

Genuine Relativism & Expressivism

The received view is that there are two alternatives to contextualism about taste judgments: genuine relativism and expressivism (or as it is sometimes called, noncognitivism). In the taste literature, these two views are generally assumed to be *genuine alternatives* to one another – that to accept genuine relativism is to reject expressivism, and vice versa. In this section, I will argue that rigidly distinguishing genuine relativism from expressivism is a serious mistake – and one made all too often. More specifically, I will argue that genuine relativism is best conceived as a specific kind of expressivism – minimalist expressivism. To be sure, there often are substantive differences between genuine relativism and what gets called “expressivism,” but as I will demonstrate, these differences are tantamount to the dispute that obtains between newer and older versions of “noncognitivism;” the divide between genuine relativism and “expressivism” is *identical* to the divide between minimalist expressivism and “old-school” noncognitivism.

I am not merely alleging that various philosophers reject some set of linguistic conventions that I think they should accept: there are real costs to locating major differences between genuine relativism and expressivism. For one thing, the distinction between the two is often maintained by systematically misusing ‘expressivism’ to refer to old-school varieties of noncognitivism (to wit, emotivism and prescriptivism) – views which have not been seriously defended in the metaethics literature since the 1970s. To reserve the term ‘expressivism’ for these views is to completely overlook all recent work within the expressivist (i.e. noncognitivist) tradition. As a consequence, the taste literature has unnecessarily retraced many of the problems, arguments, and positions that were thoroughly addressed and resolved in the metaethics literature decades ago.¹ There are distinctions worth making (e.g. between *expressing* and *reporting*) and questions

¹ For example: Early noncognitivists like Stevenson (1944) were the first to suggest that normative disagreements might be glossed as disagreements in some noncognitive *attitude* (i.e. *attitude disagreement*) rather than disagreements in *belief* (i.e. *propositional disagreement*). Beginning with Blackburn (1984) and Gibbard (1990), contemporary expressivists have increasingly distanced themselves from this “old-school” conception of disagreement in order to accommodate the fact that moral discourse is replete with talk of moral beliefs, justification, and all the other trappings of rational discourse. Various taste contextualists (cf. Marques (2014), Sundell (2011), Huvenes (2012, 2014), Gutzmann (2016), and others) and even some taste genuine relativists (viz. Berškytė & Stevens (2022))

worth answering (e.g. what are evaluative sentences used for?) which expressivists have been seriously grappling with for decades, but genuine relativists remain virtually silent on.

Conversely, genuine relativists have made serious advances (e.g. augmenting the disagreement problem with the *faultlessness* criterion) and developed formal tools (e.g. distinguishing context-sensitivity from assessment-sensitivity) which expressivists would be wise to learn from and borrow. In brief, major interdisciplinary opportunities are squandered by misdiagnosing antagonism between expressivism and genuine relativism. I hope to expose how unnecessary and frivolous the division between these two views is, and to begin to remedy it.

In Section 4.1, I will set the stage by reviewing the genealogy of expressivism and genuine relativism – specifically, how each view positioned itself as *the* minimal but necessary departure from contextualism in their respective literatures (the metaethics and taste literatures, respectively) for nearly the same reasons. Proceeding to Section 4.2, I will demonstrate how expressivism is consistently misrepresented (most often, as emotivism) in the taste literature. Then, in Section 4.3, I will show how strikingly similar the formal semantic programs of the earliest genuine relativists (viz. Lasersohn (2005, 2017) and Kölbel (2004)) are to the semantic programs advanced by the minimalist expressivists (viz. Gibbard (1990) and Horgan & Timmons (2000, 2006)). Finally, in Section 4.4, I will argue that the differences which remain between expressivism from genuine relativism reflect shortcomings for both accounts. Resolving these shortcomings collapses both views into the same thing: evaluative nondescriptivism.

4.1 The “convergent evolution” of minimalist expressivism and genuine relativism

In the metaethics literature, there is no such thing as “genuine relativism” about moral judgments; or rather, nothing goes by that name. To be sure, there are moral relativists. The sentiment that morality “varies from culture to culture” or “depends on social customs” is extremely old (at least as old as written history). But only more recently have broad observations about the diversity of human social mores been paired with precise (or even approximately

endorse an *attitude disagreement* account of taste disagreements – often explicitly citing the early noncognitivists – without paying heed to why contemporary expressivists have steadily returned to a *propositional disagreement* reading. As a consequence, the taste literature has needed to revisit the shortcomings of *attitude disagreement* accounts (see Eriksson (2019) for a thorough treatment).

precise) semantic accounts of moral language. To the best of my knowledge, every metaethicist sympathetic to moral relativism has analyzed the semantic content of moral judgments as descriptions of the moral norms (or attitudes) of some contextually salient person or group (or else become a noncognitivist). I follow Ayer (1936) and others in calling this view – the idea that moral sentences contain a hidden indexical element – subjectivism.

Subjectivism about moral judgments is, of course, analogous to (mental state) contextualism about taste judgments. But whereas genuine relativists like Kölbel (2002) and Lasersohn (2005) parted ways with contextualism in the taste literature, there is no analogous rift among moral relativists in the metaethics literature. To my knowledge, there is no historical tradition of moral relativists who (i) deny that the propositions expressed by moral sentences vary with the context of *use*, and yet (ii) maintain that moral propositions can only be true or false relative to different contexts of *assessment*. Thanks to the work of Kölbel (2002), MacFarlane (2003), and Lasersohn (2005) on taste predicates and future contingents, many contemporary moral relativists now make the distinction between contexts of use and contexts of assessment. Even so, it is striking that so few defend forms of genuine relativism about morality – and that none did so prior to 2002. For example, Jesse Prinz (2007) carefully distinguishes “content relativism” (i.e. subjectivism) from “truth relativism” (i.e. genuine relativism), but ultimately defends the former. Earlier moral relativists, such as Dreier (1990) and Harman (1996, 2000), do not draw this distinction and, in any case, also defend versions of subjectivism (i.e. content relativism).

It is against the backdrop of moral subjectivism that noncognitivism first emerged in the metaethics literature. Classical, or “old-school” noncognitivism denies that moral judgments are *cognitive* – that they express propositions, are truth-apt (i.e. have truth-values), that moral attitudes are really beliefs, etc. The first “old-school” noncognitivists were emotivists like Ayer (1936) and Stevenson (1944), who distinguished their view from subjectivism about moral judgments purely on a semantic basis (since both noncognitivism and subjectivism are forms of anti-realism). Whereas subjectivism maintains that moral judgments express propositions *about* some person or group’s moral attitudes (or norms), emotivism denies this. As Ayer (1936) wrote,

[A]lthough [emotivism] might fairly be said to be radically subjectivist, it differs in a very important respect from the orthodox subjectivist theory. For the orthodox subjectivist does not

deny, as we do, that the sentences of a moralizer express genuine propositions. All he denies is that they express propositions of a unique nonempirical character. His own view is that they express propositions about the speaker's feelings. If this were so, ethical judgements clearly would be capable of being true or false. They would be true if the speaker had the relevant feelings, and false if he had not. And this is a matter which is, in principle, empirically verifiable. Furthermore they could be significantly contradicted. For if I say, 'Tolerance is a virtue,' and someone answers, 'You don't approve of it,' he would, on the ordinary subjectivist theory, be contradicting me. On our theory, he would not be contradicting me, because, in saying that tolerance was a virtue, I should not be making any statement about my own feelings or about anything else. I should simply be evincing my feelings, which is not at all the same thing as saying that I have them. ... [W]hereas the subjectivist holds that ethical statements actually assert the existence of certain feelings, we hold that ethical statements are expressions and excitants of feeling which do not necessarily involve any assertions. (pp.68-69)

In brief, Ayer (1936) wanted to avoid subjectivism's (bad) prediction that moral sentences and explicit descriptions of people's moral attitudes can contradict one another. So, rather than claim (as subjectivism does) that people use moral judgments to *describe* moral attitudes, emotivism instead claims that people use moral judgments to *evince* (or "express") moral attitudes. Emotivism preserves the relation that moral assertions bear to moral attitudes at the cost of "downgrading" it from *description* (as in subjectivism) to *expression*.

Stevenson (1944) took a similar turn to emotivism (from subjectivism) for a closely related reason. Whereas Ayer (1936) highlighted subjectivism's problematic tendency to identify contradictions (or disagreements) where there are intuitively none, Stevenson (1944) was motivated by subjectivism's inability to correctly diagnose the presence of disagreements in other cases. As Stevenson (1944) wrote,

[Subjectivists] emphasize disagreement in *belief about* attitudes, but not disagreement *in* attitude. ... Let us illustrate the distinction in a way that shows its full bearing on the present case. Both A and B believe that X will satisfy more of A's appetencies than it frustrates, and both believe that it will *not* satisfy more of B's appetencies than it frustrates. To that extent they agree in belief about attitudes; but we clearly must not infer that they will agree *in* attitude. (p.10)

Here, Stevenson (1944) invites us to consider a dialogue like the following:

(103) A: Although X frustrates *your* moral desires, it satisfies *mine*. X is good.
B: I agree that X frustrates my moral desires and satisfies yours, but X is not good.

Intuitively, A and B are engaged in a moral disagreement about whether X is good. But speaker subjectivism fails to identify a moral disagreement (in belief) here, since according to speaker

subjectivism, the dialogue in (103) is identical with the one in (104), where A and B fully agree (in belief):

(104) A: Although X frustrates *your* moral desires, it satisfies *mine*. X satisfies my moral desires.

B: I agree that X frustrates my moral desires and satisfies yours, #but X frustrates my moral desires.

This suffices to show that speaker subjectivism is false. To be sure, subjectivists could opt for different readings of (103), according to which A and B's moral judgments describe the attitudes of some other person or group (besides each other's), but Stevenson (1944) anticipates this:

In much the same way, both [A and B] may believe that X will satisfy more appetencies than it frustrates for both of them, taken together (A's preponderantly satisfied appetencies outnumbering B's preponderantly frustrated ones), or that it will do so for every member of some broad group of people (A being a member of this group, and B not), or that it will do so for the majority of all people (A being among the majority and B not), and so on. Here there will be agreement in belief about attitudes, in varying degrees of comprehensiveness; but again we cannot infer that A and B will agree *in attitude*. ... One can have much agreement in belief about attitudes—all that [subjectivists] can hope for—and find, *conceivably* at least, that agreement in attitude is still to be obtained. (p.10)

The problem for any variant of subjectivism, then, is that two parties can still intuitively *disagree* about moral judgments, even if they first *agree* about what every person or group's attitudes (i.e. moral desires, appetencies, etc.) are. Stevenson (1944) took this as evidence that moral disagreements are *attitude disagreements*, but the important thing is that it shows that moral disagreements cannot be glossed in the way that subjectivists hope – as mere disagreements *about* people's attitudes.

Whereas Ayer (1936) identified a *false positive* disagreement problem for subjectivism, Stevenson (1944) drew attention to a *false negative* disagreement problem for subjectivism. Either way, Ayer's (1936) and Stevenson's (1944) shift to emotivism was justified as *the* minimal departure from subjectivism necessary to (i) maintain an important relation between moral assertions and attitudes, and to (ii) avoid various disagreement problems for subjectivism. As Schroeder (2008) tells the story, appealing to the expression of attitudes rather than to their description “is the minimal departure from speaker subjectivism required in order to solve the modal and disagreement problems.” (p.73)

The emotivists were the first moral noncognitivists (or nondescriptivists); they were the first to deny that moral statements are typically used to describe facts.² In lieu of describing attitudes, emotivism claims that the meaning of moral statements is provided by the expression of noncognitive attitudes, like approval and disapproval. Contemporary noncognitivists – also known as *expressivists* – have since differentiated their view from emotivism (i.e. “old-school” noncognitivism) in two critical respects: (i) the reformulation of the *expression* relation, and (ii) the progressive adoption of *minimalism*. Both refinements occurred (roughly) simultaneously in the 1980s and early 1990s – principally in Blackburn (1984, 1993) and Gibbard (1990), and both were spurred on by attempts to resolve the big problems for emotivism. First the problems, and then the refinements.

There were two closely related problems with emotivism. In some ways, the two problems come to the same thing. First, Peter Glassen’s (1959) original problem – which I called the Grammatical Problem in *Chapter 1*. At a glance, the problem is the challenge of explaining why there is so much overlap between the grammatical and logical properties of moral and ordinary descriptive judgments (and discourse). Moral judgments (e.g. ‘stealing is wrong’) and ordinary descriptive judgments (e.g. ‘grass is green’) are both declarative sentences and appear in the indicative mood. They can both be re-worded as questions (e.g. ‘is stealing wrong?’ and ‘is grass green?’) and prefixed by propositional attitude verbs (e.g. ‘I believe that stealing is wrong’ and ‘I believe that grass is green’). It is natural for competent language users to reply to either kind of statement with sentences like ‘that’s true,’ ‘that’s false,’ ‘you are correct,’ and ‘you are mistaken.’ Moral and ordinary descriptive predicates can both be transformed into abstract property nouns (e.g. ‘wrongness’ and ‘greenness’). And so on. If moral judgments are not typically used to assert propositions, are neither true nor false, and do not express beliefs (as the emotivists maintained), then the extensive overlap between the grammatical and logical properties of moral and ordinary descriptive judgments (and discourse) would seem to be a tremendous coincidence.

² I am slightly oversimplifying here, since according to Stevenson’s (1944) first pattern of analysis for moral judgments, they describe the speaker’s attitudes *in addition* to evincing a noncognitive attitude. Stevenson counts as an emotivist since he thought the expressive function of moral judgments was their distinctive feature.

The second problem for emotivism was the one posed by Searle (1962) and Geach (1965), following a point by Frege (1879). Moral judgments (e.g. ‘stealing is wrong’) can appear in unasserted contexts, as when they appear in questions (e.g. ‘is stealing wrong?’), in negated sentences (e.g. ‘it is not the case that stealing is wrong’), and when they appear embedded in conditionals (e.g. ‘if stealing is wrong, then getting your little brother to lie is wrong’). In each of these unasserted contexts, moral judgments undoubtedly retain their *usual meanings*, since ‘stealing is wrong’ intuitively (i) answers the question ‘is stealing wrong?’, (ii) contradicts the negated sentence ‘it is not the case that stealing is wrong’, and (iii) combines with the aforementioned conditional to entail ‘getting your little brother to lie is wrong.’ The problem for emotivism is that none of *these* sentences are used to express the same attitude as ‘stealing is wrong’ – say, disapproval of stealing. Since (1) ‘stealing is wrong’ has its usual meaning when embedded in more complex sentences, and (2) these complex sentences are not used to express disapproval of stealing, it follows that the meaning of ‘stealing is wrong’ is *not* adequately described by the emotivist platitude that the meanings of moral sentences are given by the attitudes they are used to express. Following a strategy first proposed by Hare (1970), noncognitivism can only be salvaged if noncognitivists can say what attitudes more complex sentences are used to express *as a compositional function* of the attitudes they say atomic moral sentences are used to express. This has proven to be an extraordinarily difficult task. This problem goes by many names: the Frege-Geach problem, the embedding problem, and the compositional problem.

As mentioned, the Grammatological and Frege-Geach problems are closely related. The first is essentially just the challenge of answering *why* moral and ordinary descriptive judgments have so many of the same grammatical and logical properties if (as emotivism claims) moral judgments do not have truth-values, do not express propositions, do not express beliefs, etc. The Frege-Geach problem, by contrast, is really just the challenge of answering *how* to give a noncognitivist explanation for each of these grammatical and logical properties.

The critical thing to note is that straightforward solutions to the Frege-Geach problem contribute *nothing* to answering the Grammatological Problem. If I were to ask a man *why* all the objects in his house – his books, furniture, clothing, silverware, etc. – are identical with those of some other person (whom he denies ever meeting), the man would not answer my question just

by providing isolated explanations for *how* he came to own each item. To adequately answer my question, the man would have to admit to knowing the other person, or to visiting their home, or to living in the same town and shopping at all the same stores, or to the other person being his long lost twin, etc. Analogously, some explanation of *why* there is so much overlap between the grammatical properties of moral and ordinary descriptive judgments is required – and straightforward noncognitive accounts of each property (which *leave mysterious* why these are the same properties that ordinary descriptive judgments have) contribute nothing to the explanation.

“Old-school” noncognitivism – in other words, emotivism – was abandoned due to both of these problems. Or rather, emotivism was updated and refined into *expressivism*, beginning with Blackburn (1984, 1993) and Gibbard (1990). As mentioned, two critical refinements occurred: (i) the reformulation of the expression relation, and (ii) the progressive adoption of minimalism. The expression relation was retooled and clarified in response to the Frege-Geach problem. Rather than say that people use some judgments to express appraisal attitudes, the expression relation was formally generalized, such that it relates *any* sentence ‘*p*’ to what it is to think that *p* (whether ‘*p*’ is a moral sentence or not). The mental state *expressed* by a declarative sentence ‘*p*’ is just the mental state required to agree with it. Abstracting and generalizing the expression relation in this way trades in the emotivist theory of meaning for the new and improved expressivist theory of meaning:

Emotivist theory of meaning: At least part of the meaning of some sentences ‘*m*’ is provided by the noncognitive mental state they are used to *express* – roughly, the appraisal attitudes people evince (or emote) by uttering ‘*m*’.

Expressivist theory of meaning: At least part of the meaning of *any* sentence ‘*p*’ is provided by the mental state it *expresses* – roughly, what it is to think that *p*.

