

4

Aesthetic Negotiation

Tim Sundell

In the lively and expanding literature on predicates of personal taste, a great variety of theories compete to describe a tiny handful of terms. Philosophers and linguists who work in this area, looking to avoid unnecessary or distracting philosophical baggage, have deliberately eschewed discussions of philosophically stress-inducing terms like ‘beautiful,’ ‘sublime,’ ‘unified,’ ‘sentimental,’ or, for that matter, ‘good.’^{1,2} This is an entirely reasonable choice, and the very liveliness of the literature is proof positive that the meanings of words like ‘tasty’ and ‘fun’ are quite complicated enough, thank you very much.³

Still, it is worth inquiring into the contours of whatever natural kind ‘tasty’ and ‘fun’ are meant to be members of. It is not obvious that such a category will include terms of traditional aesthetic interest like ‘beautiful,’ or terms that pick out aesthetic concepts in Sibley’s sense, such as ‘unified’

¹ Many thanks for helpful correspondence, discussion, and commentary to Chris Barker, David Davies, Andy Egan, Stacie Friend, Manuel García-Carpintero, Shen-yi Liao, Dan López de Sa, Teresa Marques, Aaron Meskin, Dominic Lopes, Louise McNally, David Plunkett, Isidora Stojonovic, Nick Zangwill, and to audiences at the International Summer School in Philosophy workshop on “Disagreement” in Miglieglia, Switzerland, the LanCog workshop on “Values in Context” at the Universidade de Lisboa, the London Aesthetics Forum Workshop on “Aesthetic Adjectives” at University College London, and especially to the participants and audiences at the two workshops on “The Semantics of Aesthetic Judgment” at the University of Victoria and to the organizer of those workshops, James Young.

² See, for example, Lasersohn (2005, 645); Stephenson (2007, 490), Egan (2010, 253), and Sundell (2011, 268).

³ In this chapter, I’ll use italics for emphasis and to introduce technical terminology, single quotes to mention linguistic expressions, and double quotes for quoting other authors, scare quotes, and simultaneous use/mention.

or 'sentimental.' On first pass, words like 'tasty' and 'fun' seem, far more than 'beautiful' or 'unified,' to deserve their label "predicates of *personal taste*." When we describe the chili as tasty or the roller coaster as fun, there is a natural sense in which our claim is an expression of simple personal preference. Contrast that with the critic's assertion that *Death of a Salesman* is sentimental or that Van Hoyt's *Boy with Apple* is beautiful. *Prima facie*, it is much less plausible that in these cases what we're really talking about is ourselves. Claims of this latter sort have what Railton (1998) calls "objective purport," and whether in the end they differ categorically from claims about personal taste, it would be good to have an explanation for this seeming difference.⁴

Note that this distinction differs from, and in fact seems to cross-cut, the distinction between *substantive* and *verdictive* aesthetic judgments (Zangwill 2001). (To keep things focused on linguistic expressions, we can focus on the distinction between substantive and verdictive aesthetic terms.) 'Beautiful,' at least in many cases, is used to express a judgment of overall aesthetic merit. It is—in those cases—a verdictive term. By contrast, 'sentimental'—though it may carry evaluative content—picks out a particular aesthetic quality. It is therefore a substantive term. Yet judgments about what's "beautiful" or what's "sentimental" have in common whatever it is that gives something objective purport. Contrast that pair with another pair of terms, 'tasty' and 'sexy.'⁵ Those terms differ with respect to one of the distinctions and are alike with respect to the other. 'Tasty'—unlike 'rich,' or 'balanced,' or 'delicate'—does not pick out any specific gustatory quality. Like 'beautiful,' it is used to express judgments of overall success in the relevant domain. 'Sexy,' like 'sentimental' and unlike 'tasty,' does pick out a particular quality. Yet claims about what's "tasty" or what's "sexy" are alike in that they express

⁴ I'll use *predicates of personal taste* to refer specifically to words like 'tasty' and 'fun'. For lack of better options, I'll use *properly aesthetic predicates* to refer to words like 'beautiful' and 'unified', and *predicates of taste* (without the "personal") or sometimes just *aesthetic predicates* to refer to the whole class, without assuming that any of those labels is particularly apt. Since my topic here is whether there is a semantic distinction between words like 'tasty' and words like 'beautiful,' I explicitly disavow the pre-theoretical assumption that "predicates of taste," as I use it, picks out a natural kind.

⁵ Thanks to James Young, who in his briefing for the workshop giving rise to this volume, observes that judgments about a movie star's "sexiness" (or about vegemite's "tastiness") differ in important ways from properly aesthetic judgments.

judgments that seem closely tied to the preferences or standards of individual speakers or critics.

So what distinguishes ‘beautiful’ and ‘unified’ on the one hand, from ‘tasty’ and ‘sexy’ on the other? The answer has something to do with the normative demands that an assertion involving those terms places on those around us. When we say that something is tasty or sexy, we don’t typically place a great deal of pressure on our listener to agree with our claim, whatever such agreement might amount to. By contrast, if we describe something as unified, or beautiful, we *do* typically place pressure on our interlocutor to agree or to defend her own judgments. With such an assertion, we “lay claim to correctness” (Zangwill 2005). This distinction has been characterized in a variety of ways over the years. But to avoid premature theoretical commitments in a discussion specifically about recent debates in semantics, I’ll introduce a term of art here. Let’s call this a difference in *aesthetic pressure*.⁶

So as we look to expand our focus from predicates of personal taste to aesthetic predicates more generally, we find on the one hand the distinction between substantive and verdictive, and on the other the distinction between high pressure and low pressure. I’ve claimed that the two distinctions are orthogonal. But my argument here does not depend on this or any other claim about the relation between the distinctions. We should expect *all kinds* of heterogeneity within a category broad enough to contain words as varied as ‘tasty,’ ‘unified,’ and ‘beautiful.’ My goal here is to argue that one aspect of that heterogeneity is less profound than it seems. I argue that, while the difference between high pressure and low pressure captures something real and important about our use of these classes of terms, no categorical *semantic* distinction is necessary to explain that difference.

