this dynamic vary for the different kinds of claims under the Direct Evidential umbrella, and as a result the *strength* of the opposition varies for the different kinds of claims. By the 'strength' of the opposition, I mean the extent to which a simple, unmarked claim excludes the possibility of being felicitous when the speaker does not have direct evidence for it. There is no opposition at all for claims outside the Direct Evidential domain: the simple, unmarked form simply carries no evidential information. But within that domain, the opposition's strength seems to increase in rough proportion to the degree to which a speaker needs direct experience in order to be confident that her claim is true. The more a claim type depends on a speaker's idiosyncratic standards or circumstances, whose relationship with the object in question is hard to predict in advance, the stronger the opposition will be.

For instance, I said earlier that "Margot's house looks orange under the new streetlights" is an example of a Direct Evidential claim, in that it will generally implicate that the speaker has seen Margot's house under the new streetlights herself. I also said that this is not because of any idiosyncrasy in the speaker's visual system, but rather because of the special viewing circumstances that are relevant to the claim. It is plausibly an objective claim, which has the evidential implications it does because it is generally taken for granted that it is the speaker herself who has witnessed the special circumstances at issue. However, this claim's direct evidential implicature is cancelable. There is nothing strange about saying, "Margot's house looks orange under the new streetlights—I haven't been there yet but she was complaining about it the other day" (and of course there could be umpteen similar examples). I take it that this is because the standard relevant to the claim is an objective standard, and that claims like this therefore make it into the Direct Evidential club only by the skin of their teeth.

Another example I discussed earlier presents an interesting contrast. I said that claims involving personality traits like intelligence make it into the Direct Evidentials because such

¹²⁷ Incidentally, the mechanism underlying the *indirect* evidential implications of *must*-ified claims is itself not clear. Von Stechow and Gillies remark that it ought to be due to conversational implicature, given that epistemic necessity modals in every language they've come across carry the indirect evidential signal. However, they can't come up with a Gricean calculation for it. As a result, they tentatively suggest presupposition instead—with the caveat that they would far prefer to be shown that it is a Gricean implicature after all. (2010: §5)

traits, although somewhat objectively measurable, are vague in a way that requires people to make complex judgments about them on the basis of multidimensional criteria—judgments that are clearly influenced by people’s ‘take’ on what counts as (say) intelligent in the context. In general, people who make a claim about intelligence must take a stand on *intelligent*’s cut-off, and it is presumed that they do this as a result of their own judgment. I also mentioned, though, that in some cases there is probably no implication that a speaker who makes such a claim has assessed its subject’s intelligence herself. The example I gave was, “Einstein was intelligent.” Given that Einstein is commonly used as a *benchmark* for supreme intelligence, this is not surprising: Einstein is commonly understood to be way beyond any reasonable cut-off for ‘intelligent’.

The direct evidential implicatures of claims about intelligence, then, are clearly cancelable in some contexts. But when the subject of an intelligence claim is not famous like Einstein, contextual cancellation of the usual implicature generally does not arise, and explicit cancellation usually sounds unnatural. Compare:

[A and B, discussing job candidates:]

A: Margot is pretty smart.

B: Yeah? What’s she done?

A: ??Well, I’m not sure. Margaret had a look at her resume and that’s what she told me, but I haven’t seen it yet.

with:

[C and D, also discussing job candidates—but the job they want to fill is ‘Shampoo Commercial Model’:]

C: Margot has long, dark hair.

D: Yeah? Straight or wavy?

C: Well, I’m not sure. Margaret had a look at her headshot and that’s what she told me, but I haven’t seen it yet.

As I discussed in the last chapter, personal taste claims and aesthetic claims are even more dependent on speaker idiosyncrasy. There is no kind of objective measurability for these, as there is (at least in the way that I roughly outlined) for intelligence. They depend on subjective scales—even when, as for beauty, there is also some claim to universal validity in play. The effect of this seems to be a kind of ossification for them of the evidential opposition between the simple, unmarked form (“That’s beautiful”) and forms that indicate indirect sources. Whereas many Direct Evidential claims seem to have cancelable implicatures, cancelability for personal taste and aesthetic claims seems pretty much absent across the board, given the stronger association of the unmarked form with direct evidential information. The exclusion of indirect evidence as a source for any simple, unmarked claim has become more or less complete.