The main upshot of the expressivist theory of meaning is that it can (partially) sidestep the compositional problem (or Frege-Geach problem). To see how, notice that at least part of the meaning of an ordinary descriptive judgment like ‘it is raining’ is given by what it is to think that it is raining. If Mary sincerely asserts ‘it is raining’ – all she asserts is that it is raining, but one can still reasonably gather from her assertion that *Mary thinks* it is raining. To be sure, Mary does not *say* that she thinks it is raining, but her mental state (i.e. Mary’s belief that it is raining) is nevertheless what she gives voice to when sincerely asserting ‘it is raining.’ Mary does not

speak *about* her mind, but in some not insignificant sense, *Mary speaks her mind*. The generalized expression relation captures this distinction in the following simple way: to assert ‘it is raining’ is to *report* that it is raining and to *express* the belief that it is raining. Both components contribute to the meaning of ‘it is raining.’

Since there is no compositional problem for ordinary descriptive sentences like ‘it is raining,’ it follows that there is no general problem for the expressivist theory of meaning. Expressivists are still on the hook to say *how* the (noncognitive) mental states expressed by complex moral sentences are composed from the mental states expressed by their parts (if, as some expressivists suppose, moral sentences do not express beliefs, but rather some noncognitive attitude), but formally generalizing the expression relation in the way that expressivists do inspires confidence that it *can* be done, since there is no compositional problem for ordinary descriptive judgments.

The second major refinement that occurred in the leap from emotivism to expressivism was the (progressive) adoption of minimalism. *Minimalism* – or *deflationism*, as it is equally well known – was originally a view about truth, with its origins in Frege’s (1918) redundancy theory of truth, and in Tarski’s (1935) formal semantic work. According to the redundancy theory, the truth predicate is semantically redundant. That is, for any sentence ‘*p*’, the sentence ‘it is true that *p*’ expresses exactly the same proposition which ‘*p*’ does. Tarski (1935) formalized a closely related idea in his now famous equivalence schema:

Equivalence Schema of Truth: ‘*p*’ is true iff *p*.

Armed with the equivalence schema of truth, the redundancy theory purports to deflate the metaphysical extravagance assigned to the truth predicate by various “inflationist” theories (viz. the correspondence, coherence, pragmatic, and identity theories), reducing the truth predicate to a purely formal, logical role.

Ayer (1936) was aware of and echoed Frege’s redundancy theory about the truth predicate in general, but (as Dreier (2004) notes) neglected to synthesize it with his views about moral language. Had Ayer done so, he might well have been persuaded that there is nothing aberrant or confused in calling moral judgments “true” – as in ‘it is true that stealing is wrong.’

For according to the redundancy theory, this expresses nothing more than ‘stealing is wrong’ does – a basic moral judgment, the sort of which Ayer (1936) considered ripe for explanation.

Following a renewed interest in minimalism about truth in the 1980s and 1990s (in particular, the work of Horwich (1982, 1990) and Wright (1992)), Blackburn (1984, 1993) and Gibbard (1990) were the first to rectify Ayer’s (1936) mistake. They integrated limited versions³ of the equivalence schema – that is, a *minimalist* conception of truth – into their respective versions of expressivism. This immediately helped them explain why it is appropriate (and common, in moral discourse) to reply to moral statements with expressions like ‘that’s true,’ ‘that’s false,’ and ‘you are mistaken.’ They likewise employed minimalist notions of propositions and properties –

Minimalist Schema of Proposition: ‘*p*’ expresses the proposition that *p*.

Minimalist Schema of Property: object *x* has property *F* iff *x* is *F*.

– to accommodate various natural language phenomena, including the prefixing of moral judgments with propositional attitude verbs (e.g. ‘I hope that we did the right thing’), the transformation of moral predicates into abstract property nouns (e.g. ‘wrongness’), and more.

Expressivist accommodations for “moral beliefs” and “moral assertions” were not far behind. Citing the work of Crispin Wright (1992), Horgan & Timmons (2000, 2006) adopt minimalist notions of belief and assertion:

Minimalist Schema of Belief: *S* believes that *p* iff *S* has the mental state expressed by the declarative sentence ‘*p*’ (i.e. iff *S* agrees with ‘*p*’). Beliefs are the mental states that are expressed by (i.e. are required to agree with) declarative sentences.

Minimalist Schema of Assertion: Assertions are agreement-signaling (or commitment-like) utterances of declarative sentences.

³ Moral expressivists are interested in distinguishing *realist* discourse (e.g. talk of the temperatures at which various cooking oils smoke) from *anti-realist* discourse (e.g. talk of morality, by their lights). For this reason, they tend to adopt a dual view of truth – where minimalism explains the truth-talk in evaluative discourse, and the correspondence view of truth explains the truth-talk in realist discourses. To only make use of the minimalist notion would potentially collapse the very distinction that anti-realists about a particular discourse try to maintain – namely, that its subject matter is *uniquely* unreal. Minimalism is attractive to expressivists as a means to say what is common between moral and ordinary descriptive discourse. But the point of discourse-specific anti-realism is nevertheless to differentiate the two, and the exclusive use of minimalism about truth threatens to obliterate that difference.

Since people have the mental states expressed by moral judgments (i.e. declarative sentences), it trivially follows that people have moral beliefs. Likewise, since there are agreement-signaling (or commitment-like) utterances of moral judgments, it trivially follows that there are moral assertions. By this method, Horgan & Timmons (2000, 2006) hope to rescue the last part of natural language originally denied to the emotivists – namely, why competent language users ordinarily talk of moral assertions and of moral beliefs, why nobody bats an eye at embedding a moral judgment in a ‘believes that’ clause (e.g. ‘I believe that stealing is wrong’), and why there is nothing unnatural about conditionals with descriptive antecedents and nondescriptive consequents (e.g. ‘If Bob stole the coat, then he should return it’), and vice versa.

In each progressive instance of minimalism, a deflated version of a traditionally more metaphysically loaded concept is wheeled out in order to accommodate the grammatical properties that moral judgments have – in other words, the kinds of expressions and logical relations that competent language users regularly make use of in ordinary moral discourse. Once minimalism has been fully onboarded, even the grammatical problem completely vanishes. In one respect, this is a triumph over the problem, and in another respect, a complete surrender. The original problem was to answer *why* moral and ordinary descriptive judgments have so many of the same grammatical properties *if* (as the emotivists said) they don’t have truth-values, don’t express propositions or beliefs, and so on. And with the progressive integration of minimalism into emotivism – characteristic of expressivism – it should be clear what has changed: the expressivists have conceded that moral judgments *do* have truth-values, *do* express propositions and beliefs, *can* be asserted, etc. Glassen (1959) primarily took the grammatical problem to show that moral judgments (and discourse) are *cognitive*, and the progressive adoption of minimalism by expressivists ultimately produced Horgan & Timmons’s (2000, 2006) *cognitivist* expressivism. So much for the surrender – but what of the triumph? Glassen presupposed that cognitive judgments (and discourse) *just are* descriptive, but the steady adoption of minimalism is precisely what empowers expressivists to reject this. The various minimalist concepts suffice to explain the cognitivity of moral discourse, and – owing to their lack of metaphysical substance – are helpfully silent on whether moral judgments describe metaphysically rich properties (e.g. objective wrongness). Armed with (non-global) minimalism, it becomes coherent to say that moral discourse is cognitive but – for all that – nondescriptive. Perhaps unsurprisingly, this is precisely Horgan & Timmons’ (2000, 2006) conclusion.

In review, the progression from moral subjectivism to contemporary minimalist expressivism occurred in two phases. First, moral subjectivism (the view that moral judgments are used to *report* moral attitudes) gave way to emotivism (the view that moral judgments are used to *express* moral attitudes). This shift was *the* minimal departure from moral subjectivism necessary to resolve various disagreement problems. Finally, emotivism gave way to expressivism (the view that moral judgments *express* attitudes in just the same way that all judgments express attitudes). The expression relation was reformulated and various minimalist concepts were deployed. These refinements were justified as *the* minimal departure from emotivism necessary to accommodate the cognitive grammatical features of moral discourse. In a phrase, the need to account for natural language use – including intuitions of disagreement, ordinary truth ascriptions, the rational character of moral discourse, and the ubiquity of “moral belief” talk – ushered moral subjectivism toward emotivism, and, ultimately, minimalist expressivism.

With the origins and “evolution” of minimalist expressivism in view, how did genuine relativism about taste come to be? The first genuine relativists were Max Kölbel (2002, 2004), John MacFarlane (2003), and Peter Lasersohn (2005). Whereas MacFarlane developed the view to make sense of future contingents (e.g. ‘there will be a sea battle tomorrow’), Kölbel and Lasersohn simultaneously developed the view in direct response to taste contextualism – which, as we have already seen, is perfectly analogous to moral subjectivism. Both contextualism and subjectivism are anti-realist, relativize the content of judgments to the context of use, and claim that these judgments contain a hidden indexical element, such that they covertly describe the attitudes of a contextually-salient person or group.

Echoing Ayer (1936) and Stevenson (1944), Kölbel (2002, 2004) and Lasersohn (2005) were persuaded to reject taste contextualism – and to differentiate their own burgeoning view from it – principally thanks to the (faultless) disagreement problem. I reviewed the problem extensively in Section 3.5.1. For my present purposes, it is only worth noting how the argument from faultless disagreement differs from Stevenson’s (1944) disagreement problem for moral subjectivism. The principal difference is the novelty of the *faultlessness* component. To be sure, Stevenson made do without it by highlighting the fact that two speakers who agree (to any arbitrarily great extent) about what the moral attitudes of various people and groups are, can still

(intuitively) disagree about any moral judgment. By contrast, the argument from faultless disagreement functions by showing that contextualist readings of taste disagreements can secure *disagreement* (in same-subject readings), or *faultlessness* (in different-subject readings), but not both. This is a bad result for taste contextualism, since taste disagreements seem to be paradigmatically faultless.

Despite rejecting essentially the same view for nearly the same reasons, Stevenson (1944) concludes that moral disagreements are *non-propositional* (viz. attitude disagreements), whereas Kölbel (2002, 2004) and Lasersohn (2005) conclude that taste disagreements are *propositional*. What gives? I mentioned earlier that the shift to emotivism was *the* minimal departure from subjectivism necessary to resolve the disagreement problem, and I stand by that. It is no accident that the most common contextualist reply to the faultless disagreement problem highlighted the possibility of glossing taste disagreements as attitude disagreements (hence *disagreement*) involving the description of different subjects (hence *faultlessness*).⁴ This is *perfectly* analogous to Stevenson's (1944) first pattern of analysis for moral judgments, which supplements speaker subjectivism with *attitude disagreement*.⁵

There are good reasons to hold both of these hybrid *attitude disagreement* accounts of (faultless) disagreement suspect,⁶ but the important thing to notice is that these accounts *can* accommodate faultless disagreements. What they cannot accommodate is the possibility of faultless *propositional* disagreements – but whereas taste disagreements are transparently faultless, it takes more work to show that they are also *propositional* (as I took pains to do in Section 3.5.3). So, there are really two parts (or phases) to this anti-contextualist argument. First, it is argued that taste disagreements are faultless (as in Sections 3.5.1 and 3.6), where this demonstrates the inadequacy of plain old taste contextualism. *The* minimal change necessary to account for the possibility of faultless disagreements is to augment contextualism with a non-propositional account of disagreement (as in Section 3.5.2) – just as Stevenson (1944)

⁴ Cf. Marques (2014), Sundell (2011), Huvenes (2012, 2014), and Gutzmann (2016)

⁵ According to Stevenson's (1944) first pattern of analysis, 'this is good' means 'I approve of this; do so as well.' (p.21) The first part ('I approve of this') corresponds to speaker subjectivism, whereas the imperative in the second part ('do so as well') is the emotive component which (supposedly) procures attitude disagreement.

⁶ See Eriksson (2019) for the most comprehensive treatment.

augmented subjectivism with disagreement in attitude. The second phase of the argument is to show that at least some faultless taste disagreements are *propositional* (as in Sections 3.5.3 and 3.6). This suffices to show that *any* variety of taste contextualism is inadequate – but how does this lead to genuine relativism?

If taste disagreements involve two parties that believe inconsistent (or contradictory) propositions, *and* these disagreements are faultless, then either (a) it is possible for a pair of inconsistent (or contradictory) taste propositions to both be true, or (b) it is possible to believe a false taste proposition without “being at fault,” or both. In other words, either **(CF)** is false, or **(FF)** is, or both of them are:

Contradiction Falsehood (CF): If propositions p and q contradict one another, then either p is false, or q is false, or both are false.

Faultful Falsehood (FF): If a person asserts a sentence expressing a false proposition, or they believe a false proposition, then they have made a mistake (or are at fault).

Kölbel’s (2002, 2004) work emphasizes the need to refine **(FF)** for taste discourse, while Lasersohn’s (2005, 2017) work emphasizes the need to refine **(CF)** for the same discourse. To be sure, neither philosopher accepts (a) or (b) at face-value – doing so would be incoherent. Instead, both philosophers *relativize* truth (for taste discourse, at least) to yield assessment relativized versions of **(FF)** and **(CF)** that are more palatable. Kölbel’s explicit relativizations of **(FF)** are:

Believing something that is not true according to one’s own perspective constitutes a mistake. (2002, p.33)

It is a mistake to believe a proposition that is not true in one’s own perspective. (2004, p.70)

By contrast, Lasersohn’s (2005) relativizes (a), indicating a relativization of **(CF)**:

If you say roller coasters are fun, and I say they are not, I am negating the very same sentence content which you assert, and directly contradicting you. Nonetheless, both our utterances can be true (relative to their separate contexts [of assessment]). (p.684)

In brief, genuine relativism about taste was conceived of by Kölbel (2002, 2004) and Lasersohn (2005) as the only way to accommodate the possibility of faultless (propositional) disagreements. For there to be faultless propositional disagreements, either **(FF)** must be false, or **(CF)** must be false, or both of them must be false. Rather than abandon either principle outright, Kölbel and Lasersohn pursued *the* minimal departure from both principles necessary to accommodate

faultless propositional disagreements: at the very least for taste discourse, truth and falsity were deflated – in particular, relativized to perspectives (i.e. contexts of assessment). Similarly, in lieu of abolishing the relation between taste assertions and subjective attitudes, the genuine relativists chose to deflate it: rather than claim that taste assertions *describe* subjective attitudes (as taste contextualists do), the genuine relativists instead suggest that taste assertions are *evaluated with respect to* subjective attitudes.

The remarkable similarity between genuine relativism's deflated notion of "relative-truth" and minimalism about truth is not a coincidence. In Max Kölbel's (2002) *Truth Without Objectivity*, where genuine relativism began, Kölbel explicitly relativized truth to perspectives as an extension of the groundwork on minimalism about truth by Horwich (1998), Williams (1999), and others. Crispin Wright's (1992) *Truth and Objectivity* deserves special mention, since it features centrally both in Kölbel's (2002) decision to relativize (FF), where Kölbel responds to Wright's (1992) criterion of cognitive command (which assumes (FF)) – and in Kölbel's (2002) relativization of truth to perspectives, where Kölbel explicitly extends Wright's (1992) brand of minimalism about truth.

It is hardly a coincidence, then, that Horgan & Timmons's (2000, 2006) minimalist expressivism and Kölbel's (2002) genuine relativism are so similar. Apparently acting in parallel, both paid special deference to Wright (1992), and both extended (and modified) Wright's (1992) brand of minimalism about truth to their respective discourses. Indeed, the resemblance between Horgan & Timmons's (2000) and Kölbel's (2002) adaptations of truth pluralism (inspired by Wright (1992)) is uncanny:

Our view is in the minimalist spirit though we would insist [that] to understand truth minimalistically in one discourse does not commit one to minimalism in relation to every mode of discourse. ... Thus, we are inclined to advocate a kind of pluralism about truth according to which there is a univocal notion of truth even though truth ascription may involve more or less robust metaphysical commitments in relation to different areas of thought and discourse. (Horgan & Timmons, 2000, p.147)

My conclusion is therefore that our use of the predicate 'true' is associated with two distinct notions of truth, one with a restricted range of applicability, and the other globally applicable. It may be that the global use is associated with a metaphysically neutral deflationary truth notion, while the selective use is associated with a metaphysically loaded correspondence notion. (Kölbel, 2002, p.73)

The resemblance of Kölbel's (2002) view to minimalist expressivism does not stop there. Indeed, the evolution of Kölbel's own thought seems to anticipate the minimalist expressivists' quintessential détente with truth-conditional semantics. In the opening paragraph of the preface to *Truth Without Objectivity*, Kölbel (2002) writes:

I first started thinking about the metaphysical consequences of truth-conditional semantics in 1993[.] Originally I thought truth-conditional semantics was misguided because of these consequences, and in my MPhil thesis I therefore defended a form of expressivism that leads away from truth-conditional semantics. Later on, however, I realized that truth-conditional semantics is a kind of Kuhnian paradigm. Most theorists who work in semantics ... work within this framework. In fact, for many theorists 'semantics' just means truth-conditional semantics. So I started exploring the advantages of truth-conditional semantics, and enquiring whether these advantages couldn't be had without the unwanted metaphysical consequences. This book is the outcome of that enquiry. (ix)

The evolution of Kölbel's thought mirrors the development of minimalist expressivism, which arose contemporaneously with (if not slightly prior to) Kölbel's own shift from "expressivism" to genuine relativism. Kölbel's former "expressivism" is more aptly described as emotivism,⁷ since like Ayer, Kölbel seems to have thought "truth-conditional semantics was misguided" on account of its (apparent) metaphysical consequences. Terminology aside, the essential point is just that Kölbel began with an expressive view averse to truth-conditional semantics (concomitant with emotivism), and made his transition to genuine relativism principally to make good on "the advantages of truth-conditional semantics, and [to enquire] whether these advantages couldn't be had without the unwanted metaphysical consequences." I could scarcely produce a more apt description of the historical development of minimalist expressivism.