To begin, consider a pair of proposals that is constructed to capture exactly this difference. As noted, high-pressure terms aim, in some sense or other, at a kind of objectivity; they make claims that go well beyond the tastes of the speaker. By contrast, low-pressure aesthetic terms are quite naturally thought of as, perhaps primarily, concerned with facts

⁶ I’ll sometimes refer to expressions as high or low pressure, sometimes to *uses* of expressions as high or low pressure, and sometimes to assertions as high or low pressure. The question of what exactly is the bearer of these properties is part of what’s at issue here, so to pin things down more precisely at this stage would be counterproductive.

about the person doing the evaluating.⁷ For this reason, the preponderance of theories in the literature on predicates of personal taste lend themselves perfectly to the low-pressure terms. To simplify dramatically, but in ways that shouldn't matter here: The contextualist thinks that an utterance of 'the chili is tasty' expresses the proposition that the chili satisfies the standard of the person who makes the assertion. The assessment relativist thinks that such an utterance expresses a proposition that is true just in case the chili satisfies the standard of the person evaluating that utterance (MacFarlane 2014). The centered world's relativist thinks that it consists in a bid to add to the common ground the property of being disposed to like the chili (Egan 2010). And the expressivist thinks that it does not express a proposition at all, but rather directly expresses the speaker's liking for the chili. Every one of these theories builds reference to some experiencer or standard of taste directly into its account of the semantics of 'tasty.'⁸

Let's consider the contextualist in particular, though analogous points can be made about the other views. For the contextualist, assertions of 'the chili is tasty' express a proposition that is true just in case the contextually relevant experiencer—most often, though not necessarily, the speaker or some group containing the speaker—likes the chili. This goes a long way towards explaining why this type of assertion is low pressure. While of course people can disagree over tastiness claims—and the contextualist has a range of mechanisms for explaining this when it happens⁹—we all know that standards differ. If the truth of claims of this type is bound up with facts about some standard, we will always have

⁷ There may not be a deep metaphysical distinction here, or if there is, it may take more work to properly identify it. An object that I like is, after all, an object that has the property of being liked by me. But the data to be explained here do not come prepackaged with a precise characterization of the difference either: what's to be explained is that certain classes of otherwise-similar expressions seem to differ in some way or other having to do with objectivity or normative demand. To build more theoretical machinery pretheoretically into the *explanandum* is thus not only unnecessary but methodologically problematic.

⁸ MacFarlane (2014, Ch 1), in advocating a form of relativism, claims that the connection between using the word *tasty* and knowing that you find the object's taste pleasing is *the* crucial feature about our use of the term.

⁹ See Marques (2014), García-Carpintero (2008), López de Sa (2008), Huvenes (2014), Stojanovic (2007), Björnsson and Finley (2010), and Sundell (2011), among others. It remains controversial whether any of these proposals in fact do the job better than the alternatives to contextualism, but despite some claims to the contrary, the contextualist has an embarrassment of resources to draw in explaining this type of disagreement.

the possibility of difference in standards immediately at hand to resolve the dispute.¹⁰

Assertions involving high-pressure terms seem, by contrast, to involve a claim that is in some sense more objective. We might naturally aim to capture this feature of those assertions by supposing that for high-pressure terms, in contrast to low-pressure terms, we should adopt some type of objectivist (MacFarlane 2014, p. 2) or invariantist semantic theory. Setting aside the context sensitivity present in any gradable adjective—*how much* beauty or unity it takes to qualify—‘beautiful’ (or ‘unified’), would on this view denote the same property context to context, no matter what the tastes of the speakers. Having no place in the semantics for experiencers or standards of taste, this type of term is insensitive to the differences that can make superficially contradictory claims about “tastiness” consistent. This goes a long way towards explaining why this type of assertion is high pressure. Differences in individual standards of taste are irrelevant to the truth of propositions expressed with this type of term. When two people each use the expression ‘beautiful,’ they’re talking about the same thing, no matter the difference in their tastes.

Explaining the difference in this way provides a ready-made explanation for the difference in normative demand between assertions about personal taste and assertions expressing properly aesthetic judgments. Many of our disagreements about taste seem to be, in some sense, *faultless*. But many of the disagreements we see in critical discourse about art—disagreements about particular aesthetic qualities or overall aesthetic success—seem to run deeper, to be more comparable to ordinary “genuine” disagreement, and thus less happily conceived of as faultless. On the view here, this is easy to explain. Two speakers can utter in turn ‘The chili is tasty’ and its negation while both speaking truly. That’s because ‘tasty’ involves a semantically inbuilt call for standards and those standards may differ speaker to speaker. But if one speaker utters ‘The painting is beautiful’ while the other utters its negation, they inevitably contradict one another. *Faultlessness* in our disputes about taste is a

¹⁰ I follow Plunkett and Sundell (2013a) and (2013b) in using *disagreement* to pick out a rational conflict in attitudes and *dispute* to pick out a linguistic exchange that purports to express such a conflict.

matter of consistent semantic contents, while the *genuine disagreement* expressed in our properly aesthetic disputes is a matter of inconsistent semantic contents.

The view we're considering is one on which a simple form of contextualism is true of predicates of personal taste, while some kind of invariance is true of properly aesthetic predicates. Of course, I've described both components of that view in only the sketchiest of terms. Nevertheless, I submit that there is reason for pause even at this early stage. The difference between high-pressure terms and low-pressure terms is real. But to encode it in this way—as a difference in the semantic structure of the relevant expressions—is to make that difference categorical, a binary distinction between two entirely distinct categories of expression. And that kind of explanation just can't be right. The worry is not a matter of faultless disagreement. It's a matter of *faultiness* and *disagreementlessness*. Expressions of mere difference in taste have the potential to evince or create deep and persistent conflict. And disagreements about aesthetic qualities or aesthetic merit have the potential to disappear with reference to differing standards.

To see this, consider first the dialogue in (1).