Even examples from this class that are most like the Einstein example in appearing to invite cancelability do not seem to follow through on that invitation. The *Mona Lisa*’s famous beauty, for instance, does not lead to cancelability for the implicatures arising from claims about it. This is perhaps not surprising: the degree of objectivity that exists for claims about intelligence is not matched by that of the critical standards relevant to beauty, and objectivity is in any case not a

feature of the aspect of beauty claims that requires them to reflect the speaker's own ability to experience an object as beautiful. (Someone may know that an object deserves to be thought beautiful without being disposed to find it beautiful herself.) The fame of the *Mona Lisa's* beauty, then, does not fully impose an external standard on a speaker who makes a claim about it in the way that Einstein's famous intelligence imposes such a standard: in most contexts it is simply not reasonable to refrain from committing to Einstein's intelligence, but the claim that the *Mona Lisa* is beautiful always requires both the faith that it deserves appreciation, and the speaker's own propensity to appreciate it herself. Both requirements, and the latter especially, tie the claim to the speaker's own possibly idiosyncratic standards in a way that the *Mona Lisa's* fame does not weaken; at most, it lends some degree of likelihood both to the *Mona Lisa's* aesthetic worth and to the speaker's responding to the *Mona Lisa* in the same way that others do.

We might wonder, though, about natural beauty. Hawaii's famous beauty makes it sure enough that my propensity to find it beautiful matches most people's; it is just not necessary to see it myself in order to find out about this. We might therefore expect claims about natural beauty to exhibit cancelability behavior that is similar to the Einstein example's. Yet this does not seem to be the case.

I have two things to say about this. The first is that—as I've mentioned in the last two paragraphs, but will bring out more explicitly now—the natures of the standards that are relevant in the two cases are different. In most contexts there is simply no scope for a speaker to reposition *intelligent's* cut-off in a way that puts it beyond Einstein's place on the *intelligent* scale. This is why I called the standard here 'external': the speaker is fully constrained in this matter, in the way that speakers are generally constrained by objective standards. But, while I am sure that my standard of beauty would smile upon Hawaii, it is still *my* standard of beauty that is relevant to my claim about it. This does not lead inexorably to the uncancelability of the implicatures of beauty claims, but it is of a piece with my conclusions in the last chapter, where I said that a beauty claim has a particularly tight relationship with a speaker's experience of a beautiful thing.

The second point I want to make is that many aesthetic claims' direct evidential GCIs do at least come close to being cancelable. Since embarking on this project, I have sometimes been brought up short by my own linguistic behavior. For instance, I once said to my young daughter, who was curious about a friend's vacation photo, "That's in Hawaii. It's a very beautiful place." (I really haven't ever been to Hawaii.) On the one hand, I suppose we might explain my statement by saying that I am entitled to make it on the basis of having seen pictures of Hawaii, or that in any case many of the normal rules of language use are simplified when speaking to young children.¹²⁸ But on the other hand, I can imagine making a similar statement to an adult in a context in which the nature of my own experience is less important than it perhaps usually is—and without having seen pictures. If a friend is trying to choose a vacation destination, say, I may baldly tell her, "The Seychelles are very beautiful" (after she has expressed an interest in the Indian Ocean and fishing), because I am thinking of a different friend who has recently been there and told me about it. I might be challenged—"Oh, you've been there?"—and this would

¹²⁸ If anything, though, this latter explanation would presumably encourage the diagnosis of implicature. The idea would be that sophisticated speakers would pick up on the implicature but that young children would not, and that my understanding of this was what made me think the simple, unmarked form preferable to a more complicated form that would have conveyed the 'correct' implicature. People are surely much less likely to 'simplify' language for young children in a way that perverts its core meaning.

show that my attempted cancelation was not successful; and in fact I may originally have attempted it more through a kind of distraction than through any justifiable linguistic strategy: through being simply reminded of my other friend's claim, rather than speaking wholly according to form.