In brief, Max Kölbel found his way to genuine relativism *from* the starting point of emotivism's anti-realist metaphysics, and *by* gradually warming to the cognitivity (i.e. truth-functional character) of taste judgments and discourse. By contrast, Peter Lasersohn took something like the opposite route – beginning *from* an unyielding commitment to truth-conditional semantics, and making his way to genuine relativism *by* more seriously grappling with the genuine subjectivity of taste. As Lasersohn (2017) recounts,

At first I took [the disagreement problem] as an argument that taste (and aesthetic and moral) sentences really did make objective claims, whose truth was a matter of fact—not

⁷ This loose use of 'expressivism' (where to be an expressivist is *defined* as rejecting truth-conditional semantics) is unfortunately common among genuine relativists. I will return to this point in more detail in the next section.

opinion—despite our initial intuitions to the contrary. With no first-person element to their meanings, *fun* (and *beautiful* and *wrong*) would presumably express ordinary properties, the same for everyone, which hold of their arguments absolutely, objectively and factually. But this idea never sat quite right with me, especially as it concerned the taste examples. ... Some years later, conversations with linguists [reluctant] to acknowledge any role for real-world truth and denotation in semantic theory led me to reconsider what, from my own perspective, seemed like problematic examples for truth-theoretic semantics[.] ... It seemed to me that there must be some way to treat such sentences in a truth-and-denotation-based semantic theory, without portraying them as though they made simple claims of objective fact. (pp. x-xi)

Note that Lasersohn's commitment to the truth-theoretic semantic paradigm was so strong, that he initially assumed the falsity of taste contextualism entailed taste realism – as if contextualism were the only form of anti-realism worth considering. Indeed, Lasersohn (2017) presents his initial turn towards genuine relativism as something like a course of last resort – undertaken only to protect the “central explanatory role” of truth in semantics:

Relativism about truth has until recently been such an unpopular position in analytic philosophy that it may be difficult to see the presentation of a relativist semantic theory as anything other than a defense of relativism. Yet when I began work in this area, my purpose was not so much to promote relativism as to defend logically oriented, truth-theoretic semantics against the claim that it could provide no natural account of sentences which express subjective judgment, as opposed to making objective claims. If truth-theoretic semantics really has nothing to say about such examples, they might be seen as motivation for displacing the concept of truth from its central explanatory role in semantic theory, and lend support to the kinds of theories which treat truth conditions as a matter of only peripheral—or even illegitimate—interest to semantics. (pp.2-3)

In contrasting the different paths Kölbel and Lasersohn took, I do not intend to carve any important distinction between their views. My point is just to highlight the fact that they arrived at *the same view* from quite different background assumptions and sympathies. This fact is instructive, since it suggests that genuine relativism is the natural end result of reconciling (a) the subjectivity of taste judgments with (b) truth-conditional semantics, against the backdrop of (c) the disagreement problem for less exotic, indexical views (viz. contextualism).

Take note of how in-line with minimalist expressivism this particular combination is. Emotivism emerged from subjectivism (i.e. the indexical view of moral judgments) in order to sidestep (c*) the disagreement problem for subjectivism, whilst preserving (a*) the irreducibly evaluative character of moral language. Not yet expressivism, emotivism flouted truth-conditional semantics and suffered the consequences. Contemporary expressivists

progressively adopted minimalism in order to address emotivism's problems – in other words, to make good on (b*) the advantages of truth-conditional semantics.

It might be objected that in spite of every glaring resemblance – (i) the common root in subjectivism and contextualism, and (ii) the similarity of the disagreement problems which motivated an equally minimal departure from that common root, (iii) the need to reconcile truth-functional semantics with evaluative language, and (iv) the employment of minimalism about truth to lubricate that reconciliation – and every coincidence⁸ – that genuine relativism is, for all that, a different view than minimalist expressivism. But even this last redoubt is undermined by the historical record (to say nothing of the formal semantic overlap, to be reviewed in Section 4.3). For Kölbel (2002) qualifiedly *invites* the reader to favorably compare his view to one of the most prominent (if not *the* most prominent) forms of minimalist expressivism: Alan Gibbard's norm-expressivism (1990). Indeed, Kölbel (2002) dedicates an entire section (the last section of his sixth chapter) to the matter:

In Chapter 4, I deferred discussion of Alan Gibbard's norm-expressivism (Gibbard 1990: chapter 5). ... I shall argue that Gibbard fails to clarify one crucial issue, and that makes it difficult to decide whether he should be interpreted as a [subjectivist] or a [genuine] relativist. ... The core claim of Gibbard's semantics of normative sentences is that 'the meaning of normative terms is given by what judgments normative sentences express— what states of mind they express' (1990: 84). ... There are obvious analogies between Gibbard's factual-normative worlds and my perspectives. Just as a perspective is an evaluation of all contents, objective or not, a factual-normative world is an evaluation of all contents, normative or factual. Just as Gibbard's contents can be construed as sets of factual-normative worlds, mine could be construed as sets of perspectives. ... [L]et me highlight two differences between the [genuine] relativist [interpretation] of Gibbard's [norm-expressivism (1990)] and my own. ... These differences may not amount to much. Thus my view may well be very close to [an] elaboration, of Gibbard's view. I regard that as a desirable outcome. (pp.110-115)

It is plain that Gibbard (1990) is not a subjectivist. So, if Gibbard (1990) counts as a genuine relativist by Kölbel's (2002) lights, then Kölbel (2002) counts as a minimalist expressivist, too. I submit, then, that the genealogical record – the facts about how minimalist expressivism and genuine relativism came to be – points to a remarkable degree of overlap between the two views.

⁸ To list just a few: (i) the conspicuous historical absence of genuine relativism about morality, and of contemporary minimalist expressivism about taste. (ii) Kölbel's having been a self-described "expressivist" who adopted minimalism about truth to become the first genuine relativist. (iii) The striking degree to which Kölbel (2002) and Horgan & Timmons (2000) both rely on Wright's (1992) minimalism, and draw almost identical views about how to integrate a global minimalist notion of truth with a (slightly) more restricted correspondence notion.

If they are not exactly the same view, then they are extraordinarily similar; minimalist expressivism and genuine relativism are two sides of the same coin.

4.2 Strawman “expressivism”

My goal to show that genuine relativism is a form of expressivism would amount to little more than a quibble about linguistic conventions if the popular conception of “expressivism” did not invite so much confusion. It has become commonplace in the taste literature to treat ‘expressivism,’ ‘emotivism,’ and ‘noncognitivism’ as equivalent terms for the same thing: the view (in the mold of Ayer (1936)) that a class of judgments X (i) do not assert propositions, (ii) do not express beliefs, and (iii) do not have truth-values or truth-functional meanings (i.e. “are neither true nor false”). Instead, X are supposed to (iv) solely have meaning in virtue of the noncognitive attitudes that speakers express when using X sentences.

What’s so bad about calling this set of commitments “expressivism”? Well, for a start, the philosophers who first accepted these four commitments (beginning with Ayer (1936)) called themselves emotivists (and their view emotivism), and did not call themselves expressivists (or their view expressivism). The first philosophers to call themselves expressivists (and their view expressivism) were Blackburn (1984, 1993), Gibbard (1990, 2003), and Horgan & Timmons (2000, 2006) – and they explicitly rejected most of the commitments outlined above. In that light, the risks involved in calling what is essentially Ayer’s (1936) brand of emotivism “expressivism” should be obvious: (i) it invites confusion about what Blackburn, Gibbard, Horgan, Timmons, and other contemporary expressivists believe, and (ii) to an even greater extent, it risks completely overlooking their contributions. For if colorful oversimplifications of emotivism suffice to show what “expressivism” is (e.g. ‘stealing is wrong’ means ‘boo to stealing!’), what reason is there to engage with the subtler and more complex frameworks of Blackburn (1984, 1993), Gibbard (1990, 2003), and Horgan & Timmons (2000, 2006)? With the exception of Max Kölbel and a few others, most contemporary work in the taste literature has neglected to engage with (or even mention) the work of the minimalist expressivists.

Worse still, it is common to argue against taste “expressivism” by dusting off the cobwebs on the old problems for emotivism – the grammatical problem and the Frege-Geach problem – as though expressivists have yet to seriously grapple with (or make

progress towards solving) these problems. In this section, I will demonstrate that it is extraordinarily common in the taste literature to define and argue against “expressivism” as though it were emotivism. This habit likely predates Kölbel (2002, 2004) and Lasersohn (2005), but since their work (which established genuine relativism) is among the most influential and widely cited in the entire taste literature – I will start with them, beginning with Kölbel (2002).

Kölbel (2002) is a slightly awkward and puzzling place for this problem to occur: before making his turn towards genuine relativism, Kölbel identified as an expressivist. Indeed, Kölbel (2002) evinces an in-depth understanding of (and appreciation for) the tradition of noncognitivism, including the work of Blackburn (1984, 1988, 1998) and Gibbard (1990). Be that as it may, Kölbel (2002) nevertheless characterizes “expressivism” essentially as emotivism:

Expressivism is widely known as a thesis that semantically complements non-cognitivism in meta-ethics: if there are no moral facts to be known, if moral judgements or statements [(iii)] are not capable of being true or false, then the meaning of morally evaluative sentences cannot centrally consist in their having a truth-evaluable content. Non-cognitivists are therefore called upon to offer an alternative theory of meaning for moral sentences. What they frequently offer is expressivism, the view that [(iv)] the meaning of moral sentences must be analysed in terms of special kinds of illocutionary act, for the performance of which these sentences serve. To utter the sentence ‘Gambling is bad.’, for example, is [(i)] not to assert the truth-evaluable content that gambling is bad (there is no such truth-evaluable item), but rather [(iv)] to condemn gambling and thereby to *express* one’s moral attitude towards gambling. Whether or not ‘expressivism’ is a good label for this view (‘speech-act analysis’ might be a better one), there are highly analogous views about sentences other than moral ones [(viz. taste judgments)], which we might conveniently label in the same way. (p.43)

I have marked in bold the instances in which Kölbel (2002) explicitly assimilates one of the emotivist’s core commitments to “expressivism.”⁹ Particularly interesting is the final sentence in the above quotation, where Kölbel (2002) openly considers whether ‘expressivism’ is an appropriate label for speech-act theories like emotivism. Whatever his misgivings, Kölbel (2002) persists in using ‘expressivism’ to refer to what is essentially emotivism. This use is particularly strained when discussing Gibbard (1990):

I am not, in this chapter [entitled “Expressivism”], discussing Alan Gibbard, who in his excellent book *Wise Choices, Apt Feelings* (1990) defends what seems to be a version of expressivism. The last section of Chapter 6 is devoted especially to Gibbard’s theory. ... I believe that it is unclear whether Gibbard’s theory is a form of expressivism in the sense in which I have defined it, namely as the view that some class of problematic sentence, which

⁹ In full fairness to Kölbel, his use of ‘expressivist’ is at least partly understandable, as he drafted a version of *Truth Without Objectivity* (2002) as early as 1996, relatively contemporaneous with when the first expressivists were differentiating their views from emotivism.

cannot be analysed as having truth conditions, should be analysed as having a special kind of illocutionary force. On my own interpretation Gibbard does not fit this description. (Kölbel, 2002, p.46)

Defining ‘expressivism’ to preclude Gibbard (1990) is strained not least because Gibbard (1990) identifies as an expressivist, but also because Kölbel (2002) later calls Gibbard’s (1990) view by its name, *norm-expressivism*. It is not hard to see why Kölbel’s (2002) use of ‘expressivism’ invites confusion: among other things, it requires one to say that Alan Gibbard (among the first to identify as an expressivist) is not an expressivist – and that his view, *norm-expressivism*, is not a form of expressivism.

Since by “expressivism,” Kölbel (2002) refers almost exclusively to emotivism, it is not surprising that the objection he presents against expressivism (and recounts in exquisite detail) is just the Frege-Geach problem:

The plan of this chapter is to re-examine and to generalize a certain line of objection against expressivism, a line prominently taken by Searle (1969) and Geach (1960, 1965). (p.44)

Since my point in this section is just to establish the extent to which “expressivism” is referred to and argued against as though it were emotivism, I proceed now to Kölbel’s second and even more famous work on taste discourse – “Faultless Disagreement” (2004).

Kölbel (2004) has this to say about “expressivism”:

Expressivism: Restricting (ES) and (T). Those who want to maintain the possibility of faultless disagreement need to reject the assumptions that allowed us to generate the proof [that if two people possess contradictory beliefs, their disagreement cannot be faultless] in Section II above. ... Thus we now need to consider the possibility of dropping either the commitment to instances of (ES) [“it is true that p iff p ”] or that to (T) [“it is a mistake to believe a proposition that is not true”]. ... Expressivists might want to reject [the proof] on the grounds that [(iii)] it wrongly presupposes that matters about which we might disagree faultlessly are within the range of applicability of ‘true’. [(iv)] The judgement that Matisse is better, say, is just a sentiment or attitude and therefore [(iii)] not a candidate for truth or falsehood. [(iii)] The sentence ‘Matisse is better’ does not express a truth-apt proposition, [(iv)] but rather some non-truth-apt content of judgement. Thus, they might want to restrict the range of instances of (ES) to which they regard themselves as committed. (p.64-65)

As before, I have bolded the instances in which Kölbel (2004) assimilates “expressivism” to emotivism, replete with its antagonism to truth-functional semantics. Although Kölbel (2004) main focus here is to work out which of (ES) or (T) to give up or modify (since they jointly

entail that there are no faultless disagreements), he also argues that “expressivism” cannot permit faultless disagreements. Kölbel (2004) writes,

I do not think that expressivism helps us escape from the conclusion that there can be no faultless disagreement. Expressivists face a dilemma: either they are Ayer-style expressivists and do not believe that there can be anything like proper disagreement on evaluative matters. Or they are Blackburn-style expressivists, in which case a version of [the proof] can be re-constructed in terms they cannot object to. I shall discuss each horn of the dilemma in turn. (p.65)

One might be tempted to think that Kölbel’s (2004) division between “Ayer-style expressivists” and “Blackburn-style expressivists” carves the same distinction I have been making between emotivism and minimalist expressivism, but one would be mistaken. His point about Ayer is peculiar to him: Ayer (1936) denied (unlike Stevenson (1944)) that there are moral disagreements. Ayer (1936) instead thought that the appearance of moral disagreement depended on disagreement over background facts, and as Kölbel (2004) points out, “this [factual] dispute is not faultless.” (p.65)

As for “Blackburn-style expressivists,” Kölbel (2004) is singling out a peculiarity of Blackburn’s expressivism. Namely, that Blackburn is “prepared to say that [moral beliefs] do enter into logical relations, even though they are not strictly speaking truth-evaluable.” (pp.65-66) Blackburn is peculiar in that he attempts to re-engineer the inference properties of logical connectives for noncognitive attitudes, without giving up emotivism’s (or “old-school” noncognitivism’s) core commitment to denying the truth-evaluability of moral sentences. Owing to being the first expressivist (i.e. the first to reformulate the *expression* relation and integrate some form of minimalism), Blackburn presents something of a vague, borderline case. Prior to his (1984), Blackburn did not invoke minimalism about truth at all, and neatly fit in among the “old-school” noncognitivists. By his (1984), he begins to invoke minimalism, but it is mostly relegated to explaining truth-talk, and does not extend to beliefs or assertions. Crucially, minimalism plays no role in Blackburn’s (1984) explanation of the inferential properties of logical connectives – where Blackburn takes it upon himself to independently reconstruct the inferential properties of logical connectives (e.g. disjunction and the conditional). Many of the expressivists to come after Blackburn (1984) – including Gibbard (1990, 2003) and Horgan & Timmons (2000, 2006) – adopt minimalism to a greater extent, and forgo Blackburn’s “independent reconstruction” approach.

At any rate, Kölbel (2004) tells us that if Blackburn-style “expressivists” succeed in reconstructing the inferential properties of logical connectives *without* appealing to the notion of truth, that they “will [still] want to accept” a truth-less version of (T), namely (T*):

(T*) If not-p, then it is a mistake to believe that p. (Kölbel, 2004, p.66)

Kölbel (2004) then points out that (T*) still suffices to prove that no disagreement can be faultless:

A3	(1) A judges that p	... Assumption
	(2) B judges that not-p	... Assumption
	(3) p	... Assumption
	(4) not-not-p	... 3
	(5) B has made a mistake	... 2,4, (T*)
	(6) Not-p	... Assumption
	(7) A has made a mistake	... 1, 6, (T*)
	(8) Either A or B has made a mistake	... 3-7, CD (Kölbel, 2004, p.66)

Kölbel (2004) pretty much leaves the argument there, with the apparent problem for “Blackburn-style expressivism” being that the view has no tools to further modify (T)/(T*) to permit only some disagreements to be faultless.

The problem seems to be that “Blackburn-style expressivism” is under-equipped to distinguish objective truths (propositions, assertions, and beliefs) from their non-objective counterpart(s). Horgan & Timmons (2000, 2006) explicitly distinguish the two, and it is fairly clear how they can modify (T) to achieve the result that some disagreements are faultless, and some not:

(T-T) If *p* is FALSE, then it is a mistake to believe that *p* is TRUE.

(T-B) If *p* is FALSE, then it is a mistake to BELIEVE that *p*.