- (1) Alphie: Vegemite is tasty.
 Betty: No, you're wrong. Go try it again. Have you noticed how salty it is?
 Alphie: I have. I still think it's tasty.
 Betty: You're nuts. You only like it because of your fond memories of your trip to Australia.
 Alphie: That's not true. Have you noticed how the maltiness of the vegemite complements the flavor of the bread?
 Betty: Yes, but I don't want something that makes the bread even maltier. I want something that adds some richness or sweetness, like good old butter or jam. Can't we agree about anything?

Alphie and Betty are aware, as we all are, that tastes differ. Nevertheless, we can imagine Alphie and Betty continuing this debate for some time, drawing on all kinds of specific features of the vegemite in making their arguments, and their debate, in the end, becoming quite heated. Some—not all! but some—of our arguments about tastiness are serious, and heated, and persistent.

Note that this feature of these debates seems unrelated to the consistency or inconsistency of the semantic content of the speakers' utterances, irrespective of our preferred theory of the semantics of 'tasty':

- (2) Alphie: Vegemite tastes good to me.
 Betty: Well, it doesn't taste good to me. Go try it again. Have you noticed how salty it is?

[...]

Once Alphie and Betty make clear that they are talking about "tastiness to me," there is no question that the propositions they assert are consistent. But there is clearly *still* disagreement of some kind here, and in fact the dialogue in (2) does not feel all that different from the dialogue in (1). There are differences between (1) and (2). Betty can't respond in (2) with the speech act of denial. She can't respond "No, you're wrong, it doesn't taste good to me" or "That's false," etc. But in this context that's beside the point. The difference in available linguistic devices does not mean that Alphie and Betty don't disagree here. Clearly they do, and in a way that does not feel radically different from the dialogue in (1).¹¹ But in (2), their disagreement is expressed quite happily with utterances that uncontroversially express consistent propositions.

Now consider a couple of disputes involving high-pressure terms.

- (3) Alphie: *Titanic* is a terrible movie.
 Betty: I can see why you feel that way, but to me it's beautiful.
 Alphie: Fair enough.
- (4) Alphie: That Ornette Coleman solo is mercurial and dynamic.
 Betty: I can see why you feel that way, but to me it is meandering and chaotic.¹²
 Alphie: Fair enough.

The dialogues in (3) and (4) are *disappointing*. We may feel that the speakers could have learned more about the works and had a vastly more rewarding exchange if they'd been willing to dig in their heels even a little bit. What the dialogues in (3) and (4) are not, however, is

¹¹ See Stevenson (1937) for an early discussion of the notion of disagreement in attitudes, and Huvenes (2014) for a more recent discussion of that notion as it applies specifically to these issues.

¹² I borrow the dynamic/chaotic contrast from Sibley (1959, 428).

bizarre. There is nothing confused, or incompetent, or semantically infelicitous about them.

What we see in (3) and (4) is that even with high-pressure terms—verdictive or substantive—we can still defuse the argument with reference to personal preference or personal standards. Contrast that with the dispute in (5).

- (5) Alpie: There are three people in my office.
 Betty: I can see why you feel that way, but to me there are only two.
 Alpie: Fair enough.

However objective we may aim to be in our aesthetic verdicts or even in our substantive claims about aesthetic qualities, we nevertheless have available a strategy for defusing the argument that simply is unavailable in other kinds of disagreement.¹³ We may lay claims to correctness, but those claims vary in strength.

On the view under consideration, we explain the difference between high-pressure and low-pressure predicates of taste by implementing in the semantics the notion that low-pressure terms call for some standard or experiencer, while high-pressure terms are, in that respect, context-invariant. On such a view, disagreements involving low-pressure terms can be *faultless* in virtue of the fact that—when speakers have different standards in mind—they express propositions that are logically *consistent*. Disagreements involving high-pressure terms can be *genuine* in virtue of the fact that—no matter what standards the speakers may have in mind—they express propositions that are logically *inconsistent*. But the view doesn't work. Why? The problem is this: aesthetic pressure is dynamic and a matter of degree. But logical consistency is static and not a matter of degree. You can't rely on a feature of a word that doesn't change to explain a feature of its use that does.

¹³ To emphasize, whether such moves are ultimately legitimate is a separate question. Perhaps a strong form of realism about aesthetic value or properties will turn out to be right and that we therefore should take an agreement to disagree about sentimentality to be confused in a way not all that different from agreement to disagree about the number of people in the office. Even so, the manoeuvre has a *prima facie* legitimacy that is lacking in the office case. That stands in need of explanation and is enough to cause problems for this account of the difference in normative demand.

In our discourse about personal taste, we sometimes do not place any pressure on those around us to agree with the preferences we express. But sometimes we do. Sometimes it varies over the course of a conversation. A conversation that begins low pressure—simply a sharing of facts about personal preference—can turn high pressure, with participants digging in their heels and demanding that others share their preferences or defend their own. Conversely, in our discourse about aesthetic qualities or aesthetic merit, we often make claims that place a high degree of pressure on our interlocutor to agree or to defend her own position. But sometimes we do not. Sometimes it varies over the course of a conversation; even in a debate about a quality like unity or sentimentality, we can, under the right circumstances, revert to differing standards to defuse the argument.

Considerations like these suggest that whatever the difference is between *typically* low-pressure and *typically* high-pressure terms, it should not be directly encoded as a difference in their meanings. More generally, considerations like these suggest that disagreement itself—at least in our disputes about taste, aesthetic qualities, and aesthetic merit—should not be measured by the consistency or inconsistency of the propositions expressed in the dispute. Aesthetic pressure comes and goes, it increases or decreases in strength. It gets entrenched in debates about personal preference and it vanishes in debates about the most objective seeming aesthetic qualities. Ordinary disagreement over the truth or falsity of propositions is simply the wrong model for a practice like this.