But it is nevertheless suggestive, I think, that such a context can pull one's speech in this direction. The hypothetical Seychelles claim is similar in this respect to something I did actually catch myself saying once. Someone asked me about a movie; I had just read a review of it but had never seen it, and I said, "It's awful." It is, incidentally, much easier to be sure a movie is awful on the basis of testimony than to be sure it is fantastic: being fantastic requires an immense artistic enterprise to come together successfully, whereas being awful merely requires this enterprise to collapse in one of any number of ways. (It is, if you like, a law of artistic thermodynamics.) I'm pretty sure that this fact affected my willingness to dispense with the proper evidential form.

Once again, of course, I could have been challenged: "When did you see it?" And in such a case I would have had to explain myself: "Oh, I was just reading about it—but it sounds really bad." In this case, as in the Seychelles case, such a challenge shows that I was speaking distractedly, or loosely, and I am forced to correct my claim. But what I think these examples also show is just how fine a line there is between cancelation for these kinds of implicatures, and correction. They show a tendency to want to use the simple, unmarked form even in the absence of personal experience, when one is sure of one's claim and one has a clear reason to make it that is directed away from oneself.

8.3 Aesthetic implicature and aesthetic belief

One comment I have received on drafts of this work, from more than one person, goes something like this: "It makes sense that aesthetic claims in conversation carry evidential information, but what does that have to do with aesthetic belief? Can someone who's never been to Hawaii *really believe* that Hawaii is beautiful, or not?"

I have already discussed Chierchia's suggestion that a sentence's typical scalar implicature is retained, as part of what is believed, when that sentence is embedded in a belief context. I noted, though, that because such implicatures are typically cancelable, we can still attribute this kind of belief to someone, when he does not also believe the implicature, by explicitly canceling the implicature in the belief context: e.g. "John believes that some—indeed, all—of the students are waiting for him." I also noted that such implicatures in belief contexts can be contextually canceled as well (e.g. in B's statement from the movie theater context above: "C thinks she's got ten dollars").

But what about when a sentence's typical GCI *can't* be canceled, or at least approaches uncancelability? Well—do we all believe that either $2+2=5$ or $2+2=4$? What people say about this probably depends on whether they've had any introductory logic. I haven't done any surveys, but my prediction is that the relatively formal-logic-free would balk at the question, in something like the way we balk at questions with presupposition failures ("When did you stop cooking the books?"). There is simply no reason, in most cases, for attributing a disjunctive belief to someone when that person is known to believe exactly one of the disjuncts. If that is the case, she should be said to believe just that. If she is known to believe both disjuncts, and both are relevant to the conversation, then attribution of belief in their conjunction would seem to be

appropriate. Attribution of a disjunctive belief of this kind is generally only appropriate when the believer thinks one of two options is true (or possibly both), but doesn't know which.

That is, because the attribution of disjunctive beliefs follows the pattern for assertion of disjunctive beliefs, we perhaps do not have clear intuitions about our beliefs when they supposedly take such a weird disjunctive form. (It is one of the effects of experience in formal logic that it influences these intuitions in a particular direction.) In my arithmetic example, it is of course obvious that we all believe just one of the disjuncts, so the strangeness of the disjunctive belief attribution echoes the strangeness of Lauer's earlier example of "Mary broke all her arms": in both cases, the typical implicature (the subject doesn't know which disjunct is true; Mary has more than two arms) is clearly false, but rather than being easily contextually canceled, it renders the claim not quite felicitous. And attempting to cancel any disjunctive belief attribution's implicature explicitly would of course be strange too (unlike with "John believes that some—indeed, all—of the students are waiting for him").

The conclusion I am heading for concerning the original question—"Can someone who's never been to Hawaii really believe that Hawaii is beautiful?"—is, I hope, clear. We just don't *put* it that way, given the uncancelable implicature that such an attribution would carry. There is, moreover, no equivalent in the aesthetic case to the simple truth-functional semantics that allows us to say, if pressed for a yes-or-no answer: yes, I do believe that either $2+2=5$ or $2+2=4$. As I have discussed, the uncancelability for aesthetic claims is in any case even more pronounced than for Lauer's disjunction cases. We therefore balk so entirely at attributing aesthetic beliefs to people in the absence of their experience of the object in question that it is natural to think the experience is more intimately connected to the content of that belief than it actually is.