In both of the equivalent principles above, the basic maneuver is to differentiate minimalist notions of truth, falsity, and belief from their respective correspondence notions – where the correspondence notions are represented in all capital letters as TRUTH, FALSITY, and BELIEF. A minimalist belief that *p* can just be defined (as in Section 4.1) as the mental state expressed by the declarative sentence ‘*p*.’ That is, minimalist belief that *p* can just be defined as having the mental state required to agree with ‘*p*.’ Now, note that the equivalence schema (ES) only trades on *minimalist* truth. It tells us that to believe *p* is true is just to believe that *p*. By contrast, to

believe p is TRUE is to believe that p corresponds to the way the world is (or equivalently, to BELIEVE that p). From here it is relatively clear what expressivists should deny to permit faultless taste disagreements without rendering every disagreement faultless. They should deny that taste sentences can be either TRUE or FALSE. That way, taste disagreements can be faultless, whereas disagreements about matters that are either TRUE or FALSE will always be faultful.¹⁰

The details aside, the basic contour of the minimalist expressivist maneuver above is practically identical to the genuine relativist's. The genuine relativist makes conceptual space for faultless disagreements (without rendering every disagreement faultless) by differentiating two notions of truth: one *with* and one *without* objectivity. And so can the expressivist. The genuine relativist calls these two notions 'truth' and 'truth in a perspective' (or 'truth relative to a context of assessment') while the expressivist calls them 'TRUTH' and 'truth.' There may be substantive differences between the two approaches (as I will explore in Section 4.3), but then again, there are substantive differences between different forms of minimalist expressivism.

As concerns the present section, the important takeaway is that Kölbel (2004) uses 'expressivism' exclusively to refer to views which strictly deny that a judgment class has truth-values or truth-functional meanings. This includes Ayer's (1936) emotivism, as well as Blackburn's (1984) attempts at reconstructing the inference properties of logical connectives without employing a univocal, minimalist notion of truth. Either way, Kölbel (2004) artificially defines 'expressivism' to exclude more recent developments where minimalism (and especially truth pluralism) has been adopted and employed to a greater extent. It should come as no surprise, then, that Kölbel's (2004) argument against "expressivism" only serves to demonstrate why expressivists need more minimalism than Blackburn (1984) permits (and why they have increasingly come to rely on it since).

¹⁰ The minimalist expressivist will still have some work to do. Even if 'sushi is tasty' is neither TRUE nor FALSE, it is still intuitively some kind of mistake for S to believe that sushi is tasty if S does not enjoy sushi. So, even if (T-T) / (T-B) are sufficient to permit some disagreements to be faultless and some to be faultful, it cannot (yet) successfully diagnose the mistake involved in having a taste belief misaligned with one's preferences.

I turn now to Lasersohn's (2005) treatment of "expressivism," which is (in my estimation) more characteristic of the norm in the taste literature. Without Kölbel's misgivings about how to make sense of contemporary minimalist expressivists, Lasersohn (2005) writes,

Expressivism: ... [Expressivists] claim that [(iii)] the problem is in assuming that [taste sentences like 'this is fun'] are true or false at all. If we encounter trouble assigning a truth value, maybe the trouble comes from our assumption that these examples even have truth values. ... The categories of truth and falsity apply only to certain illocutionary acts (assertions, or statements), so the obvious way to pursue [expressivism] is to deny that utterances of [taste sentences] involve the performance of such acts. [(iv)] The most obvious alternative illocutionary act we might appeal to here is what we might call an act of "affective expression"; this gives us [expressivism]: Treat utterances of sentences like *This is fun* or *This is tasty* as non-assertive acts of affective expression. [For instance]:

- (28) Whee!
- (29) Mm-mm.
- (30) Oh, boy!

[(i)] Utterances like these don't make assertions; [(iv)] they just express some inward mental or emotional state. [According to expressivism], we would be assimilating sentences like *This is fun* or *This is tasty* to these. (pp.656-657)

Across the various bolded points in the excerpt above, Lasersohn (2005) attributes nearly all of emotivism's core commitments to "expressivism." It should come as no surprise, then, that he appeals to truth-talk and the Frege-Geach problem as the primary reasons to reject "expressivism." As Lasersohn (2005) writes,

An obvious problem with this idea is that we can deny sentences like *This is fun* or *This is tasty*, say they're not true, etc.; but you can't do the same thing with these non-assertive utterances of affective expression. There is nothing very strange about (31), for example:

- (31) John: This is fun!
- Mary: That's not true! This isn't fun at all!

Example (32), in contrast, is pretty bizarre:

- (32) John: Whee!
- Mary: ??That's not true! This isn't fun at all!

Mary's utterance is anomalous even though John's utterance of *Whee!* strongly implicates that whatever activity they're engaged in is fun. Additionally, sentences like *This is fun* or *This is tasty* can appear embedded under truth-functional connectives and other logical operators, and participate in the usual logical consequence relations which such embeddings give rise to. One would like to preserve the idea that (33) is an ordinary example of Modus Ponens, for example:

- (33) If there is a loop, the roller coaster is fun.
- There is a loop.
- Therefore, the roller coaster is fun.

But it is quite hard to see how to maintain this idea if sentences like *The roller coaster is fun* do not have truth values. See, e.g., Geach (1965) for additional discussion. (p.657)

Truth-talk and the validity of *modus ponens* arguments were originally issues for Ayer's (1936) brand of emotivism. Indeed, as I reviewed in Section 4.1, these (and related) problems for emotivism were what motivated noncognitivists' turn away from emotivism and towards expressivism in the first place. Whether or not expressivism has adequately answered these problems, to recycle them as original problems for “expressivism” – *without qualification* – is either to misrepresent contemporary expressivism, or else to completely ignore it.

Since genuine relativism about taste began with Kölbel (2002, 2004) and Lasersohn (2005), I have basically shown that expressivism has been characterized and argued against as if it were emotivism from the outset of the genuine relativist tradition. The work that remains is to show that this remains a *live* mistake. Expressivism continues to be misrepresented (essentially, *as* emotivism) in the work of contemporary contextualists (e.g. Zouhar (2019)) and genuine relativists (e.g. Lasersohn (2017)) alike. Even more surprisingly, the misrepresentation of expressivism is perpetuated by philosophers who explicitly endorse (at least partly) “expressivist” accounts (e.g. Clapp (2015), Gutzmann (2016), and Berškytė & Stevens (2022)) – a development which underscores the depth of the issue. I will briefly show how the misrepresentation of expressivism appears in each of these five works, proceeding chronologically.

Clapp (2015) defends what he calls a “broadly expressivist,” non-alethic account of the semantics of taste sentences. (p.1) His account is fairly unique in the taste literature – mixing some of emotivism’s characteristic features with an account of disagreement adapted from genuine relativism. Clapp (2015) also distinguishes between “metaphysical truth conditions” and “mere semantic truth conditions” in a manner reminiscent of the truth pluralism invoked by the minimalist expressivists. (p.19) He joins Kölbel in being one of the few philosophers in the literature to even mention some of the minimalist expressivists (viz. Blackburn (1984, 1988) and Gibbard (1990, 2003)), whose views Clapp (2015) distinguishes from his own.

For all that, the two “broadly expressivist” commitments of Clapp’s (2015) account – which he calls “the *negative semantic thesis*” and “the *positive pragmatic thesis*” – are recognizably emotivist. (p.1) As Clapp (2015) writes,

[The problems for contextualism and genuine relativism] are suggestive of a *broadly expressivist* approach that incorporates both the *negative semantic thesis* that [(iii)] utterances of simple sentences containing PPTs do not encode truth-conditional content, and the *positive pragmatic thesis* that [(iv)] the illocutionary point of asserting and denying such sentences is to *pragmatically express* [(ii)] non-doxastic mental states. (pp.14-15)

Clapp's (2015) "*negative semantic thesis*" is most directly comparable to Blackburn's (1984) limited use of minimalist resources. For as Clapp (2015) later clarifies, he only denies that taste sentences encode "metaphysical truth conditions" – by contrast, he claims that *they do* encode "mere semantic truth conditions." Unlike Blackburn (1984), however, Clapp's (2015) account is still recognizably an illocutionary speech-act theory (akin to emotivism) – a feature which the expressivists shed in their turn away from emotivism. That is, Clapp (2015) affirms the old-school, emotivist idea that *expressing* is something *people do* when they use normative sentences. As I reviewed in the previous section, the expressivists reformulated and generalized the expression relation, defining it as a relation borne between *sentences* and *mental states* (rather than defining it as a speech-act). Schroeder (2010, 2023) makes this point quite clearly:

The distinction between the ordinary sense[s] of 'express' and the expressivist's use of 'express' as a theoretical term for the relationship between 'grass is green' and the belief that grass is green – whatever that relationship turns out to be – is useful for helping us to understand why even though some of the early emotivist theories *anticipated* expressivism in important ways, none of them really fully appreciated the fundamental idea of expressivism. ... [E]ven Ayer ... described his view by saying that moral sentences 'express' states of mind. So if we were to classify his view casually, ... we might be tempted to call Ayer an expressivist. But Ayer did not yet understand the fundamental idea of expressivism, and his suggestions about the meaning of 'stealing is wrong', though colorful and suggestive, therefore do not take the form of an expressivist theory of meaning. ... [I]t is better to think of [expressivism] as a significant departure from the earlier generations of noncognitivist theories. Ayer, Stevenson, and Hare all aspired to tell us the meaning of what [moral sentences] are used to do. Such [emotivist] theories can be thought of as essentially *speech act theories*[.] (Schroeder, 2010, pp.73-74)

In that light, Clapp's (2015) account should be regarded as a kind of hybrid emotivism – it is "broadly expressivist" in the sense that it is broadly emotivist.

In "If expressivism is fun, go for it!" (2016), Gutzmann deliberately attempts to defend a version of the "expressivism" which Lasersohn (2005) rejects. Since we have already seen that the "expressivism" which Lasersohn (2005) anticipates is essentially just emotivism, it should come as no surprise that Gutzmann's (2016) "expressivism" is as well. As Gutzmann (2016) writes,

One of the alternative options that Lasersohn reject [sic] is *expressivism*, according to which [(iv)] utterances involving PPTs are akin to expressive speech acts (Searle 1969). The aim of this article is to show that despite of [sic] Lasersohn's (2005) dismissal, such an approach is not only viable, but even fares [sic] better than his approach in some respects. ... [T]he problem of faultless disagreement arises from the assumption that PPTs are true or false. Expressivism solves the problem by [(iii)] denying just that basic assumption of a truth-conditional analysis of PPTs. This corresponds to the fourth option for a solution to faultless disagreement that Lasersohn discusses (and rejects): ... Option 4: [(iii)] Deny that truth and falsity are involved. (Lasersohn 2005: 656) ... [O]ne way to spell this out ... is [(i)] to assume that PPT-utterances are not considered as proper assertions [of] a proposition as true. Instead [taste utterances] are considered "as non-assertive acts of affective expression" (Lasersohn 2005: 656). That is, [taste utterances (iv) are] akin to purely expressive utterances of interjections [like 'Whee!' and 'Mm-mm.']. (pp.4-5)

To be sure, Gutzmann (2016) defends a hybrid theory which augments a basic contextualist truth-conditional semantics with "expressivist" resources (viz. *attitude disagreement*). I discussed this approach within the broader strategy of augmenting contextualism with non-propositional forms of disagreement in Section 3.5.2. In the previous section, I also reviewed Stevenson's (1944) emotivism, and mentioned that (according to his first pattern of analysis) he augments a basic subjectivist semantics with *attitude disagreement*. In this respect, Gutzmann's (2016) hybrid "expressivism" is practically identical to Stevenson's emotivism – a comparison which Gutzmann (2016) invites:

However, as correctly identified by Lasersohn (2005), [a purely] indexical analysis leaves the intuition about disagreement completely unaccounted for. ... We can keep the indexical analysis on the truth-conditional layer [to account for the intuition of faultlessness] and account for the intuition of contradiction [i.e. disagreement] in the use-conditional dimension of PPT-utterances. ... That is, much like in emotive approaches to normative predicates like *good* or *bad* (Stevenson 1937), the expressive component of a PPT-statement contains an affective expressio[n] of a deontic attitude. ... Stevenson (1937: 23) argues that expressive, or as he calls it, *emotive* meaning, "is a tendency of a word, arising through the history of its usage, to produce (result from) affective responses in people." (pp.16-17)

Take note of (i) the near total overlap between Gutzmann's (2016) hybrid "expressivism" and Stevenson's (1944) emotivism, and (ii) Gutzmann's (2016) casual indifference between his own use of 'expressive meaning' and Stevenson's (1937) use of 'emotive meaning.' These are among the most transparent examples of how 'expressivism' is used to refer to emotivism in the taste literature.

Next up is Lasersohn's (2017) full-length book. There is not much to say about Lasersohn's (2017) discussion of "expressivism" – and that fact is itself noteworthy. Lasersohn

(2017) dedicates just three and a half pages to discussing and dismissing “expressivism,” and the content is largely recycled from his (2005). Lasersohn (2017) writes,

Expressivism ... [Expressivists] claim that [(iii) taste sentences] lack truth values entirely. ... [Perhaps] we were wrong to assume that sentences like *Roller coasters are fun* or *Licorice is tasty* even have truth values to begin with. ... [Expressivists] claim that [(iii)] these sentences lack truth values because [(iv)] they are used to perform some other sort of direct illocutionary act [(i)] besides making assertions, and that the assignment of truth values and truth conditions only makes sense for assertive speech acts and not for other kinds. ... One obvious possibility is [(iv)] that they simply express some aspect of the speaker’s affective or emotional state. ... [C]onsider some examples which less controversially serve merely to express affective state—including expressing a sense of fun or tastiness:

(29) a. Whee!
b. Mm-mm.
c. Oh, boy!

Utterances like these do not make assertions; they serve as conventional expressions of what the speaker is feeling. In claiming that sentences like *Roller coasters are fun* or *Licorice is tasty* likewise [(iv)] are used to perform acts of affective expression [(i)] rather than make assertions, [expressivists] would be assimilating them more-or-less to examples like those in (29). (pp.35-36)

Published more than a decade after Lasersohn’s (2005) initial foray into the taste literature, Lasersohn’s (2017) treatment of “expressivism” is practically identical. Now, as then, Lasersohn attributes emotivism’s core commitments to “expressivism,” without any apparent concerns about how to accommodate contemporary expressivism. Lasersohn’s (2017) arguments against “emotivism” are mostly familiar: truth-talk and the Frege-Geach problem, presented more-or-less as they appear in Lasersohn (2005). Lasersohn’s (2017) original contribution is the addition of a third problem:

In addition to the concern that taste sentences appear to function in logic as though they had truth values, we may add the concern that they function *grammatically* just like sentences which are used to make direct assertions about matters of fact. ... [S]entences expressing matters of opinion have the same internal syntactic structure as sentences expressing matters of fact, have the same external syntactic distribution, show the same alternations [sic] in mood and sentence type, appear as complements to the same speech act and attitude verbs—in fact, as far as the grammatical system of the language is concerned, they appear to behave in every way like ordinary truth-value-bearing clauses. A linguistic analysis [viz. “expressivism”] which treats them as fundamentally different would seem to be missing a key point. (p.38)

Perhaps unsurprisingly, this addition is unmistakably a version of Glassen’s (1959) grammatical problem for emotivism. So concludes Lasersohn’s (2017) treatment of “expressivism.” In a phrase, Lasersohn (2017) persists in defining “expressivism” and objecting to it as if it were emotivism, and refrains from even mentioning the minimalist expressivists.

In “On the insufficiency of taste expressivism” (2019), Zouhar takes an explicitly anti-“expressivist” stance. For all that, he commits the now familiar mistake of reducing expressivism to emotivism by fiat. As Zouhar (2019) writes,

It is sometimes claimed, however, that [(iii)] there are no objective facts that would make taste utterances true or false, but they [(iv)] are rather supposed to express non-propositional content. Different theories provide different explanations of what this kind of content is: they invoke affections or evaluative mental states, or evaluative attitudes toward something, to name just a few notable examples. When one utters “This is tasty,” [(i)] instead of expressing a proposition that the thing designated by “this” instantiates the property expressed by “is tasty,” [(iv)] one expresses a content that consists in one’s manifestation of one’s positive evaluative [(ii)] non-doxastic) attitude toward, or affection for, the thing in question. This approach can be labeled *taste expressivism*, since it is akin to expressivist approaches developed for some other kinds of discourse (most notably, moral discourse). (pp.5-6)

In the bolded instances above, Zouhar (2019) attributes all four of emotivism’s core commitments to expressivism. Zouhar (2019) can hardly be blamed, however, since he is responding directly to Clapp (2015) and Gutzmann (2016), who – as we have already seen – define their own “expressivism” in those terms. It should come as no surprise that Zouhar’s (2019) primary objection to both views is that they cannot accommodate truth-talk.

Finally, I turn to consider Berškytė & Stevens (2022), who liken their view to a “combination of Assessment-Sensitive Relativism with an expressive component.” (p.61) That is, Berškytė & Stevens’s (2022) view augments genuine relativism with *attitude disagreement*, analogous to the manner in which Gutzmann (2016) augments contextualism with *attitude disagreement*. Berškytė & Stevens’s (2022) primary reason for supplementing genuine relativism with *attitude disagreement* is to sidestep a disagreement problem they develop for pure (i.e. unhybridized) genuine relativism.¹¹ The result is what Berškytė & Stevens (2022) call “*Expressive-Relativism*.” As they write,

Our proposal holds that when a speaker uses a PPT, not only do they express a truth-conditional content (which is evaluated relative to an individual), they are also [(iv)] expressing a non-descriptive attitude. It is the clash of attitudes between the speakers that provides the basis for disagreement, rather than a contradiction in the descriptive content. In short, the relativist semantics accounts for the faultlessness [intuition], whereas the

¹¹ I reviewed and responded to Berškytė & Stevens’s (2022) disagreement problem for genuine relativism in Section 3.7.3. At a glance, their worry is that propositions without objective truth-conditions cannot be contradictory in the strong sense that requires objectivity. My response was to point out that nothing precludes non-objective propositions from being “contradictory” in a weaker sense. That they are non-cotenable (i.e. cannot both be accepted within the same perspective) suffices to show that they are “contradictory” in a weaker sense.

expressivist semantics will account for the disagreement [intuition]. Furthermore, these two semantic aspects are not independent[.] PPTs, in expressing an attitude ... commit speakers to a restricted range of evaluations of the content expressed. Cases of faultless disagreement are cases where interlocutors commit themselves to complement sets of parameters that the truth-conditions are relativised to. Disagreement over matters of taste ... does not involve contradicting one another, it involves refusing to endorse any shared perspective on the subjective judgements expressed. (Berškytė & Stevens, 2022, p.33)

Merely in conceiving of genuine relativism as compatible with “expressivism,” Berškytė & Stevens (2022) markedly improve on the depiction of expressivism by most genuine relativists. Having said that, Berškytė & Stevens’s (2022) invocation of *attitude disagreement* is essentially emotivist in character. Case in point, Berškytė & Stevens (2022) treat *expression* along the lines of Stevenson’s (1963) emotivism – as a speech act – as an illocutionary act to be performed by speakers, rather than as a technical relation obtaining between sentences and mental states.