If ordinary disagreement is the wrong model for our practice of aesthetic debate, then what is the right model? I submit that aesthetic debate is better conceived of on the model of a *negotiation of standards* rather than as a disagreement as that notion is commonly conceived.¹⁴ Hints of the motivation for this idea can be found in popular culture. Consider first a famous quote from the film *High Fidelity*, itself a work deeply concerned with aesthetic debate. Rob Gordon is out on a date and remarks in voice-over, “I agreed [with her] that what really matters is what you like, not what you *are* like. Books, records, films—these things

¹⁴ Plunkett and Sundell (2013a) offer a notion of “disagreement” that is consistent with the thought that many disagreements are expressed in ways that don’t fit the traditional model.

matter. Call me shallow but it's the fuckin' truth, and by this measure I was having one of the best dates of my life." What Gordon's observation highlights is that much of the time we *want* to be similar in taste to those around us. We care deeply about how similar we are to those we care about, where we differ when we do, and how—where possible—to synchronize our tastes.

Of course, things can also run in the other direction. The internet meme Hipster Kitty stops liking things once they're popular and remarks, "Arcade Fire won a Grammy? Dead to me" (<http://memegenerator.net/instance/10290147>).

Alexander Nehemas (2002, 211) expresses a distinct but related sentiment:

Although I say this with serious discomfort, a world in which everyone liked Shakespeare, or Titian, or Bach for the same reasons—if such a world were possible—appears to me no better than a world where everyone tuned in to Baywatch or listened to the worst pop music at the same time. What to me is truly frightful is not the quality of what everyone agrees on, but the very fact of universal agreement. Even the idea of two individuals whose aesthetic judgments are absolutely identical sends shivers down my spine.¹⁵

The point of both of these quotes is that we care (among other things, of course) a great deal not just about our own aesthetic judgments, but about the degree to which those judgments overlap or conflict with the judgments of others. Sometimes it's important to us that our tastes *differ* from those around us, or from a particular segment of those around us, or even that some difference in taste exists in general. But this observation is congenial to the more general point. We care about how similar or different we are in our tastes and aesthetic standards. Much of the time we want to reduce differences, and occasionally we want to augment them. So we need a strategy for sussing those differences out and, where appropriate, attempting to align our standards.

It's crucial at this point to observe an easy to miss fact about this kind of view. The idea that aesthetic debate consists in a kind of negotiation of standards is *metaphysically neutral* in the following sense: To describe an exchange as a negotiation of standards is not yet to say anything at all about the rule by which the results of that negotiation can be measured.

¹⁵ Thanks to Aaron Meskin for emphasizing this point, and for directing me to the Nehemas passage.

Suppose some extreme form of meta-aesthetic anti-realism turns out to be true. That's entirely consistent with the negotiation view for obvious reasons. But suppose things turn out the opposite way, and some strong form of aesthetic realism turns out to be true. Realism about aesthetic value is a metaphysical thesis, not a linguistic one. And it is entirely consistent with the general claim that the aim of aesthetic debate is the alignment of standards, and with the particular implementation suggested below.

If realism is true, then there will be metaphysically robust facts of the matter about whether speakers have aligned their standards *in the right way*. Speakers who sync up on a standard of taste that corresponds to the objective aesthetic joints in the universe will have done it right. Speakers who do not, will have made a mistake. Their mistake will not consist, on the view I advocate here, in the literal expression of a false proposition, but that's neither here nor there. If you're *already* a metaphysical realist, then the goal of evaluating objects by a standard that lines up with the aesthetic joints is every bit as important as the expression of true propositions. Perhaps, when we seek to align our standards, we are doing so purely for its own sake, or as an expression of personal preference, or social power, or what have you. Or perhaps we seek to align our standards in a way that corresponds to what the world really is like aesthetically. Most likely, our aesthetic debates range widely among these and other aims. The claim that in aesthetic debate we seek, in the first instance, to align our standards is neutral among all of these options.

The idea that aesthetic debate consists largely in an attempt by speakers to coordinate or align their aesthetic standards has been raised elsewhere in the literature on predicates of taste. Stephenson (2007) suggests that “[participants in a conversation about matters of taste] are trying to align their world views, not only with regard to factual beliefs such as whether Bill works on Fridays, but also with regard to subjective matters such as what is tasty” (p. 510). Egan (2010) puts it like this:

I propose that we should think of this effect of successful aesthetic assertions, and successful resolutions of aesthetic disputes, of inducing mutual self-attribution of certain dispositions to have a particular sort of response to a particular (kind of) object, as the central business of assertions and disputes about taste, and not as a mere side effect.¹⁶ (p. 260)

¹⁶ There is a subtle difference Egan's claim and the one I advocate here. It has to do with what's meant by the phrase “central business”. My claim is that aesthetic debates are *in the*

Stephenson and Egan advocate distinct versions of a relativist semantics for predicates of taste. Those accounts do an excellent job of capturing the sense in which disputes involving ‘tasty’ or ‘fun’ play a role in coordinating our standards for tastiness and fun. Indeed, I don’t have a negative argument to present against those views here. Nevertheless, I submit that there is another option, one that maintains Stephenson’s and Egan’s insight about the importance of coordination, that calls for a simpler and more traditional semantic apparatus, and that more easily scales from disputes about taste or subjective matters to the most objective seeming claims about aesthetic qualities or merit.

When two speakers work to sync up on their standards of taste, their activity is less like a disagreement about some factual proposition and more like a negotiation. When we engage in an activity like this, we push and pull, we demand concessions, and we make compromises. Keeping in mind those cases of “faultless disagreement,” a negotiation is also the kind of activity where we can find ourselves deeply at odds without necessarily taking one another to be factually mistaken. Yet, as noted directly above, negotiation of standards doesn’t happen in a vacuum. Perhaps we want to align our standards because we have to make joint decisions about what to order at a restaurant, or because we have to choose which paintings to include in the show.¹⁷ But it need not be so concrete. We may want to align our standards because we hope to convey cultural cachet on objects that deserve it. Or to reward with our esteem objects that create in humans a distinctive kind of response. Or we may, like Rob Gordon, simply want to make it the case that we have a maximal amount of evaluative overlap with a person we care about or want to know better. What is the mechanism for this kind of

first instance attempts at the coordination of standards. As I emphasize throughout, that leaves open the question of why we attempt to coordinate standards, or how the standards on which we align are to be evaluated. Those questions might receive answers corresponding closely to more traditional views of the aims of aesthetic debate, for example, that the standards we align on should be evaluated by how well they reflect the objective aesthetic facts. Thus my claim is meant to be *consistent* with a range of existing views about the aims of aesthetic debate. Egan, by contrast, could be read as making a slightly different claim, namely that the central business of aesthetic debate is to coordinate standards *as opposed* to getting at the aesthetic facts or some other alternative. Thanks to an anonymous referee for pressing this point.