There is one more factor that presumably contributes to this phenomenon. If we are asked whether we believe X or Y, we have a range of possible responses, depending on our epistemic relationship with the disjuncts:

- (1) "X, yeah."
- (2) "Y, yeah."
- (3) "They're both true."
- (4) "Yeah—don't know which, though."
- (5) "Nope."

We know that the disjunction is a complex form, and generally just don't care about having a straight-up answer (yes or no) to the question about the disjunction itself, as long as the answerer makes her relationship with the disjuncts clear. But because (e.g.) "Hawaii is beautiful" has a simple, unmarked form, we assume that it is the most basic case and that it should therefore be a straightforward matter to decide whether people believe it or not: we want to insist on a straightforward, unproblematic answer to the question, "Can someone who's never been to Hawaii really believe that Hawaii is beautiful?" What I have been arguing, however, is that, although "Hawaii is beautiful" is grammatically simple, it is effectively not evidentially neutral, and so there is no evidentially neutral way to make such a claim about Hawaii or to attribute such a belief about it. As a result, there is no straightforward answer to the question about Hawaii: 'yes' and 'no' are more or less equally misleading, and both answers require further explanation. We may of course easily say, though, that someone who's never been to Hawaii can believe that it must be beautiful.

REFERENCES

- Aikhenvald, A. Y. (2006). Evidentiality in grammar. In *Encyclopedia of Languages and Linguistics*, ed. Keith Brown. Oxford: Elsevier, 320-325.
- Allison, Henry (2001). *Kant's Theory of Taste*. Cambridge: Cambridge University Press.
- Allison, Henry (2003). Reply to the comments of Longuenesse and Ginsborg. *Inquiry* 46(2): 182-194.
- Bach, Kent (1999). The myth of conventional implicature. *Linguistics and Philosophy*, 22(4): 327-366.
- Barker, Chris (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1): 1-36.
- Barker, Stephen (2000). Is value content a component of conventional implicature? *Analysis*, 60(3): 268-279.
- Bar-On, Dorit and Chrisman, Matthew (2009). Ethical neo-expressivism. In *Oxford Studies in Metaethics*, Vol. 4, ed. R. Shafer-Landau. Oxford: Oxford University Press, 133-165.
- Boisvert, Daniel R. (2008). Expressive-assertivism. *Pacific Philosophical Quarterly*, 89(2): 169-203.
- Budd, Malcolm (2003). The acquaintance principle. *British Journal of Aesthetics*, 43(4): 386-392.
- Cappelen, Herman and Hawthorne, John (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Chierchia, Gennaro (2004). Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. In *Structures and Beyond*, ed. A. Belletti. Oxford: Oxford University Press, 39-103.
- Copp, David (2001). Realist-expressivism: A neglected option for moral realism. *Social Philosophy and Policy*, 18(2): 1-43.
- Eckardt, Regine (2007). Licensing *or*. In *Presupposition and Implicature in Compositional Semantics*, eds. S. Uli and P. Stateva. London: Palgrave Macmillan.
- Egan, Andy (2010). Disputing about taste. In *Disagreement*, eds. R. Feldman and T. Warfield. Oxford: Oxford University Press, 247-286.
- Fara, Michael (2005). Dispositions and habituals. *Noûs*, 39: 43-82.
- Finlay, Stephen (2005). Value and implicature. *Philosophers' Imprint*, 5(4). Retrieved from <http://hdl.handle.net/2027/spo.3521354.0005.004>.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Ginsborg, Hannah (1998). Kant on the subjectivity of taste. In *Kant's Aesthetics*, ed. H. Parret. Berlin: Walter de Gruyter, 448-465.
- Ginsborg, Hannah (2003). Aesthetic judging and the intentionality of pleasure. *Inquiry*, 46(2): 164-181.
- Ginsborg, Hannah (2014). Aesthetic judging and the intentionality of pleasure. In *The Normativity of Nature: Essays on Kant's Critique of Judgement*. Oxford: Oxford University Press, 94-110.
- Ginsborg, Hannah (forthcoming). Kant's 'Young Poet' and the subjectivity of aesthetic judgment. In *Proceedings of the 12th International Kant Congress: Nature and Freedom*, eds. Violetta L. Waibel and Margit Ruffing. Berlin: Walter de Gruyter.
- Grice, H. P. (1989a). Logic and Conversation. In *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