I submit that I have made my point – perhaps excessively so. In the taste literature, it is strikingly common to define and argue against “expressivism” as if it were synonymous with emotivism. Worse still, “expressivism” is generally defined in ways that explicitly exclude the minimalist expressivists (to wit, Blackburn (1984, 1993), Gibbard (1990, 2003), and Horgan & Timmons (2000, 2006)) from qualifying as expressivists at all. This is problematic not least because they were the first to identify as expressivists, but also because it reflects a pervasive misunderstanding of what distinguishes expressivism from emotivism. Specifically, the reformulation of the expression relation and the progressive adoption of minimalism – the hallmarks of metaethical expressivism – are practically absent from the taste literature’s treatment of “expressivism.” As a result, the taste literature has *de facto* blindfolded itself to many of the most significant achievements of metaethical expressivism. The costs of this are considerable: powerful tools developed decades ago remain overlooked, while old problems and untenable positions are unnecessarily revisited. Most troublingly, the taste literature has largely convinced itself that genuine relativism and expressivism are incompatible views. In the previous section, I showed that they are strikingly similar programs designed to address the same problems. Proceeding now to Section 4.3, I examine the overlap between the formal semantic apparatuses of genuine relativism and minimalist expressivism.

4.3 Formal semantic overlap

In Section 4.1, I demonstrated that genuine relativism and minimalist expressivism emerged under strikingly similar circumstances. Both views aim to reconcile (i) the irreducible “subjectivity” of evaluative language with (ii) its cognitive grammatical features (e.g. its pervasively truth-functional character) against the backdrop of (iii) persistent disagreement problems for more conventional, context-sensitive, or indexical analyses. While addressing similar problems with similar resources is one thing, advancing near-identical formal semantic frameworks is quite another.

In this section, I will compare the formal semantic apparatuses of the first two genuine relativists – Kölbel (2002) and Lasersohn (2005, 2017) – with those of two early minimalist expressivists, namely Gibbard (1990) and Horgan & Timmons (2000, 2006). My aim is to show that these four semantic frameworks exhibit marked similarities. If not variants of one another, then they are substantially compatible approaches. I will proceed chronologically, beginning with the formalism at the heart of Gibbard’s (1990) *norm-expressivism*.

4.3.1 Gibbard’s *Norm-Expressivism*

Strictly speaking, Gibbard’s (1990) formalism only accounts for a narrowly-defined set of normative judgments – specifically, judgments about what is rational, or what “it makes sense” to do, think, or feel. Gibbard (1990) then accounts for *moral* judgments derivatively, in terms of normative judgments. Since *taste* judgments are distinct from both, however, my purposes are best served by outlining Gibbard’s (1990) formal semantic apparatus for normative judgments (and bypassing Gibbard’s (1990) more complex account of moral judgments).

In order to model evaluative (i.e. normative) mental states, Gibbard (1990) differentiates between evaluative (i.e. normative) mental states and ordinary descriptive (i.e. fact-representing) mental states. He models the set of an ordinary person’s descriptive mental states (i.e. factual beliefs) as B , and the set of norms (i.e. normative beliefs) a person accepts as N . (p.93) Together, the set of a person’s factual and normative beliefs can be modeled as an ordered pair, $\langle B, N \rangle$. Ordinary people aren’t fully opinionated; they don’t assign truth-values to (i.e. evaluate) every possible factual and normative statement. Gibbard (1990) models a consistent evaluation of

every factual statement – a *possible world* – as w , whereas he models a consistent evaluation of every normative statement – a *complete system of norms* – as n . Taken together, a consistent evaluation of all factual and normative statements is modeled as the ordered pair $\langle w, n \rangle$ – a *factual-normative world*.

Gibbard's (1990) formal apparatus explicitly piggybacks on standard possible world semantics – which models the contents (or intensions) of sentences as the *sets* of possible worlds in which they hold. So, it should come as no surprise that Gibbard (1990) models the contents (or intensions) of normative sentences as *sets* of factual-normative worlds. In particular, Gibbard (1990) tells us that “[n]o matter how complex a normative statement is, we can still represent its content by a set of factual-normative worlds—by the set of all factual-normative worlds for which the statement holds.” (p.96) Now, to hold in a possible world w is one thing, but what does it mean for a sentence to hold in a factual-normative world $\langle w, n \rangle$? Gibbard (1990) cleverly defines holding in a factual-normative world in terms of the more basic notion of holding in a possible world. That way, so long as we understand what it means for a factual sentence to hold in a possible world, we will understand what it means for a sentence to hold in a factual-normative world. As Gibbard (1990) writes,

What is it for S to hold in $\langle w, n \rangle$? Earlier I spoke of the descriptive predicate that, for a given system of norms N , N -*corresponds* to a normative predicate. Take the normative predicate ‘rational’, for instance: its N -corresponding descriptive predicate is ‘rational according to N ’, or ‘ N -permitted’. Now to settle whether normative statement S holds in factual-normative world $\langle w, n \rangle$, we do the following. Replace each normative predicate in S with its n -corresponding descriptive predicate. That yields a purely descriptive statement S_n . Then normative statement S holds in $\langle w, n \rangle$ if and only if S_n holds in w . (p.96)

So, to check whether the normative sentence ‘donating to charity is rational’ holds in the factual-normative world $\langle w, n \rangle$, we would simply check whether the descriptive sentence ‘donating to charity is permitted by n ’ holds in w .

On this analysis, basic factual sentences (e.g. ‘the sky is blue’) will hold in a factual-normative world $\langle w, n \rangle$ if and only if they hold in w – regardless of the normative component, n . By definition, basic factual sentences contain no normative predicates. Hence, the operation of replacing a basic factual sentence’s (nonexistent) normative predicates with their “ n -corresponding” descriptive predicates always returns the original factual sentence, regardless of the normative component, n . It follows that the content of a basic factual sentence like ‘the

sky is blue' is given by the set of all factual-normative worlds which combine (i) any possible world w in which it holds, with (ii) any complete system of norms n .

By contrast, whether basic normative sentences (e.g. 'donating to charity is rational') hold in a given factual-normative world will generally depend on both parameters (i.e. including the possible world component w).¹² That is, even if a basic normative sentence like 'donating to charity is rational' holds in some factual-normative world $\langle w_1, n_1 \rangle$, it *might not* hold in a different factual-normative world $\langle w_2, n_1 \rangle$ which contains the same complete system of norms n_1 . For instance, consider the complete system of norms n_s which permits alleviating world hunger and forbids funding organized crime. Now suppose that charities in w_h exclusively work to reduce world hunger, while charities in w_c exclusively fund organized crime. In that case, 'donating to charity is rational' holds in $\langle w_h, n_s \rangle$ but not in $\langle w_c, n_s \rangle$.

I mentioned earlier that Gibbard's (1990) formal apparatus derives "holding" in a factual-normative world from the standard notion of "holding" in a possible world. So, if we know what it means to "hold" in a possible world, we will know what it means to "hold" in a factual-normative world. But what does it mean for a proposition to "hold" in a possible world? It is just for it to be *true* in (or according to) that world. What should we make of this notion of truth – of "holding" in a possible world? Is this notion of truth the ordinary, correspondence notion – the one invoked if I were to ask, for instance, whether it's *true* that the sky is blue? Not quite. For an ordinary proposition to be correspondence-true (or TRUE, if we prefer) is for it to hold in the *actual* world. So, "holding" in a possible world – and hence, "holding" in a factual-normative world – is a more primitive notion of truth. But if not the correspondence notion of truth, what is it for a proposition to "hold" in a world, or factual-normative world? It is *de facto* a minimalist and *relative* notion of truth; for a proposition to hold in a world is for it to be true *relative* to that world. Correspondence-truth (i.e. TRUTH) can then be defined as "holding" in the *actual* world – or in Gibbard's (1990) case, as "holding" in the set of all factual-normative worlds indexed to the actual world $w_@$ (i.e. holding in $\langle w_@, n \rangle$, for all values of n).

¹² The only exceptions are (i) tautological normative sentences, such as 'it is rational to do what is rational' and 'if giving battle is forbidden, then giving battle is not permitted', and (ii) *categorically* normative sentences, such as 'lying is always irrational' and 'murder is categorically forbidden.'

Strictly speaking, Gibbard's (1990) formal apparatus is compatible with realism about rationality (i.e. realism about what "it makes sense" to do, think, or feel) if and only if there are normative facts, such that at least some normative statements hold in $w_{@}$. Although Gibbard's (1990) apparatus allows for the possibility of normative facts, Gibbard (1990) himself ultimately defends skepticism about normative facts. Specifically, he appeals to a kind of Darwinian socio-biological argument to the effect that it is unlikely that there are any normative facts – and *a fortiori*, that no normative statements hold in all factual-normative worlds indexed to the actual world. As Gibbard (1990) writes,

In explaining why we make the normative judgments we do, I found normative facts superfluous. In the end, to be sure, I do deny that there are normative facts—but only in the end. Normative facts, if there were any, would be the facts of the special kind represented, naturally or artificially, by normative judgments. That is what would pick them out as normative. As it turns out, I claimed, our making these judgments can be explained without supposing they represent facts of any special kind. Thus at the end of the argument we can conclude that, at least in this sense, there are no normative facts. (p.122)

One reason to formally differentiate between factual and normative evaluations (i.e. between possible worlds w and complete systems of norms n) – as Gibbard's (1990) formalism does – is to accommodate skepticism about normative facts. For if Gibbard (1990) thought that there were normative facts, there would be little reason to differentiate between factual evaluations (w) and normative evaluations (n) – as the latter would be a subset of the former. Bifurcating the set of all sentences in this way allows Gibbard (1990) to reconcile the truth-functional character of normative sentences with skepticism about normative facts. For according to Gibbard's (1990) apparatus, normative sentences may be minimalistically-true relative to a set of factual-normative worlds, and for all that, neither TRUE nor FALSE, since normative statements are neither true nor false in the set of all factual-normative worlds indexed to the actual world (i.e. neither true in every $\langle w_{@}, n \rangle$, for all n , nor false in every $\langle w_{@}, n \rangle$, for all n).

In review, Gibbard's (1990) formal semantic apparatus is an extension of possible-world semantics. He distinguishes consistent evaluations of all factual statements (i.e. possible worlds w) from consistent evaluations of all normative statements (i.e. complete systems of norms n). Taken together, a consistent evaluation of all factual and normative statements is a *factual-normative world* $\langle w, n \rangle$. The semantic content (or intension) of a statement S is given as the set of factual-normative worlds in which it holds – as the set of all $\langle w, n \rangle$ relative to which S

counts as minimalistically-true. In possible world semantics, a statement S is correspondence-true (i.e. TRUE) iff S holds in the actual world $w_@$, and correspondence-false (i.e. FALSE) iff $\neg S$ holds in $w_@$. Analogously, in Gibbard's (1990) system, a normative statement N would be TRUE iff N holds in all factual-normative worlds indexed to the actual world (i.e. in $\langle w_@, n \rangle$, for all n), and FALSE iff $\neg N$ holds in all factual-normative worlds indexed to the actual world (i.e. in $\langle w_@, n \rangle$, for all n). On Gibbard's (1990) view, however, there are no normative facts. In other words, for any normative statement N , neither N nor $\neg N$ holds in all factual-normative worlds indexed to the actual world (i.e. in $\langle w_@, n \rangle$, for all n). It follows that every normative statement is neither TRUE nor FALSE. By contrast, each factual statement is either TRUE or FALSE, since whether any factual statement holds in all factual-normative worlds indexed to the actual world (i.e. in $\langle w_@, n \rangle$, for all n) depends only on whether it holds in $w_@$. By definition, $w_@$ is an evaluation of all factual statements; each factual statement either holds in $w_@$ or its negation does. Therefore, factual statements are TRUE iff they hold in the actual world $w_@$ – and FALSE otherwise.

4.3.2 Horgan & Timmons's *Cognitivist Expressivism*

With Gibbard's (1990) norm-expressivism in view, I proceed now to review Horgan & Timmons's (2000, 2006) cognitivist expressivism. In “Nondescriptivist Cognitivism” (2000), they defend the coherence of nondescriptive beliefs, but largely refrain from proposing a formal language. Horgan and Timmons introduce one later in “Cognitivist Expressivism” (2006), where they refine and extend their earlier work. Unlike Gibbard (1990), Horgan & Timmons (2000, 2006) directly theorize about moral statements – or, at the very least, about moral *ought* statements of the form ‘it ought to be that p .’ They hope to account for other kinds of moral statements (e.g. ‘donating to charity is permissible’ and ‘stealing is wrong’) derivatively, on the basis of their account of ought statements.

Horgan & Timmons (2000) challenge what they call the *semantic assumption* (SA) – the assumption that “all genuinely cognitive content is descriptive content.” (p.124) It follows from this assumption that if moral beliefs and sentences have cognitive content (which the grammatical problem indicates they do), that they are necessarily descriptive – that they are “in the business of representing some (putative) state of affairs or stating some (putative) fact.”

(p.124) Denying the semantic assumption makes conceptual space for Horgan & Timmons (2000) to say that moral judgments are genuinely cognitive (i.e. belief-eligible and assertible) *without* being descriptive.

The challenge involved in denying the semantic assumption is to explain what it is for beliefs and judgments to be cognitive, if not for them to be descriptive. Horgan & Timmons (2000) answer this challenge by appealing to minimalism. In their view, beliefs count as minimalistically cognitive – and hence, minimalistically *as* genuine beliefs – by dint of the facts that they (i) are commitment-states to propositions, which (ii) have a “distinctive *constitutive inferential role* in an agent’s cognitive economy[.] … In particular, their role is to combine in a distinctive way with other beliefs (other commitment-states) to inferentially yield further beliefs (further commitment-states).” (pp.136-137) This minimalistic conception of belief permits Horgan & Timmons (2000) to distinguish between two different kinds of genuine beliefs by distinguishing two different kinds of commitment-states to propositions: *is-commitments* and *ought-commitments*.

According to Horgan & Timmons’s (2000) semantic framework, descriptive beliefs are *is-commitments* to so-called “core descriptive contents” (i.e. propositions, or ways the world might be), whereas moral beliefs are *ought-commitments* to “core descriptive contents.”¹³ Descriptive and moral beliefs both qualify minimalistically as *genuine* beliefs, since both are commitment-states with respect to core descriptive contents (i.e. propositions), and both have the aforementioned “constitutive inferential role” of beliefs. Similarly, both descriptive and moral assertions qualify minimalistically as *genuine* assertions, since both are “*stance-taking* speech act[s]” used to express genuine beliefs. (p.135) In brief, since moral beliefs minimalistically qualify as genuine beliefs, and speech-acts used to express genuine beliefs minimalistically count as genuine assertions, it trivially follows that moral assertions are genuine assertions.

¹³ Do not be alarmed by the fact that Horgan & Timmons (2000) say that *nondescriptive* beliefs (and sentences) have “core *descriptive* content,” as overall content (i.e. content *simpliciter*) and “core descriptive content” refer to different things. When Horgan & Timmons (2000) say that moral beliefs have “core descriptive content,” all they are saying is that moral beliefs are ought-commitments to a particular way the world might be – where the “core descriptive content” supplies the (desired) way for the world to be. For example, suppose I believe that everyone should be nice to Tim. The “core descriptive content” of my belief is *that everyone is nice to Tim*. The overall content is nondescriptive: *that it ought to be that everyone is nice to Tim*.

Although Horgan & Timmons (2000) do not propose a formal language or truth-theoretic apparatus, they do advocate for a particular kind of truth pluralism (minimalist-truth for moral discourse and correspondence-truth for descriptive discourse) which constrains the kind of view they intend – and later defend in “Cognitivist Expressivism” (2006). As Horgan & Timmons (2000) explain,

[T]o understand truth minimalistically in one discourse does not commit one to minimalism in relation to every mode of discourse. ... Thus, we are inclined to advocate a kind of pluralism about truth according to which there is a univocal notion of truth even though truth ascriptions may involve more or less robust metaphysical commitments in relation to different areas of thought and discourse. (pp.146-147)

In an earlier footnote, Horgan & Timmons (2000) clarify that according to their framework, moral discourse involves a “morally engaged” use of the truth predicate, whereas ordinary descriptive discourse employs the more metaphysically robust, correspondence-truth notion:

[M]any terms—including the truth predicate—are subject to contextually variable semantic standards. In the case of moral thought and discourse, which is nondescriptive in overall declarative content, typically the contextually operative semantic standards governing the truth predicate dictate a morally engaged use—the use we have just explained. But in some contexts the semantic standards dictate a morally detached use of the truth predicate, under which ‘true’ signals language-world correspondence; on this usage, only statements whose overall declarative content is descriptive are either true or false. (p.146)

Horgan & Timmons (2000) explain the “morally engaged use” of the truth predicate as follows:

When one thinks or remarks, ‘The claim that slavery ought to be stopped is true’, what is one doing? The appropriate answer involves noting that such a truth ascription constitutes a *morally engaged* semantic appraisal: one that is infused with one’s own moral commitment. The main idea can perhaps be conveyed by saying that truth ascriptions to moral statements involve a kind of appraisal in which semantic and moral are ‘fused’—which is to be expected, since ordinary uses of the truth predicate operate in accordance with schema T [i.e. the schema according to which ‘p’ is true iff p]. (p.146)

In brief, Horgan & Timmons (2000) envision truth ascriptions to moral sentences *M* as consistent with the minimalist equivalence schema, and therefore “morally engaged” in just the same way that sincere utterances of *M* are. In either case, one’s assertion is “infused with one’s own moral commitment.” Horgan & Timmons (2000) do not have the terminology of the genuine relativists at their disposal, but it is fairly clear what they intend: for a minimalist truth ascription to a moral statement to be “infused with one’s own moral commitment” is for one to assess that moral statement as true from within – that is, *relative* to – their own committed moral outlook.