¹⁷ See Marques (2014) and García-Carpintero and Marques (2014) for discussion of the connection between disagreements about taste and solving coordinated action problems.

activity—not an argument about the truth of a proposition, but the negotiation of a shared standard?

Recent work on the notion of *metalinguistic* uses of linguistic expressions provides, I submit, the perfect candidate for this kind of activity. In his (2002), Chris Barker introduces the notion of a *metalinguistic* or *sharpening* use of a gradable adjective like ‘tall.’¹⁸ That notion is taken up in García-Carpintero (2014), Sundell (2011), Sundell and Plunkett (2013a), and in Barker’s own (2013) as the kind of usage that can play a role in disputes that center not on the truth of propositions that are the literal content of speakers’ utterances, but on choices for how the expressions employed in that very dispute should be used.

To get a sense of the relevant notions, consider the following pair of scenarios. Both scenarios, along with the substance of my discussion of them, are taken from Sundell and Plunkett (2013b). First, suppose that Alphie has just arrived at the Antarctic research station for the first time. He stares at the known-to-be-accurate thermometer, shivering, and asks his new co-worker the question in (6).

Antarctic Research Station

(6) Alphie: Is this cold?

Betty: Nope, I’m afraid this isn’t cold.

When Betty says to Alphie “I’m afraid this isn’t cold,” she has not provided Alphie with new information about the temperature. Alphie already knows what the temperature is. Rather, Betty has provided Alphie with information that is, in the first instance, about language. She has informed him that in his new context the threshold for the gradable adjective ‘cold’ is lower than the current outside temperature. In other words, the temperature would have to be even lower than it already is to be deserving of the label ‘cold,’ as that label is employed around here. Betty has used (not mentioned) the word ‘cold’ in such a way as to communicate information to Alphie about how that very expression is used in the present context. This is what Barker (2002) calls a *metalinguistic usage* of an expression.

Why should Alphie care about this information if it is just about language? He’s concerned about the temperature, after all, not about mere definitions or contextual parameter settings. The answer should be

¹⁸ Similar observations go back much further. See Hare (1952), pp. 112–13.

obvious. Although the information Betty conveys is, in the first instance, about language, it has relevance far beyond that. By learning how ‘cold’ is used around here, Alphie can infer non-linguistic information about precisely the issue he’s most concerned with—facts about the range of temperatures he can expect to be typical in his new work environment.

Now consider a second type of case. This time, let’s imagine that Alphie and Betty no longer work in Antarctica, but rather in an office in Chicago. Alphie often feels chilly while Betty most of the time does not. The two frustrated coworkers have the dispute in (7) while looking together at their shared thermostat.

Office Thermostat

- (7) Alphie: It’s cold in here.
 Betty: No, it’s not cold in here.

Like Antarctic Research Station, Office Thermostat involves the meta-linguistic usage of the expression ‘cold.’ Alphie and Betty’s dispute does not express a disagreement about what the temperature is. Again, they both know what the temperature is. But Office Thermostat differs from Antarctic Research Station in an important way. In Antarctic Research Station, there is a settled fact of the matter about how ‘cold’ is used locally. The information Betty conveys in Antarctic Research Station is thus essentially descriptive: “Here’s how people use ‘cold’ around here.” Office Thermostat is not like that. In Office Thermostat, Alphie and Betty are not arguing about some settled fact of the matter regarding people’s use of the word ‘cold’ around the office. Rather, they are trying themselves to *settle the question*.

If Betty succeeds in pushing Alphie to adopt a usage of ‘cold’ according to which the office doesn’t qualify, she will have *made it the case* that the threshold for ‘cold’ is lower than the current office temperature. Why would she bother to do so? Why argue about how to use a word? Because how we use words matters. A word like ‘cold’ plays a certain functional role in our practice of decision-making and coordinated action regarding thermostats. In a way that has nothing to do with what’s *analytic* about ‘cold,’ users of that word systematically agree to turn up the heat if everyone involved agrees that the word applies to the current temperature. Arguing about how to use the word ‘cold’ is one way to argue about whether to turn up the heat.

To emphasize, it's not just that in Office Thermostat Alphie and Betty are not having a factual disagreement about the temperature. They're not even having a factual disagreement about language. Rather, they are *negotiating* the matter of how one bit of context-sensitive terminology will be made precise for purposes of their conversation.¹⁹ They are doing so because of the functional role that 'cold' plays in our climate control practices, a role that it plays irrespective of which particular temperature property it picks out once made precise. Settling this linguistic question thus has consequences for the speakers' beliefs and actions regarding non-linguistic matters. Metalinguistic disputes of this particular kind—where speakers do not disagree about how some expression is in fact used, but rather negotiate how that expression will be used—are what Sundell and Plunkett (2013a) call *metalinguistic negotiations*.