- Grice, H. P. (1989b). Further Notes on Logic and Conversation. In *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Guyer, Paul (1997). *Kant and the Claims of Taste*, 2nd Edition. Cambridge: Cambridge University Press.
- Haslanger, Sally (2014). The normal, the natural and the good: Generics and ideology. *Politica & Società*, 3: 365-392.
- Hopkins, Robert (2000). Beauty and testimony. In *Philosophy, the Good, the True, & the Beautiful*, ed. A. O'Hear. Cambridge: Cambridge University Press.
- Hopkins, Robert (2011). How to be a pessimist about aesthetic testimony. *Journal of Philosophy*, 108(3): 138-157.
- Hume, David (1965). Of the standard of taste. In *Hume's Ethical Writings*, ed. A. MacIntyre. London: Collier-Macmillan, 275-295.
- Janaway, Christopher (1997). Kant's aesthetics and the 'empty cognitive stock'. *The Philosophical Quarterly*, 47(189): 459-476.
- Kant, Immanuel (2000). *Critique of the Power of Judgment*. Cambridge: Cambridge University Press. Translated by Paul Guyer and Eric Matthews.
- Karttunen, Lauri (1972). Possible and must. *Syntax and Semantics*, 1: 1-20.
- Kennedy, Christopher (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30: 1-45.
- Kratzer, Angelika (1991). Modality. In *Semantics: An International Handbook of Contemporary Research*, eds. A. von Stechow and D. Wunderlich. Berlin: Walter de Gruyter, 639-650.
- Lassiter, Daniel (2014). The weakness of must: In defense of a Mantra. *Semantics and Linguistic Theory* 24: 597-618. DOI: <http://dx.doi.org/10.3765/salt.v24i0.2985>
- Lauer, Sven (2014). Mandatory implicatures in Gricean pragmatics. In *Proceedings of the Formal & Experimental Pragmatics Workshop*, eds. J. Degen, M. Franke and N. Goodman. Tübingen: 21-28.
- Levinson, Stephen C. (2001). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.
- Liao, Shen-Yi; McNally, Louise and Meskin, Aaron (2016). Aesthetic adjectives lack uniform behavior. *Inquiry*, 59(6): 618-631.
- Liao, Shen-Yi and Meskin, Aaron (2017). Aesthetic adjectives: Experimental semantics and context-sensitivity. *Philosophy and Phenomenological Research*, 94(2): 371-398.
- MacFarlane, John (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth, NY: Penguin.
- Manley, David and Wasserman, Ryan (2008). On linking dispositions and conditionals. *Mind*, 117(465): 59-84.
- Martin, M. G. F. (2010). What's in a Look? In *Perceiving the World*, ed. B. Nanay. Oxford: Oxford University Press, 160-224.
- Maslova, Elena (2003). *A Grammar of Kolyma Yukaghir*. Berlin: de Gruyter Mouton.
- Mayol, Laia and Castroviejo, Elena (2013). How to cancel an implicature. *Journal of Pragmatics*, 50: 84-104.
- McDowell, John (1998a). Aesthetic value, objectivity, and the fabric of the world. In *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 112-130.