In review, Horgan & Timmons's (2000) framework formally distinguishes between ordinary descriptive and moral discourse. Ordinary descriptive sentences are said to have descriptive content and express *is-commitments*, while moral sentences have nondescriptive content and express *ought-commitments*. Both sorts of sentences are truth-evaluable – but whereas truth ascriptions in descriptive discourse involve the correspondence notion of truth (i.e. TRUTH), truth ascriptions in moral discourse involve a “morally engaged,” minimalist notion of truth.

Horgan & Timmons's (2000) framework is largely compatible with Gibbard's (1990) formal, truth-theoretic apparatus. By modeling the set of descriptive beliefs which a person accepts as D , and the set of moral beliefs they accept as M , the set of a person's descriptive and moral beliefs can be modeled as the ordered pair $\langle D, M \rangle$ – where this is analogous to Gibbard's (1990) $\langle B, N \rangle$ (the set of a person's factual and normative beliefs). What remains to be seen is how Horgan & Timmons (2000) might truth-theoretically represent the semantic contents (or intensions) of descriptive and moral sentences, and their logical consequence relations.

Thankfully, they gesture at the answer in yet another clarifying footnote:

[T]here is an intelligible notion of logical consequence that applies to beliefs and assertions whether or not their overall declarative content is descriptive. [I]t is surely plausible that this is so, i.e., that logic governs psychological commitments and sociolinguistic stances, even those with nondescriptive declarative content. ... If truth-value assignments are extended to encompass base-case ought-statements in addition to atomic statements, then the resulting truth-theoretic account of logical truth and of the logical-consequence relation can be interpreted (i) as applicable to morally engaged uses of the truth-predicate, hence (ii) as also applicable to morally engaged thought and discourse (which conforms to schema T), hence (iii) as applicable to beliefs and assertions even when their overall declarative content is not descriptive. Moreover, presumably the possible-world semantics of deontic logic could be smoothly incorporated into such a truth-theoretic approach, with possible worlds construed as specifiable by Carnap-style ‘state descriptions’—maximal consistent sets of atomic statements and negations of atomic statements. (p.149)

A maximal consistent set of descriptive sentences could be modeled as d , and a maximal consistent set of moral sentences as m . A consistent assignment of truth-values to all descriptive and moral sentences could then be modeled as the ordered pair $\langle d, m \rangle$: call this a *descriptive-moral valuation*. The parallels here to Gibbard's (1990) formal apparatus are obvious: just as factual-normative worlds $\langle w, n \rangle$ are consistent completions of actual people's credal-normative states $\langle B, N \rangle$, descriptive-moral valuations $\langle d, m \rangle$ are consistent completions of actual people's descriptive-moral belief states $\langle D, M \rangle$.

Before drawing any overly speculative conclusions, let's turn to Horgan & Timmons's (2006) formal apparatus. In the formal language they develop, a nonsentential formula (*nsf*) is a syntactic representation of what they have so far been calling a "core descriptive content" (i.e. an ordinary descriptive that-clause, or proposition). Sentences themselves are defined as the syntactic items produced by slotting nsfs into the bracketed slot(s) of a sentential-formula forming operator (*sff operator*).¹⁴ The two basic sff operators are **I**[] and **O**[], which express *is-commitment* and *ought-commitment* to the nsf inserted into their bracketed slot, respectively. Logically complex constructions of sff operators are also sff operators. For example, where C is the nsf 'that charities fund organized crime,' and D is the nsf 'that John donates to charity,' $\neg\mathbf{I}[C] \rightarrow \mathbf{O}[D]$ represents the complex sentence 'if it is not the case that charities fund organized crime, then it ought to be the case that John donates to charity.'

Proceeding to semantics, Horgan & Timmons (2006) introduce a *valuation* as "an assignment of the truth-values T and F to some (but not necessarily all) of the sentences and closed nonsentential formulas of the formal language[.]" (p.280) More precisely, a valuation **V** is defined as a pair $\langle \mathbf{N}, \mathbf{S} \rangle$, where **N** is a valuation of at least some (but not necessarily all) nsfs, and **S** is a valuation of at least some (but not necessarily all) sentences. I will consider each in turn, beginning with **N**.

A valuation of nsfs **N**, can assign T or F to any subset of nsfs, provided that (i) no nsf is assigned both T and F, and (ii) every logically related nsf is consistently evaluated. For example, iff **N** assigns T to an nsf A, it must assign F to $\sim A$; iff it assigns T to a disjunction of nsfs $A \vee B$, it must assign T to A, or T to B, and so on. **N** may assign *neither* T nor F to an nsf, reflecting either *indifference* (i.e. uncertainty) or *indecision* (i.e. incompleteness). For instance, **N** may assign neither T nor F to 'that there are an even number of stars in the universe' on the grounds of *indifference*, and to 'that the moon is larger than every unripe pineapple harvested in 1823' on the grounds of *indecision*.

¹⁴ I am simplifying for the sake of brevity. To be exact, a *sentence* in Horgan & Timmons's (2006) formal language is a *sentential formula* containing no free variables. A *sentential formula* is defined as the result of "inserting the respective elements of [a sequence of n variables] into the respective left-to-right quantificational slots in [an sff operator Ω with n quantificational slots] and inserting the respective elements of [a sequence of m nsfs] into the respective left-to-right bracketed slots in [Ω 's m bracketed slots]." (p.289)

The second component of a valuation pair \mathbf{V} ($\langle \mathbf{N}, \mathbf{S} \rangle$) is \mathbf{S} , a valuation of sentences. A valuation of sentences \mathbf{S} is an assignment of T or F to sentences – that is, to sff operators ($\mathbf{I}[]$, $\mathbf{O}[]$, or some logical complex) with nsfs inserted into their bracketed slots. A valuation of sentences \mathbf{S} is constrained by the valuation of nsfs \mathbf{N} with which it is paired in \mathbf{V} . This constraint is most pronounced for is-sentences: for any nsf A, \mathbf{S} assigns the same truth-value to $\mathbf{I}[A]$ that \mathbf{N} assigns to A. By contrast, \mathbf{S} is relatively free to assign truth-values to ought-sentences, with two key constraints: (i) if \mathbf{S} assigns T to $\mathbf{O}[A]$, it must also assign T to $\mathbf{O}[B]$ for every nsf B logically entailed by A, and (ii) if \mathbf{S} assigns F to $\mathbf{O}[A]$, it must also assign F to $\mathbf{O}[B]$ for every nsf B that logically entails A. Like \mathbf{N} , \mathbf{S} may assign neither T nor F to sentences due to indifference or indecision. Finally, \mathbf{S} obeys the same consistency constraints as \mathbf{N} : no sentence is assigned both T and F, and all logically related sentences are consistently evaluated (e.g. \mathbf{S} assigns T to sentence \mathcal{Q} iff it assigns F to $\neg\mathcal{Q}$).

Thus, a valuation \mathbf{V} ($\langle \mathbf{N}, \mathbf{S} \rangle$) consistently assigns truth-values to some subset of nsfs (\mathbf{N}) and sentences (\mathbf{S}). Is-sentences receive the same truth-values as their embedded nsfs, while ought-sentences have looser constraints. If \mathbf{V} assigns T to $\mathbf{O}[A]$, it must also assign T to every ought-sentence whose nsf is logically entailed by A. Likewise, if \mathbf{V} assigns F to $\mathbf{O}[B]$, it must also assign F to every ought-sentence whose nsf logically entails B. Roughly, if A ought to be, everything logically entailed by A ought to be. And if B ought not to be, everything that entails B also ought not to be.

Actual belief states modeled as valuations \mathbf{V} are *partial* – they assign neither T nor F to at least some nsfs and sentences. A valuation \mathbf{V} is *complete* if it assigns T or F to every nsf and sentence. Complete valuations \mathbf{V} fully specify *what is the case* and *what ought to be the case*. Horgan & Timmons (2006) distinguish descriptive from moral sentences according to logical reducibility to a single is-commitment sff operator $\mathbf{I}[]$ with an arbitrarily complex nsf inserted into its bracketed slot. As Horgan & Timmons (2006) write,

[I]f B is descriptive ... then B is logically equivalent to some belief B^* that is a logically simple is-commitment (and thus is expressible by a descriptive sentence constructed by inserting some closed nonsentential formula into a single occurrence of the operator ‘ $\mathbf{I}[]$ ’). .. [I]f B is not descriptive (i.e. is not expressible by a descriptive sentence), then B is not logically equivalent to any is-commitment. (p.295)

Descriptive sentences (and beliefs) are logically reducible to a single occurrence of $\mathbf{I}[]$ with some nsf inserted into its bracketed slot, while moral sentences (and beliefs) must contain at least one occurrence of the ought-commitment sff operator $\mathbf{O}[]$ – which prevents reducibility to a single occurrence of $\mathbf{I}[]$.

This guarantees that any valuation \mathbf{V} is separable into a consistent evaluation of all nsfs and descriptive sentences \mathbf{d} and a consistent evaluation of all moral sentences \mathbf{m} . Complete valuations \mathbf{V} are thus identical with descriptive-moral valuations $\langle \mathbf{d}, \mathbf{m} \rangle$, isomorphic with Gibbard's (1990) factual-normative worlds $\langle \mathbf{w}, \mathbf{n} \rangle$.

Since sentences in Horgan & Timmons's (2006) formal apparatus are minimalistically-true relative to any valuation \mathbf{V} that assigns them T, the *semantic content* (or *intension*) of any given sentence \mathcal{Q} (i.e. the complete set of conditions in which \mathcal{Q} is considered true) is expressible by the set of all complete valuations \mathbf{V} (themselves expressible as descriptive-moral valuations $\langle \mathbf{d}, \mathbf{m} \rangle$) that assign T to \mathcal{Q} . This closely parallels Gibbard's (1990) framework, where sentence contents are sets of factual-normative worlds $\langle \mathbf{w}, \mathbf{n} \rangle$.

All that remains is to show that Horgan & Timmons's (2006) apparatus assigns correspondence truth-values (i.e. TRUE / FALSE) to descriptive sentences without assigning them to moral sentences. As they write,

[Although a valuation generally] reflects a given agent's nonmoral and moral beliefs ... using the truth predicate in a morally engaged [i.e. minimalist] way ... there will also be a valuation that assigns truth values in accordance with a morally detached correspondence-usage of 'true', and that furthermore assigns truth and falsity based on whether or not a given sentence actually corresponds to how things are or not (rather than on the basis of any particular agent's *beliefs*). In such a valuation, all sentences of the form $\mathbf{O}[A]$ will be assigned neither T nor F. (p.280)

Let's call the special valuation, which assigns truth and falsity based on correspondence to reality, $\mathbf{V}_@$. The key idea is that a sentence \mathcal{Q} is correspondence-true (i.e. TRUE) just in case \mathcal{Q} is minimalistically-true relative to $\mathbf{V}_@$ (i.e. iff $\mathbf{V}_@$ assigns T to \mathcal{Q}). Equivalently, a sentence \mathcal{Q} is TRUE iff every consistent completion of $\mathbf{V}_@$ is an element of the intension of \mathcal{Q} , FALSE iff every consistent completion of $\mathbf{V}_@$ is an element of the intension of $\neg\mathcal{Q}$, and neither TRUE nor FALSE otherwise.

Horgan & Timmons's (2006) moral irrealism is synonymous with their claim that the correspondence-truth valuation $V_{@}$ assigns neither T nor F to any sentence of the form $O[A]$. In other words, Horgan & Timmons (2006) maintain that $V_{@}$ is a complete truth-assignment to all descriptive sentences but to no moral sentences. This corresponds to the view that $V_{@}$ is identical with a special valuation of all and only descriptive sentences d – namely, $d_{@}$. If $V_{@}$ assigns neither T nor F to any ought-sentence, then every moral sentence is neither TRUE nor FALSE. By contrast, since $V_{@}$ assigns either T or F to every descriptive sentence, each descriptive sentence is either TRUE or FALSE.

In summary, Horgan & Timmons's (2006) formal semantic apparatus defines valuations V that consistently assign truth-values to nsfs and sentences. They distinguish between descriptive and moral sentences. The former are logically reducible to a single is-commitment operator $I[]$, whereas the latter are not reducible in this way, since they contain at least one ought-commitment operator $O[]$. Complete valuations V are therefore expressible as consistent evaluations of all descriptive and moral sentences – as *descriptive-moral valuations* $\langle d, m \rangle$. The semantic content (or intension) of a sentence – the complete set of conditions in which it is true – is therefore the set of descriptive-moral valuations relative to which that sentence is minimalistically-true. Correspondence-truth can then be defined as minimalistic-truth relative to a special valuation – the actual valuation $V_{@}$. A sentence Ω is correspondence-true (i.e. TRUE) iff $V_{@}$ assigns T to Ω , and correspondence-false (i.e. FALSE) iff $V_{@}$ assigns F to Ω . According to Horgan & Timmons's (2006) moral irrealism, $V_{@}$ assigns either T or F to every descriptive sentence, and neither T nor F to any moral sentence. It follows that descriptive sentences are either TRUE or FALSE, and that moral sentences are neither TRUE nor FALSE.

4.3.3 Kölbel's *Perspective-Relativism*

I proceed now to the formal semantics of the genuine relativists, beginning with Kölbel (2002). Where Gibbard (1990) and Horgan & Timmons (2000, 2006) aim to model normative and moral judgments, Kölbel (2002) attempts to account for taste judgments. As discussed in Sections 3.5.1 and 4.1, Kölbel's (2002) primary motivation is to make sense of how faultless propositional disagreements are possible. Pursuant to this goal, he relativizes the truth of sentence contents to *perspectives*. Like Gibbard's (1990) factual-normative worlds and Horgan &

Timmons's (2006) complete valuations, Kölbel's (2002) perspectives are consistent truth-value assignments to all sentence contents. Unlike these constructs, however, perspectives are intrinsically tied to individuals – each thinker is said to “possess” a perspective at any given time.

Although Kölbel (2002) is intentionally vague about what it means for a person to possess a perspective, he is fairly clear about what perspectives are, and what their possession is *not*. First, perspectives are not mental states analogous to beliefs, and so they are not possessed in the way that people possess beliefs. Nor are perspectives differing “vantage points,” or clashing “points of view” on the same objective world. A perspective is rather more like a possible world in which a person lives – in that they both specify the complete set of sentences which a person *should* believe on the threat of making a mistake. The *possession* relation, then, is largely analogous to the relation between a person and the possible world which they inhabit. To be sure, possible worlds are distinct from perspectives in a couple important respects, not least of which is that “two individuals who are causally related to one another must inhabit the same world. By contrast, two thinkers may be causally related yet possess different perspectives.” (p.101) The other essential difference is that “ordinary possible worlds are purely objective and contain no evaluative elements[,]” whereas perspectives assign truth-values to objective (i.e. descriptive) and non-objective (i.e. evaluative) sentences, alike. (p.101)

Kölbel (2002) further clarifies what perspectives are – and what differentiates objective from non-objective sentence contents (or propositions) across three constraints that he places on perspectives:

(C1) For all [contents] p : p is non-objective iff it is possible that there be thinkers A and B , such that p is true in A 's perspective and p is not true in B 's perspective. ...

(C2) For all [contents] p : p is objective iff it is not possible that there be thinkers A and B , such that p is true in A 's perspective and p is not true in B 's perspective. ...

(C3) For all thinkers t and all perspectives s : t possesses s iff for all contents p : t ought to believe p only if p is true in s . (pp. 102-103)

The first two constraints neatly bifurcate the set of all sentence contents into those which are *objective*, and those which are *non-objective*. This guarantees that any given perspective \mathbf{P} is

expressible as an ordered pair $\langle o, n-o \rangle$, where o is a consistent truth-assignment to all objective sentence contents, and $n-o$ is a consistent truth-assignment to all non-objective sentence contents.

From Kölbel's (2002) constraints on perspectives, we can also gather that an element of a perspective (i.e. a proposition) is *objective* just in case it is an element of an ordinary possible world. This follows from the second constraint – that an element of a perspective is objective iff it necessarily has the same truth-value in the perspective of every person (inhabiting the same world) – together with the fact that the possible world which a person inhabits is identical with the set of truth-assignments which are necessarily the same for everyone inhabiting that world. Hence, the objective component o of any perspective can be expressed as an ordinary possible world w . Since by Kölbel's (2002) own characterization, an ordinary possible world w is a consistent truth-assignment to all factual, *non-evaluative* sentence contents, it follows that the non-objective component $n-o$ is expressible as e – a consistent truth-assignment to all evaluative sentence contents. Therefore, perspectives \mathbf{P} are expressible as ordered pairs $\langle w, e \rangle$ – *factual-evaluative worlds*. Expressed as factual-evaluative worlds $\langle w, e \rangle$, Kölbel's (2002) perspectives are comparable to Gibbard's (1990) factual-normative worlds $\langle w, n \rangle$ and Horgan & Timmons's (2006) descriptive-moral valuations $\langle d, m \rangle$.