Two crucial features of Office Thermostat are worth dwelling on here. First, the functional role that 'cold' plays in our discourse about climate control tightly constrains the range of values that will be advocated for in typical situations. In a metalinguistic negotiation about 'cold,' in the context of a dispute about whether to turn up the heat, you will almost never find a speaker advocating for a threshold of higher than, say, 27° C. Such a speaker would be advocating for a use of 'cold' that would play a defective role in the relevant practice. It would suggest that even if the office is already at 27°, we should nevertheless turn up the heat even higher, and in practice such a suggestion is extremely rare. But—and this is the first crucial feature—nothing *semantic* about the word 'cold' dictates this feature of our usage. If we were to limit our attention to disputes about climate control, we might even come to believe that 'cold' simply can't apply to temperatures of more than 27°, as a matter of meaning. But that would be a mistake. The fact that 'cold' just so happens to be used in a wider range of scenarios—the temperature in the sauna, the temperature in the oven, the temperature in the lava—demonstrates that 'cold' can have a threshold of any temperature

¹⁹ There are cases where metalinguistic usages *are* employed to have factual disputes (rather than negotiations) about language. We could imagine, for example, a continuation of Antarctic Research Station, where Charlie, another station veteran, joins the conversation and objects to Betty's claim. "No, this *is* cold. Even around here, the temperature is normally a bit higher and you know it. Stop trying to intimidate the newbie!" Plunkett and Sundell (2013a) call both types of disputes *metalinguistic disputes*. Those particular metalinguistic disputes in which meaning is being negotiated, rather than debated as an antecedently settled matter, are the ones we call *metalinguistic negotiations*.

at all, if you cook up the right context. Even if describing such a context becomes extraordinarily forced or difficult—a context, for example, where the center of the sun qualifies as ‘cold’—that’s a fact about the ways in which we put the word to use. Not about its semantics.

The second crucial feature of Office Thermostat is this. Given (a) our actual goals in the conversation (reaching a decision about the thermostat, being comfortable at work), and (b) facts about us (the range of temperatures at which we can in fact work comfortably), standards themselves are evaluable as *better* or *worse*. The fact that Alphie and Betty’s negotiation is, in the first instance, about how to use a word in no way commits us to the view that any outcome is just as good as any other. Their disagreement is about language, but it is not “just semantics.” In the context of an argument about climate control in an office setting, a threshold of 27° is a *bad standard*. It’s not a bad standard because uses of ‘cold’ with that threshold setting will express false propositions. They might well express propositions that are true. It’s a bad standard because the speakers are trying to fix on a standard that corresponds to their ability to work comfortably in the office, and a threshold of 27° doesn’t do that.

So let’s return now to predicates of taste, and in particular let’s consider a simple form of contextualism about those expressions. On such a view—here modeled on Lasersohn (2005)’s “Option 2”—sentences containing predicates of taste make “indexical reference to some relevant individual or group, not necessarily the speaker”. For example, an utterance of an expression of the form ‘ x is tasty’ is true just in case x is tasty relative to the standards of some contextually determined experiencer. That experiencer will often be, but is not necessarily, the speaker or some group containing the speaker. Note that on this kind of view, the contextually determined standard can vary for at least two reasons: Different *individuals* can have different standards. But a single individual is also a member of *groups* that vary in size and make up. Some experiencer groups containing the speaker will be sensitive mainly to that individual speaker’s likes and dislikes. Other larger groups or even idealized experiencers will be sensitive to a much wider range of tastes and considerations.²⁰

²⁰ I don’t myself endorse this version of contextualism. I advocate a specific alternative in Sundell (forthcoming). My goal here is not to show that this version is correct, but rather to show that even this simple, widely dismissed form of contextualism can offer an account

With this type of theory in mind, consider now a couple of disputes involving predicates of taste.

Buying Cookies

(8) Alvie: Oreos are tasty.

Betty: No, Oreos aren't tasty. They're too sweet.

Talking Art

(9) Alvy: *The Seventh Seal* is a beautiful movie.

Mary: No, it's not. It's adolescent, fashionable pessimism.

Let's imagine that Alvie and Betty are standing around the grocery store deciding which cookies to bring home for movie night. Meanwhile, Alvy and Mary are walking home from dinner, debating which artists belong in the Academy of the Over-Rated. The contextualist about predicates of taste, just like the contextualist about 'cold,' can now analyze these disputes as metalinguistic negotiations. What Alvie and Betty are really up to is an attempt to synchronize their standard for what shall count as a "tasty" cookie, for purposes of the conversation. What Alvy and Mary are really up to is an attempt to synchronize their standard for what shall count as a "beautiful" movie, for purposes of the conversation.

Crucially, the very same observations that are made above about Office Thermostat can be made here about Buying Cookies and Talking Art. First, the functional role played by the expression tightly constrains the range of values that will be advocated for in typical circumstances. How does this observation apply here? On the version of contextualism we're considering, the relevant standard concerns the *experiencer*—which individual or group's tastes or standards matter for the claim. As is exemplified in Buying Cookies, the word 'tasty' typically plays a role in fairly low-stakes decisions, in domains where speakers assume a fairly high degree of interpersonal variation in standards. As a result, the word

of the difference between low- and high-pressure predicates of taste, once augmented with a story about metalinguistic negotiation. Almost any view about the semantics of predicates of taste is consistent with the independently motivated notion of metalinguistic negotiation. Thus, no view should have to posit a difference in meaning to explain the distinction between high- and low-pressure terms. How much of the motivation for more exotic views is retained once metalinguistic negotiation is in the picture is another question. On that, see Sundell (2011) and Sundell and Plunkett (2013a).

'tasty' will lend itself to uses where the relevant standards are *local*. The kind of standard that a speaker will have reason to advocate for will be one that corresponds to their own likes and dislikes or to those of a tightly circumscribed experiencer group. Alphie wants to buy the Oreos because he likes Oreos. So he'll advocate for a standard of taste that corresponds to his own preferences. At the same time, he recognizes that taste in cookies varies and that Betty has corresponding goals. So he'll be reasonably well prepared to drop the issue. Even if he's ready to dig in his heels a bit, it's just cookies after all. He and Betty will only take the argument so far. That makes it low pressure.

What is the role typically played by the expression 'beautiful' in our aesthetic discourse? That's a harder question to answer. But Alvy and Mary's argument raises issues that at least for them involve stakes much higher than a decision about which cookies to buy or what to order at a restaurant. Alvy and Mary are deciding what to make of each other: how informed they are, how subtle a critic, how immune from fashion and pretension, etc. A word like 'beautiful' is deployed in contexts where the stakes—however variable and nebulous—are higher, sensitive to a wider and more culturally loaded range of considerations, and where, rightly or wrongly, speakers assume a lower degree of interpersonal variation in standards. Alvy and Mary might even believe—and they might even be right!—that there are objective, aesthetic joints in nature, and they might aim to sync up on a standard corresponding to those joints. For these reasons, a word like 'beautiful' lends itself to standards that are *expansive*. The kind of standard that a speaker will have reasons to advocate for will be one that corresponds to the aesthetic standards of a generic experiencer, or a larger or more refined experiencer group, or even, perhaps, to an experiencer sufficiently idealized so that his standards align perfectly with the aesthetic joints in nature.²¹ The fact that 'beautiful' is used in these contexts, and thus lends itself to these standards, makes it high pressure.²²

²¹ To emphasize, nothing about the view commits us, one way or another, to the truth of any form of aesthetic realism. The semantic view here is consistent with full-blown subjectivism and also with die-hard realism. That's a feature, not a bug.