- McDowell, John (1998b). Values and secondary qualities. In *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 131-150.
- McDowell, John (2003). Subjective, intersubjective, objective. *Philosophy and Phenomenological Research*, 67(3): 675-681.
- McGrath, Sarah (2011). Skepticism about moral expertise as a puzzle for moral realism. *Journal of Philosophy*, 108(3): 111-137.
- McNally, Louise and Stojanovic, Isidora (2017). Aesthetic adjectives. In *The Semantics of Aesthetic Judgements*, ed. J. Young. Oxford: Oxford University Press, 17-37.
- Meskin, Aaron (2004). Aesthetic testimony: What can we learn from others about beauty and art? *Philosophy and Phenomenological Research*, 69(1): 65-91.
- Meskin, Aaron (2007). Solving the problem of aesthetic testimony. In *Knowing Art: Essays in Aesthetics and Epistemology*, eds. Matthew Kieran and Dominic MacIver Lopes. Dordrecht: Springer, 109-124.
- Mothersill, Mary (1984). *Beauty Restored*. Oxford: Oxford University Press.
- Nickel, Bernhard (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford: Oxford University Press.
- Ninan, Dilip (2014). Taste predicates and the acquaintance inference. *Semantics and Linguistic Theory* 24: 290-309. DOI: <http://dx.doi.org/10.3765/salt.v24i0.2413>
- Palmer, F. R. (1990). *Modality and the English Modals*, 2nd edition. Longman.
- Pettit, Philip (1983). The possibility of aesthetic realism. In *Pleasure, Preference and Value: Studies in Philosophical Aesthetics*, ed. Eva Schaper. Cambridge: Cambridge University Press, 17-38.
- Plato (1997a). *Phaedo*. In *Plato: Complete Works*, ed. J. M. Cooper. Indianapolis, IN: Hackett Publishing Company, 49-100. Translated by G. M. A. Grube.
- Plato (1997b). *Republic*. In *Plato: Complete Works*, ed. J. M. Cooper. Indianapolis, IN: Hackett Publishing Company, 971-1223. Translated by G. M. A. Grube, rev. C. D. C. Reeve.
- Potts, Christopher (2007a). The expressive dimension. *Theoretical Linguistics*, 33(2): 165-198.
- Potts, Christopher (2007b). Into the Conventional-Implicature Dimension. *Philosophy Compass* 2(4): 665-679.
- Potts, Christopher (2013). Conversational implicature: interacting with grammar. University of Michigan Workshop in Philosophy and Linguistics. <http://web.stanford.edu/~cgpotts/manuscripts/potts-interacting2013.pdf>.
- Potts, Christopher (2015). Presupposition and implicature. In *The Handbook of Contemporary Semantic Theory*, 2nd edition, ed. S. Lappin and C. Fox. Chichester, UK: John Wiley & Sons. DOI: 10.1002/9781118882139.ch6
- Richard, Mark (2004). Contextualism and relativism. *Philosophical Studies*, 119: 215-242.
- Richard, Mark (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- Ridge, Michael (2006). Ecumenical expressivism: Finessing Frege. *Ethics*, 116(2): 302-336.
- Rind, Miles (2003). Kant's beautiful roses: A response to Cohen's 'Second Problem'. *British Journal of Aesthetics*, 43(1): 65-74.
- Russell, Benjamin (2006). Against grammatical computation of scalar implicatures. *Journal of Semantics*, 23(4): 361-382.
- Sassoon, Galit W. (2013). A Typology of Multidimensional Adjectives. *Journal of Semantics*, 30: 1-46. Draft found at:

<http://www.semanticsarchive.net/Archive/zA0ZTJkY/A%20typology%20of%20multidimensional%20predicates.pdf>

- Schroeder, Mark (2009). Hybrid expressivism: Virtues and vices. *Ethics*, 119(2): 257-309.
- Scruton, Roger (1974). *Art and Imagination: A Study in the Philosophy of Mind*. London: Methuen.
- Sibley, Frank (1959). Aesthetic Concepts. *Philosophical Review* 68(4): 421-450.
- Simons, Mandy (2007). Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117: 1034-1056.
- Travis, Charles (2004). The Silence of the Senses. *Mind*, 113(449): 59-94.
- Vetter, Barbara (2014). Dispositions without conditionals. *Mind*, 123(489): 129-156.
- von Stechow, Kai and Gillies, Anthony S. (2010). *Must ... stay ... strong!* *Natural Language Semantics*, 18(4): 351-383.
- von Stechow, Kai and Iatridou, Sabine (2007). Anatomy of a modal construction. *Linguistic Inquiry*, 38(3): 445-483.
- Wasserman, Ryan (2011). Dispositions and generics. *Philosophical Perspectives*, 25: 425-453.
- Whiting, Daniel (2013). It's not what you said, it's the way you said it: Slurs and conventional implicatures. *Analytic Philosophy*, 54(3): 364-377. DOI: 10.1111/phib.12024
- Wilson, Carl (2014). *Let's Talk About Love: Why Other People Have Such Bad Taste*. New York: Bloomsbury Academic.