As mentioned earlier, Kölbel (2002) explicitly relativizes the truth-values of sentence contents to perspectives. (p.100) His primary motivation for doing so is to permit the possibility of faultless propositional disagreements. At a glance, Kölbel (2002) accomplishes this by relativizing the notions of truth and falsity at the heart of norms about truth and falsity like *Faultful Falsehood (FF)*:

Faultful Falsehood (FF): If a person believes (or asserts a sentence expressing) a false proposition, then they have made a mistake (or are at fault).

*Relativized Faultful Falsehood (RFF)*¹⁵: If a person believes (or asserts a sentence expressing) a proposition which is false according to the perspective that they possess (at that time), then they have made a mistake (or are at fault).

So long as a proposition p is non-objective (i.e. free to vary in truth-value between perspectives possessed by different people inhabiting the same world), then it is possible for the parties to a

¹⁵ **(RFF)** is logically entailed by Kölbel's (2002) third constraint (C3) on perspectives. Whereas **(RFF)** states (in effect) that a person *ought not* believe propositions which are false in their perspective, (C3) states that a person *ought only* to believe propositions which are true in their perspective.

disagreement over p to assert (and believe) contradictory statements (' p ' and ' $\sim p$ ') without either party making a mistake (by the lights of (RFF), at any rate). Thus occurs just in case the truth-value of p differs between the perspectives possessed by disputants, and each disputant asserts (and believes) a statement (' p ' or ' $\sim p$ ') which is true according to their own perspective.

Armed with this deflationary and explicitly relativized notion of truth, Kölbel (2002) proceeds to define a narrower notion of objective, or correspondence-truth (i.e. TRUTH) in terms of the relative notion. Kölbel (2002) writes,

We could even define a notion of absolute truth in terms of the notion of relative truth: a content is *objectively true* just if it is objective and true in everyone's (or equivalently: someone's) perspective. Similarly, we can define objective falsehood: a proposition is *objectively false* just if it is objective and false in everyone's (or equivalently: someone's) perspective. Objective truth and objective falsehood are not bivalent: some contents are neither objectively true nor objectively false. (p.102)

Kölbel (2002) use of 'everyone's perspective' is somewhat hasty here – it would be more accurate to say "the perspective of everyone inhabiting the same world." This is because even objective propositions (e.g. *the Moon orbits the Earth*) can be true according to a perspective possessed by someone inhabiting some possible world (say, w_1) and yet false according to a perspective possessed by someone inhabiting a different possible world (say, w_2). In that light, a proposition is *objectively true* in a possible world w just in case it is true and objective in the perspective of everyone inhabiting w . Correspondence-truth (i.e. TRUTH) follows as a special case of *objective truth* – as objective truth in the actual world $w_@$.¹⁶ To be exact, a proposition p is TRUE (FALSE) according to Kölbel's (2002) semantic framework iff it is true (false) in all perspectives possessed in $w_@$, and it is not possible for the perspectives possessed in $w_@$ to differ with respect to the truth-value they assign p .

As we have already seen within Kölbel's (2002) semantic framework, a perspective \mathbf{P} is expressible as an ordered pair $\langle w, e \rangle$, where w assigns either 'true' or 'false' to all non-evaluative (i.e. objective) sentence contents, and e assigns either 'true' or 'false' to all

¹⁶ An alternative reading of Kölbel's (2002) passage on objective truth reads 'everyone's perspective' as "every *actual* person's perspective." On this alternative reading, *objective truth* is immediately identified as the status of being true and objective in the actual world $w_@$.

On either reading, an important status is carved out: the status of being true and objective in $w_@$. Whether we call this status "objective truth" or "correspondence-truth" is not important.

evaluative (i.e. non-objective) sentence contents. Since (i) p is TRUE (FALSE) iff p is necessarily true (false) in all perspectives possessed by people in $w_@$, and (ii) all perspectives possessed by people in $w_@$ necessarily agree with respect to the set of their objective truth-assignments – itself expressible as the world which those people inhabit, $w_@$ – it follows that (iii) p is TRUE (FALSE) iff p is true (false) in $w_@$. Like all ordinary possible worlds, $w_@$ is a complete truth-assignment to all and only *non-evaluative* sentence contents. It follows that $w_@$ assigns neither true nor false to evaluative sentence contents, and therefore, that evaluative contents are neither TRUE nor FALSE. According to Kölbel's (2002) semantic framework, then, each non-evaluative proposition is either TRUE or FALSE, whereas all evaluative propositions are neither TRUE nor FALSE.

Finally, Kölbel's (2002) semantic framework yields the semantic content (i.e. intension) of each sentence S as the set of all perspectives \mathbf{P} which evaluate S as true. This follows straightforwardly from Kölbel's (2002) formal semantic apparatus, since the complete set of conditions in which a sentence S is considered true is given by the set of all perspectives \mathbf{P} which evaluate S as true. Kölbel (2002) confirms this interpretation when he compares his semantic framework to Gibbard's (1990):

There are obvious analogies between Gibbard's factual-normative worlds and my perspectives. Just as a perspective is an evaluation of all contents, objective or not, a factual-normative world is an evaluation of all contents, normative or factual. Just as Gibbard's contents can be construed as sets of factual-normative worlds, mine could be construed as sets of perspectives. In each case the set contains those perspectives or factual-normative worlds that evaluate the content as true. (Kölbel, 2002, pp.111-112)

In summary, Kölbel's (2002) formal semantic apparatus defines a perspective \mathbf{P} as a consistent assignment of truth-values to all sentence contents. Additionally, any given person is said to *possess* (or inhabit) a given perspective (at any given time), where the perspective a person possesses corresponds to the set of sentences they ought to believe (on the threat of making a mistake). Kölbel (2002) then distinguishes between objective and non-objective sentence contents. Whereas objective contents necessarily share the same truth-values in all perspectives (possessed in the same world), the truth-values assigned to non-objective contents are free to vary between perspectives (possessed in the same world). Perspectives \mathbf{P} are therefore expressible as consistent evaluations of all objective and non-objective contents $\langle o, n\text{-}o \rangle$. In Kölbel's (2002) view, all objective contents are factual (such that o is expressible as an ordinary

possible world w) and non-evaluative. Therefore, perspectives P are also expressible as consistent evaluations of all factual and evaluative contents $\langle w, e \rangle$ – as *factual-evaluative worlds*. Kölbel's (2002) basic notion of truth is minimalist in character and explicitly relativized to perspectives: truth in a perspective. He then defines *objective truth (falsity)* in a given world as necessarily being true (false) in all perspectives possessed in that world. Correspondence-truth and correspondence-falsity (i.e. TRUTH and FALSITY) follow as special cases of objective truth and falsity – as objective truth and falsity in the actual world $w_@$. Since $w_@$ assigns truth-values to all and only factual, non-evaluative contents, it follows that each non-evaluative content is either TRUE or FALSE, whereas every evaluative content is neither TRUE nor FALSE. Finally, the semantic content (or intension) of sentences in Kölbel's (2002) formal semantic apparatus can be construed as sets of perspectives (or *factual-evaluative worlds* $\langle w, e \rangle$).

4.3.4 Lasersohn's *Assessment-Relativism*

I proceed, finally, to Lasersohn's (2005) formal semantic apparatus. Lasersohn (2005) follows Kölbel (2002) in attempting to model the semantics of taste judgments – unlike Gibbard (1990) and Horgan & Timmons (2000, 2006), who aim at normative and moral judgments, respectively. Lasersohn's (2005) framework is relatively unique among these four, in that he does not directly invoke complete truth-assignments to sentence contents. This is a consequence of the fact that Lasersohn (2005) explicitly extends a form of Kaplanian semantics (as in Kaplan (1977)). Rather than directly relativize truth-values to valuations, Lasersohn (2005) instead relativizes the truth-values of sentence contents p to index triples $\langle w, t, j \rangle$, consisting of a world w , a time t , and a judge j . For the sake of brevity, I will eschew the time component,¹⁷ and refer to the resulting ordered pair $\langle w, j \rangle$ as a *context of assessment* – the term which Lasersohn (2017) later adopts in his full-length book revisiting the subject. For further conciseness, I will often represent contexts of assessment simply as a .

¹⁷ Lasersohn (2005) originally includes the time parameter merely to allow for his apparatus to account for Kaplan-style tense operators. In his later account, Lasersohn (2017) forgoes the need for such tense operators, and relegates the time parameter to the role of tracking how each judge's tastes change across time. By treating the same person at different times as different judges, however, we can eschew the need for the time parameter altogether.

In Lasersohn's (2005) formal semantic framework, $\models_{f,a} \varphi$ represents the claim that sentence φ is true relative to context of assessment a , in formalistic context f . Similarly, $\not\models_{f,a} \varphi$ represents the claim that φ is false relative to a in f . Lasersohn's (2005) formalistic context f is nearly identical with what Lasersohn (2017) later calls the *context of use*: f encodes a world w , a time t , and a speaker u ,¹⁸ which together with (the Kaplanian character of) any sentence φ determines a unique sentence content p . Accordingly, Lasersohn (2005) represents the content p of any sentence φ in formalistic context f as $\{\varphi\}_f$.

Since Lasersohn (2005) relativizes the truth-values of sentence contents to contexts of assessment a , and any sentence φ and formalistic context f determine a unique sentence content p , I will write $\models_a p$ as shorthand for the claim that content p is true relative to context of assessment a (and $\not\models_a p$ for the claim that p is false relative to a).¹⁹

In Lasersohn's (2005) formal semantic apparatus, the content of any given sentence is a function from contexts of assessment to truth-values. Lasersohn (2005) further clarifies that $\{\varphi\}_f$ (the content of any given sentence φ in formalistic context f) is the function that assigns Truth to each context of assessment a if $\models_{f,a} \varphi$, and Falsehood otherwise, if $\not\models_{f,a} \varphi$. (p.666) In short, any given sentence content $\{\varphi\}_f = p$ is the function that assigns, for each context of assessment a , Truth to a if $\models_a p$ and Falsehood to a if $\not\models_a p$. That is, for all contents p and contexts of assessment a , p is that function that assigns Truth to a if p is true relative to a , and Falsehood to a if p is false relative to a . This guarantees that for all sentence contents p and contexts of assessment a , either p is true relative to a , or p is false relative to a .

I am now in a position to show how Lasersohn's (2005) contexts of assessment a are equivalent to (and expressible as) a set of consistent truth-assignments to all sentence contents \mathbf{V}_a . As just discussed, Lasersohn's (2005) framework requires that for any given context of

¹⁸ To be sure, f encodes *more* than the context of use does. In addition to the three parameters (world, time, and speaker) needed (in addition to a sentence φ) to determine a unique sentence content p , f also encodes a judge parameter j . The judge parameter in f plays no role in determining the content $\{\varphi\}_f$ of sentence φ in f , and serves only to help Lasersohn (2005) define the notion of "truth in a context f ." Since this notion is parasitic on his more basic notion of truth relative to a context of assessment ($\models_{f,a} \varphi$), it is not directly relevant to my current project.

¹⁹ The purpose of this shorthand is just to be able to express $\models_{f,a} \varphi$ even more concisely as $\models_a p$ (so long as p is the content of sentence φ in formalistic context f).

assessment a and sentence content p , p is either true relative to a (if $\models_a p$), or else false relative to a (if $\not\models_a p$). This guarantees that any given context of assessment a assigns a truth-value to every sentence content p . In other words, since Lasersohn's (2005) sentence contents are total functions from contexts of assessment to truth-values, it follows that his contexts of assessment are total functions from sentence contents to truth-values. Therefore, any given context of assessment a in Lasersohn's (2005) semantic framework is expressible as consistent truth-assignment to all sentence contents, \mathbf{V}_a .

As we have already seen, Lasersohn (2005) defines the semantic content $\{\varphi\}_f$ (or *intension*) of a sentence as the set of all contexts of assessment a relative to which $\{\varphi\}_f$ is true. Lasersohn (2005) expresses this point casually when he writes that “[i]nstead of treating the content of a sentence as a set of time-world pairs [as in Kaplan's (1977) original language], we should treat it as a set of time-world-individual triples [i.e. contexts of assessment a].” (p.663) Since contexts of assessment a are expressible as consistent truth-assignments to all sentence contents \mathbf{V}_a , it follows that the semantic content of a sentence (or intension) in Lasersohn's (2005) semantic framework is equivalent to the set of all complete truth-assignments \mathbf{V}_a which assign truth to it.

Now, recall that contexts of assessment a are pairs of the form $\langle w, j \rangle$. In Lasersohn's (2005) semantic framework, the set of all sentence contents cleanly divides into contents whose truth-values can vary with the judge parameter j , and those which cannot. Lasersohn (2005) calls the former *subjective*, and the latter *objective*. By definition, if a subjective content is true relative to $\langle w_1, j_1 \rangle$, it may be false relative to $\langle w_1, j_2 \rangle$, whereas an objective content is guaranteed to assign the same truth-value to both contexts of assessment (since they share the same world). Where o is a consistent truth-assignment to all objective contents, and s is a consistent truth-assignment to all subjective contents, a consistent truth-assignment of all sentence contents \mathbf{V}_a (itself equivalent to context of assessment a) is therefore expressible as an ordered pair $\langle o, s \rangle$: an *objective-subjective assessment*.

Since the truth-values of objective contents do not depend on the judge parameter j , they vary only with the world w of the context of assessment. For short, we may say that contents which are objective and true (false) relative to a world are *objectively true (objectively false)*

relative to that world. Correspondence-truth and correspondence-falsity (i.e. TRUTH and FALSITY) can then be defined as objective truth and objective falsity relative to the actual world $w_{@}$. Indeed, this is precisely what Lasersohn (2017) does in his later work. As he writes,

We can also define monadic (or “absolute”) truth for sentence contents by universally quantifying on the perspective indices, and fixing the modal evaluation world index to the actual world $w_{@}$:

(141) a. Φ is [absolutely] true iff for all [contexts of assessment a]: $\Phi(w_{@}, [a]) = \text{truth}$;
b. Φ is [absolutely] false iff for all [contexts of assessment a]: $\Phi(w_{@}, [a]) = \text{falsity}$;
(p.102)

For context, ‘for all contexts of assessment a : $\Phi(w_{@}, a) = \text{truth}$ ’ states that sentence content Φ is true relative to all contexts of assessment a in the actual world $w_{@}$. Hence, a sentence content is absolutely true (i.e. TRUE), according to Lasersohn (2017), iff it is true relative to every context of assessment a whose world component is the actual world $w_{@}$. Equivalently, a sentence content is TRUE (FALSE) iff it is objectively true (objectively false) in the actual world $w_{@}$.

Since subjective contents (i.e. contents whose truth-value can vary with j) can neither be objectively true in the actual world $w_{@}$ (or in any world), nor objectively false in the actual world $w_{@}$, it follows that every subjective content is neither TRUE nor FALSE. As Lasersohn (2017) writes,

[I]n cases where an expression deals with matters of opinion, the relativization of the denotation assignment will be ineliminable, so that no absolute truth values can be assigned, only relative, parameter-dependent, truth values. Because possible worlds are “total,” resolving all matters of fact, ineliminable variation in truth value according to non-world parameters amounts to dependency of truth value on non-factual matters. In this case the content expressed by a sentence is neither true nor false *tout court*, but only relative to a particular way of resolving parameters [(i.e. the judge parameter, j)] which the facts of the world leave unresolved. (p.8)

According to Lasersohn’s (2005, 2017) formal semantics, then, each objective (i.e. matter-of-fact) content is either TRUE or FALSE, whereas every subjective (i.e. matter-of-opinion) content is neither TRUE nor FALSE.

In review, Lasersohn’s (2005) semantic framework relativizes the truth-values of sentence contents to ordered pairs $\langle w, j \rangle$, expressible as *contexts of assessment a*. The content of a sentence φ is written as $\{\varphi\}_f$ (or p) and defined as a truth-assignment to every context of assessment a . It follows that each context of assessment a assigns a truth-value to every sentence

content p . Contexts of assessment a are therefore expressible as truth-assignments to all sentence contents, \mathbf{V}_a . Since Lasersohn (2005) gives the semantic content p (or *intension*) of each sentence as the set of all contexts of assessment a relative to which p is true, it follows that the semantic content of a sentence (i.e. its *intension*) is expressible as the set of all complete truth-assignments \mathbf{V}_a which assign it truth. Lasersohn (2005) differentiates contents whose truth-values can vary with the judge parameter j from those which cannot – calling the former *subjective* and the latter *objective*. Complete truth-assignments to all sentence contents \mathbf{V}_a are therefore expressible as ordered pairs $\langle o, s \rangle$: *objective-subjective assessments*. By convention, a content in Lasersohn's (2005) framework is *objectively true (objectively false)* relative to a world iff it is objective and true (false) relative to that world. Lasersohn (2017) then defines absolute truth (i.e. TRUTH) and absolute falsity (i.e. FALSITY) as objective truth and objective falsity relative to the actual world $w_@$. Since each objective content is either true or false relative to $w_@$, it follows that each objective content is either TRUE or FALSE. But since subjective contents cannot be objectively true in the actual world (or in any world), they are neither TRUE nor FALSE according to Lasersohn's (2005, 2017) formal semantic apparatus.

4.3.5 Six key commonalities

In the previous four subsections, I reviewed the formal semantic frameworks of two forms of minimalist expressivism – Gibbard's (1990) and Horgan & Timmons's (2000, 2006) – and two forms of genuine relativism – Kölbel's (2002) and Lasersohn's (2005, 2017), and identified six features shared by all four frameworks. In this final subsection, I will summarize these six key commonalities, highlighting the striking structural parallels between the frameworks. Indeed, each framework's core commitments appear to encapsulate the same underlying theoretical structure. To the extent that peripheral commitments distinguish minimalist expressivism from genuine relativism, I will argue that comparable differences obtain between different *forms* of minimalist expressivism (and between different forms of genuine relativism). In that light, the apparent contrast between minimalist expressivism and genuine relativism would seem to draw a distinction without a substantive difference. The formal semantic evidence suggests that minimalist expressivism and genuine relativism are extraordinarily similar and compatible views – that they are two sides of the same coin.