²² A word like "terrible," as in dialogue (3) above, is used in a wide range of conversations. Whether it is low or high pressure on a particular occasion will depend on the background conditions and stakes on that occasion. Essentially, that's what I claim for "tasty" and "beautiful" as well, the only difference being whether a term lends itself to a

In our debates about climate control, we're unlikely to advocate for a standard of coldness of 27°. In our debates about cookies, we're unlikely to advocate for a standard of tastiness corresponding to the preferences of a fully idealized and impartial rational agent. In our debates about art, we're unlikely to argue for a standard of beauty that is sensitive only to the present whims of the individual speaker. Much more—indeed almost all of the interesting stuff—remains to be said about the standards we *do* advocate for. The point here is simply this. Nothing about the semantics of these words dictates that they lend themselves to the kinds of standards they do. There's nothing conceptually confused or semantically incompetent about an office worker negotiating on behalf of a standard for coldness of 27°. That person is confused, but what they're confused about is climate control. In the same way, there's nothing conceptually confused or semantically incompetent about a cookie buyer who advocates for a gustatory standard of taste corresponding to the preferences of an idealized rational agent. That person is confused, but they're confused about the point of eating cookies. Similarly, an art critic or a gallery owner or even just an art lover like Alvy or Mary is confused if they use 'beautiful' in such a way that it reflects nothing but their current unreflective preferences. But they're not confused about the meaning of 'beautiful.' They're confused about the value of art.

Metalinguistic negotiations involving predicates of taste thus share with Office Thermostat the first crucial feature: the functional role played by the relevant terms constrains the kind of standards speakers will tend to advocate for. But those functional roles can also vary from context to context or even over the course of a conversation. Most of the time, the standards of "tastiness" we have reason to advocate for will be highly local. But not always. Sometimes, they'll start local, but get more expansive over the course of the conversation. We might start out simply expressing our own gustatory preferences as a way to get to know each other, so our "tastiness" claims are very low pressure. Then, halfway through the conversation, we realize that our preferences are close *enough* that we should go to dinner together. But now we have to decide where to go and what to order, so all of a sudden we have reason to sync

certain kind of standard reliably enough to give the impression of being *characteristically* low or high pressure.

up on a standard that works for both of us, even in places where originally we differed.

Corresponding scenarios are possible (indeed, common) for ‘beautiful.’ We might start on the mutual assumption that our standards are the same, and that indeed they correspond to the preferences or responses of a very expansive group or a highly idealized experiencer—or even that they line up with the aesthetic joints in nature. Over the course of the conversation, that assumption is put to the test and we may if we find striking and persistent differences of opinion, give up for now on the goal of finding common ground. At that point, we employ standards that are increasingly local, and our assertions place a correspondingly lower degree of aesthetic pressure on our listener. In all of these cases, over the course of the entire conversation, our disputes reflect negotiations of standards. But the stakes of those negotiations, and our level of commitment to finding common ground, have the potential to vary dramatically from context to context and even from moment to moment within a conversation.

Metalinguistic negotiations like those in *Buying Cookies and Talking Art* also share with *Office Thermostat* the second feature: given the background conditions to a given negotiation, the standards themselves are evaluable as better or worse. As I’ve emphasized, aesthetic disputes do not take place in a vacuum. Our goals in those disputes are sometimes, but need not be, practical or concrete. Given our actual goals—buying cookies we’ll enjoy, listening to music that will move us, aligning our standards with the aesthetic joints in nature—and given the relevant range of facts about the world and about us—the properties of the objects we discuss, our palates and our education, our perceptual, cognitive, and affective apparatus more generally—some standards will be better than others.

Suppose Alphonse, on a quick viewing of a painting characterized by bright colors and sweeping lines, describes the painting as “garish.” And suppose further that the vast preponderance of other critics—well-educated and informed critics, familiar with the artist and with the genre to which the painting belongs—disagree with that description, preferring instead to describe the painting as “dynamic.” Nothing about the view under consideration here prevents us from saying that Alphonse is mistaken. He may not be mistaken in virtue of having expressed a false proposition. Perhaps, relative to the preferences of the experiencer or

standard indexically referenced in his use of the term ‘garish’—a standard corresponding to Alphie’s current and unreflective attitudes, perhaps—paintings characterized by precisely this type of bright color and sweeping line are deserving of the kind of negative evaluation conveyed by ‘garish,’ rather than the kind of positive evaluation conveyed by ‘dynamic.’ If Alphie is talking to himself or writing in his diary, we might even think that he is not mistaken in any interesting sense.

But if Alphie is having a dispute about this issue with fellow museum visitors, or gallery owners, or with art history textbook authors or simply art lovers, he is mistaken in a clear sense. He is mistaken not in virtue of literally expressing a false proposition but in virtue of advocating for a bad standard. In these contexts, the standard Alphie advocates for fails to accomplish what the participants in those conversations are trying to do. It fails to get at what is interesting, or valuable, or historically important, or delightful, about the bright colors and sweeping lines of the painting. To be clear, nothing about this semantic view *requires* that we say any of those things. A simple form of contextualism augmented with a story about metalinguistic negotiation is *consistent* with the most radical or naive forms of meta-aesthetic subjectivism. But for those who (like me) find that kind of view unappealing, the important point is that it sits just as happily with the idea that the painting really is dynamic, and that to say otherwise is a mistake. The relevant notion of *mistake* is metaphysically robust—if you’re a realist independently, you can even say the mistake is entirely objective. But even that does not require that the mistake in question consist in the expression of a false proposition. Alphie is objectively mistaken because he advocated for an objectively bad standard. This view thus not only accounts for the fact that low-pressure assertions *seem* more subjective while high-pressure assertions *seem* more objective. It’s fully consistent with a view on which low-pressure assertions *are* subjective while high-pressure assertions *really are* objective. Whether you choose to endorse the latter view depends—as it should—on philosophical considerations independent of the semantics of these terms.

What is the take-home message of the argument here? First, it is possible to respect the distinction between simple predicates of taste and properly aesthetic terms while denying that there is a deep or categorical *semantic* distinction between the two classes of terms. Low-pressure aesthetic terms are those terms that, given what we do with

them—given, that is, their functional role in our aesthetic discourse—lend themselves to standards that are *local*. High-pressure aesthetic terms are those that, given the role they play in aesthetic discourse, lend themselves to standards that are *expansive*. As noted, a great deal remains to be said about those functional roles and about the corresponding standards. But the crucial point is that this difference in our usage of the terms is a matter of what we do with them, and not a matter of their linguistically encoded context-invariant meaning.

Second, a big part of the action in our aesthetic discourse occurs at the metalinguistic level, and not at the level of literally expressed semantic content. The notion that an important part of aesthetic debate consists in an attempt by speakers to align their standards is an idea with independent plausibility, one that has been advocated by proponents of competing views in the literature on the semantics of predicates of taste. The phenomenon of metalinguistic negotiation provides the ideal mechanism for such an activity. Arguing about whether Oreos are tasty is a way of using the known taste of Oreos as a fixed point on which to negotiate what tastes we find pleasing or should reward with our purchasing decisions. Arguing about whether *The Seventh Seal* is a beautiful movie is a way of using the known properties of *The Seventh Seal* as the stable ground on which to negotiate what we should admire in film and what we should dismiss as fashion or pretension. None of this entails that coordinating standards is the *only* goal of aesthetic debate or even that aligning standards itself is an activity that cannot be tied to or evaluated with respect to real or even objective properties of the objects under discussion. Views on these latter points will depend on one's other views in aesthetics and metanormative theory.

Third, and finally, even a simple form of contextualism—once augmented with the notion of metalinguistic negotiation—provides the resources to explain the full range of perceived faultlessness and disagreement in aesthetic debate. Aesthetic pressure—the insistence that our interlocutors agree with our claims or the perception of aesthetic difference *as* disagreement—comes in varying strengths and comes and goes over the course of a conversation. It can persist intensely even when we recognize that our interlocutors are entirely justified in advancing the views that they do, and it can vanish just when we seemed most at odds. This behavior is different in kind from what we should expect in a dispute focused on the truth of a single agreed upon proposition. Simple

contextualism and metalinguistic negotiation can jointly explain how aesthetic debate resembles a negotiation, and negotiations *are* the kind of thing that can vary in intensity and be taken up or abandoned per the present needs or inclinations of the speakers. They are also the kind of activity where, finding ourselves deeply at odds, sometimes we attribute error to our conversational partner and sometimes we don't. These are the things we should be looking for in a semantic theory for the words we employ in those conversations.

References

- Barker, C. (2002). The dynamics of vagueness, *Linguistics and Philosophy*, 25: pp. 1–36.
- Barker, C. (2013). Negotiating taste. *Inquiry* 56 (2–3): pp. 240–57.
- Björnsson, G. & Finlay, S. (2010). Metaethical contextualism defended. *Ethics* 121 (1): pp. 7–36.
- Egan, A. (2010). 'De gustibus non disputandum est (at least, not always),' in R. Feldman & T. Warfield (eds.). *Disagreement*. Oxford: Oxford University Press.
- García-Carpintero, M. (2008). 'Relativism, vagueness, and what is said,' in M. García-Carpintero & M. Kölbel (eds.). *Relative truth*. Oxford: Oxford University Press.
- García-Carpintero, M. & Marques, T. (2014). 'Disagreement about taste,' *Australasian Journal of Philosophy*, 92 (4): pp. 701–23.
- Hare, R. M. (1952). *The Language of morals*. Oxford: Oxford University Press.
- Huvenes, T. (2014). Disagreement without error. *Erkenntnis*, 79/1: pp. 143–54.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28: pp. 643–86.
- López de Sa, D. (2008). 'Presuppositions of commonality,' in M. García-Carpintero & M. Kölbel (eds.). *Relative truth*. Oxford: Oxford University Press.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Marques, T. (2014). Doxastic disagreement. *Erkenntnis*, 79 (1): pp. 121–42.
- Nehamas, A. (2002). A Promise of Happiness: the place of beauty in a world of art. *The Tanner Lectures on Human Values*, 23.
- Plunkett, D. & Sundell, T. (2013a) Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13 (23).
- Plunkett, D. & Sundell, T. (2013b) Dworkin's interpretivism and the pragmatics of legal disputes. *Legal Theory*, 19: pp. 242–81.
- Railton, P. (1998). Aesthetic value, moral value, and the ambitions of naturalism. In: J. Levinson (ed.), *Aesthetics and Ethics: Essays at the Intersection*. Cambridge University Press.

- Sibley, Frank. (1959). Aesthetic Concepts. *Philosophical Review*, 68: pp. 421–50.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30: pp. 487–525.
- Stevenson, C. L. (1937). The emotive meaning of ethical terms” *Mind*, 46 (181): pp. 14–31.
- Stojanovic, I. (2007). Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy*, 30: pp. 691–706.
- Sundell, T. (2011). Disagreements about taste. *Philosophical Studies*, 155 (2): pp. 267–88.
- Sundell, T. (forthcoming). The tasty, the bold, and the beautiful. *Inquiry*.
- Zangwill, N. (2001). *The Metaphysics of Beauty*. Ithaca: Cornell University Press.
- Zangwill, N. (2005). ‘Beauty,’ in J. Levinson (ed.). *The Oxford Handbook of Aesthetics*. Oxford: Oxford University Press.