To begin with, all four frameworks define a consistent truth-assignment to all sentence contents, or an entity expressible as one (**feature #1**). Gibbard (1990) calls this entity a *factual-normative world*, Horgan & Timmons (2006) call it a *complete valuation*, Kölbel (2002) calls it a *perspective*, and Lasersohn (2017) calls it a *context of assessment*.

All four semantic frameworks proceed to cleanly bifurcate the set of all sentence contents into a set of factual contents and a set of non-factual contents (**feature #2**). As a consequence, each framework's complete truth-assignment entity is expressible as an ordered pair containing (1) a consistent truth-assignment to all factual contents, and (2) a consistent truth-assignment to all non-factual contents. Gibbard (1990) rigorously differentiates between *factual* and *normative* statements, and represents his factual-normative worlds as $\langle w, n \rangle$, where w and n are consistent truth-assignments to all factual and normative contents, respectively. Horgan & Timmons (2000, 2006) sharply differentiate between *descriptive* and *moral* statements, such that their complete valuations are expressible as *descriptive-moral valuations* $\langle d, m \rangle$. Kölbel (2002) strictly separates *objective* from *non-objective* contents, with his perspectives thereby expressible as $\langle o, n-o \rangle$. Moreover, he strongly implies that only objective contents are factual, and that only non-objective contents are evaluative, such that his perspectives are also expressible as *factual-evaluative worlds* $\langle w, e \rangle$. Lastly, Lasersohn (2005, 2017) rigidly differentiates between *objective* and *subjective* contents, with his contexts of assessment thereby expressible as *objective-subjective assessments* $\langle o, s \rangle$.

All four semantic frameworks accept a deflated and intrinsically relativized notion of truth as primitive (**feature #3**). Each semantic framework relativizes the truth-values of sentence contents to truth-assignment entities (i.e. factual-normative worlds, valuations, perspectives, or contexts of assessment), with the result that sentence contents are expressible as functions from these entities to deflated truth-values. Gibbard (1990) prefers to speak of *deflated* truth-values, and of sentences “holding in” factual-normative worlds. He yields the contents of sentences as functions from factual-normative worlds to deflated truth-values. Horgan & Timmons (2000, 2006) prefer to speak of *minimalist* truth (and of “morally engaged” truth-assertions “infused with one’s own moral commitment”). Sentences are assigned “non-substantive” truth-values (T or F) by valuations, with the result that contents are expressible as functions from complete valuations to minimalist truth-values. Kölbel (2002) and Lasersohn (2005, 2017) prefer to speak

of *relative truth*, and explicitly relativize truth to perspectives and contexts of assessment, respectively. On either view, sentence contents are expressible as functions from these entities to relativized truth-values. The authors of all four frameworks stress that the truth-values they take as primitive are not to be confused with objective, absolute, or correspondence notions of truth and falsity.

At this stage, each semantic framework proceeds to define narrower, correspondence notions of truth and falsity (i.e. TRUTH and FALSITY) in terms of the more primitive, deflated notions. In each case, TRUTH (FALSITY) is defined as a special case of deflated, relative truth (falsity) – as deflated truth (falsity) relative to all truth-assignment entities indexed to the actual world $w_@$ (**feature #4**). Gibbard (1990) explicitly extends possible worlds semantics, where a correspondence notion of truth (i.e. TRUTH) is defined as truth relative to the actual world $w_@$. Analogously, in Gibbard’s (1990) system, correspondence-truth follows as deflated truth relative to $\langle w_@, n \rangle$, for all values of n . Lasersohn’s (2005, 2017) framework also extends a form of possible worlds semantics, where objective truth in a world is defined as truth relative to all contexts of assessment indexed to that world. He then explicitly defines correspondence (or absolute) truth as objective truth in the actual world $w_@$ – in other words, as truth relative to all contexts of assessment indexed to $w_@$. Kölbel’s (2002) definition of objective truth is structurally identical to Lasersohn’s (2005, 2017) – defined as truth in every person’s perspective (in the same world). Correspondence truth follows as objective truth in the actual world. Just as Lasersohn’s (2017) definition of absolute (or correspondence) truth can be expressed as truth relative to $\langle w_@, j \rangle$, for all values of j , Kölbel’s (2002) can be written as truth in $\langle w_@, e \rangle$, for all values of e . Horgan & Timmons (2006) likewise define a special valuation ($V_@$) whose assignment of minimalist truth-values “actually corresponds to how things are or not.” (p.280) Since Horgan & Timmons (2006) require $V_@$ to be a consistent truth-assignment to all descriptive sentences $d_@$, it follows that a sentence is correspondence-true (i.e. TRUE) just in case it is assigned T by $\langle d_@, m \rangle$, for all values of m .

Somewhat tangentially to their main systems, the authors of all four semantic frameworks stipulate that the actual world $w_@$ assigns deflated truth-values to all and only factual contents, and that the sentence contents they set out to model (i.e. normative, moral, and taste contents) are non-factual. As a result, each author ends up assigning correspondence truth-values (i.e. TRUTH

and FALSITY) exclusively to factual contents, and to none of the sentence contents they set out to model (**feature #5**). In effect, each author claims that only descriptive sentences can be either TRUE or FALSE, and that the evaluative sentences the framework sets out to model are neither TRUE nor FALSE.

For instance, Gibbard (1990) denies that there are any normative facts. Since only factual statements can hold in $\langle w_@, n \rangle$ for every value of n , it follows (given the definitions of TRUTH and FALSITY above) that every normative statement is neither TRUE nor FALSE. By contrast, each factual statement is either TRUE or FALSE. Horgan & Timmons (2006) claim that the correspondence valuation ($V_@$) assigns minimalist truth-values to all descriptive sentences and to no moral sentences. It follows that each descriptive sentence is either TRUE or FALSE, and that moral sentences are neither TRUE nor FALSE. Kölbel (2002) and Lasersohn (2005, 2017) claim that taste contents are non-objective, or subjective: that they can vary in truth-value relative to different perspectives or judges in the same world. By contrast, they claim that factual contents are objective: that they do not vary in truth-value across perspectives or judges in the same world. It follows (on either framework) that each factual content is either TRUE or FALSE, and that taste contents are neither TRUE nor FALSE.

Finally, each semantic framework yields the *intension* of any sentence content (i.e. the complete set of conditions in which it is true) as the set of all complete truth-assignment entities which count it as minimalistically-true (**feature #6**). For Gibbard (1990), the semantic content (or intension) of any given sentence is the set of factual-normative worlds in which it holds. In Horgan & Timmons's (2006) framework, the semantic content (or intension) of a sentence is the set of all complete valuations which assign it T. Similarly, the semantic content (or intension) of a sentence in Kölbel's (2002) framework can be construed as the set of perspectives in which the content is true. Finally, Lasersohn (2005, 2017) yields the semantic content (or intension) of any given sentence content as the set of contexts of assessment relative to which the content is true.

In summary, the semantic frameworks of the minimalist expressivists (Gibbard (1990) and Horgan & Timmons (2000, 2006)) and the genuine relativists (Kölbel (2002) and Lasersohn (2005, 2017)) share six key commonalities. They each:

Feature #1) Define a consistent truth-assignment \mathbf{E} to all sentence contents.

Feature #2) Bifurcate the set of all sentence contents into a factual and a non-factual component, such that \mathbf{E} is expressible as an ordered pair $\langle \mathbf{f}, \mathbf{n} \rangle$, where \mathbf{f} and \mathbf{n} are consistent truth-assignments to all factual and non-factual contents, respectively.

Feature #3) Accept a deflated and intrinsically relativized notion of truth as primitive. Truth-values of sentence contents are relativized to truth-assignment entities, with the result that sentence contents are expressible as functions from these entities to deflated truth-values.

Feature #4) Define narrower, correspondence notions of truth and falsity (i.e. TRUTH and FALSITY) in terms of the primitive, deflated notions. TRUTH (FALSITY) is defined as deflated truth (falsity) relative to all truth-assignment entities indexed to the actual world $w_@$.

Feature #5) Stipulate that the actual world $w_@$ exclusively assigns truth-values to factual contents, and that evaluative sentence contents are non-factual. As a result, each factual content is either TRUE or FALSE, whereas every evaluative content is neither TRUE nor FALSE.

Feature #6) Yield the semantic content (or *intension*) of any sentence S as the set of all complete truth-assignments \mathbf{E} relative to which S counts as minimalistically-true.

Notwithstanding the various details needed to handle linguistic complexity (viz. compositionality, quantification, predication, etc.),²⁰ the six features above form the core of a generic semantic theory grounded in deflationary and relativized truth-conditional meaning. Since each of Gibbard's (1990), Horgan & Timmons's (2000, 2006), Kölbel's (2002), and Lasersohn's (2005, 2017) semantic frameworks embrace this core as common ground, they appear to be closely related formulations of the same underlying theoretical structure.

To be sure, a few peripheral commitments help to distinguish minimalist expressivism from genuine relativism. For instance, Gibbard's (1990) factual-normative worlds and Horgan & Timmons's (2006) complete valuations are characterized as completions (i.e. fully-opinionated idealizations) of belief states. By contrast, Kölbel's (2002) perspectives and Lasersohn's (2005, 2017) contexts of evaluation function more like independent standards at which beliefs aim, and

²⁰ I have opted to forgo comparing minimalist expressivist and genuine relativist frameworks with respect to how they handle compositionality, quantification, and predication. This is not least because not every framework I considered provides a comprehensive treatment of each, but also because the two views are unlikely to diverge with respect to these more mundane details, and ultimately because doing so would extend the discussion beyond what is relevant here.

At any rate, it is worth noting that Horgan & Timmons' (2006) and Lasersohn's (2005, 2017) compositional rules, treatments of predication, and the belief predicate are largely interchangeable.

to which beliefs are answerable. Similarly, whereas Kölbel (2002) and Lasersohn (2005, 2017) take the deflated truth-conditions of evaluative sentences to fully capture their semantic content, Gibbard (1990) and Horgan & Timmons (2006) regard this as insufficient. For them, an expressivist dimension must be added: some of an evaluative sentence's meaning is also provided by the mental state it expresses. Thus, for minimalist expressivists, the distinction between factual and non-factual contents is not merely a matter of whether some contents are relativized to non-world parameters – it is also a matter of their functional role in discourse: the kinds of mental states they express and the illocutionary acts they are used to perform.

That being said, to the extent that differences like these help to distinguish the minimalist expressivist and genuine relativist camps, comparable differences obtain between different *forms* of minimalist expressivism (and between different forms of genuine relativism).²¹ For instance, Gibbard (1990) and Horgan & Timmons (2000, 2006) diverge considerably in how they characterize moral beliefs and their relationship to truth. Whereas Horgan & Timmons (2000, 2006) claim that moral beliefs are genuinely cognitive (i.e. truth-apt and assertible), Gibbard (1990), in more standard noncognitive fashion, describes normative beliefs as intrinsically action-guiding (rather than truth-aiming) states. This difference is reflected in their accounts of disagreement: whereas Horgan & Timmons (2000, 2006) represent conflicting moral beliefs as genuinely contradictory (analogous to the manner in which Kölbel (2002) and Lasersohn (2005) do), Gibbard (1990) frames normative disagreements as a clash of attitudes. Indeed, in his later work, Gibbard (2003) reinterprets this clash as the tension between plan-like states with non-jointly-realizable goals.

Comparable differences likewise obtain between different forms of genuine relativism. For instance, it is peculiar to Kölbel's (2002) system that the difference between objective and non-objective contents is determined by *a priori* discursive conventions which “specify the topics disagreement on which indicates a mistake and warrants discussion, and other topics disagreement on which does not indicate a mistake[.]” (p.105) Note that Kölbel (2002) is not merely saying that social and linguistic conventions influence which contents people consider objective and non-objective – he is making the vastly more controversial claim that a content's

²¹ Since the magnitude of the differences involved can hardly be quantified, I trust readers to judge the matter for themselves.

objectivity (or non-objectivity) *just is* a matter of linguistic conventions. Lasersohn (2005, 2017), by contrast, grounds the distinction between his objective and subjective contents in the pre-theoretic distinction between matters of fact and matters of opinion. Lasersohn (2017) then writes:

[Although] it is a significant fact that we do “come equipped,” so to speak, with intuitions which pretheoretically identify certain kinds of sentences as concerning matters of opinion and others not[,] ... [t]his is not to claim that our intuitions are always reliable, or that the mere intuition that some sentences are only subjectively true or false is enough to establish that they are. (p.4)

In Lasersohn’s (2005, 2017) system, then, the distinction between objectivity and subjectivity is not essentially a matter to be settled by linguistic conventions (or by intuitions motivated by them), but rather only a matter as to whether a content’s truth-value is sensitive to non-world indices – a matter of fact to be settled on a case-by-case basis, independent of linguistic conventions.

Since the differences between *forms* of minimalist expressivism (and between forms of genuine relativism) are comparable to the differences between the minimalist expressivist and genuine relativism camps, I submit that although the two views *can* be differentiated, that the grounds for doing so are roughly on a par with the grounds for distinguishing Gibbard’s (1990) view from Horgan & Timmons’ (2000, 2006), or the grounds for distinguishing Kölbel’s (2002) account from Lasersohn’s (2005, 2017). That is to say that the apparent divide between the two camps would seem to carve a distinction without a substantive difference. Our linguistic conventions ought to reflect this fact: the formal semantic evidence suggests that minimalist expressivism and genuine relativism are extraordinarily similar and compatible views – that if a distinction is to be made, it is the distinction which obtains between two sides of the same coin.

4.4 Bridging the divide

In the previous section, I showed that minimalist expressivists and genuine relativists advance nearly identical semantic frameworks. Although a few commitments differentiate the two camps, I argued that they are comparable to the variations within each camp. In this final section, I will briefly argue that the minor differences between minimalist expressivism and genuine relativism reflect shortcomings on both ends, the rectification of which dissolves what

little remains of the distinction between the two views. In particular, I will argue that (i) minimalist expressivists should differentiate each person's evaluative beliefs from the evaluative sentences which are "true for them" (in order to accommodate the possibility of substantive evaluative mistakes), and that (ii) genuine relativists ought to emphasize the expressive dimension of evaluative judgments (in order to explain their functional role). In a phrase, there are compelling reasons for minimalist expressivists and genuine relativists to adopt each other's techniques, and to find common cause in evaluative nondescriptivism. I begin first with the minimalist expressivists' hesitancy to wave the banner of relativism.

4.4.1 Expressivism's need for genuine relativism

Minimalist expressivists are extraordinarily careful not to *explicitly* relativize the truth-values of evaluative sentences. Although they *implicitly* relativize the truth-values of evaluative sentence contents to different outlooks in their semantic frameworks, they nevertheless shrink from using expressions like 'relative truth' and 'truth for a person' – as though they were taboo. Granted, it is perfectly understandable why: since the inception of the noncognitivist tradition, the emotivists and their successors have conceived of their views as *responses* to the problems for traditional forms of content relativism (i.e. subjectivism / contextualism). As discussed in Section 4.1, Ayer (1936) and Stevenson (1944) established the independence of their brand of noncognitivism by resisting its assimilation into the fairly broad and much older schema of moral relativism. More recently, Horgan & Timmons (2006b) have defended their view from arguments by Shafer-Landau (2003) and Bloomfield (2003) to the effect that expressivism's irrealist "denial of moral high ground implies an insidious form of relativism." (p.512)

Although I will refrain from rehashing the dialectic of that particular debate, I mention it to draw attention to two details. First, moral relativism has something of a dirty reputation (earned or otherwise). This has had the unfortunate consequence of warping the stakes of debates about whether noncognitivism and its successors are forms of relativism, distracting from what is substantive about those debates. This brings me to my second point: the crux of the issue is (or ought to be) whether minimalist expressivism's irrealist denial of moral high ground precludes

the possibility of genuine evaluative mistakes – not whether that denial qualifies as a form of relativism.

At a glance, the important question is whether minimalist expressivism denies us (or threatens to deny us) the ability to substantively rebut pernicious moral claims and beliefs – to disqualify them on the grounds that they are genuinely mistaken. As Bloomfield (2003) writes,

If one denies that there is a “higher” moral truth at the [morally] disengaged [i.e. objective] level, then when both parties in a dispute use theories that are equally consistent or rational from their own points of view, they are both equally in the right. And if this is not an insidious form of moral relativism, then there isn’t one. (p.515)

Bloomfield (2003) provides the following example. The Taliban claim that it is not good to educate women. We disagree, and claim that it *is* good to educate women. Does moral expressivism permit us to decisively rebut the Taliban’s claim, or does it force us to concede (at some level) that both claims are “equally in the right” – that they are merely consistent with different points of view (i.e. different systems of moral beliefs), neither of which occupies “the objective moral high ground?”

It seems to me that the moral expressivist is in a bind. Intuitively, moral expressivism is at a disadvantage if it is unable to procure a *substantive* sense in which the Taliban’s claim is mistaken. It is not enough that we disagree with what they say, for by the same token, they disagree with what we say. The point is not to show that each party *considers* the other to be mistaken, but to show that one of them – the Taliban – *is* substantively mistaken. Although it is certainly possible for the Taliban’s claim to stem from false ordinary descriptive beliefs – say, about the likely consequences of educating women, or about what God believes – it is equally possible that the Taliban’s claim derives exclusively from their system of moral beliefs. Suppose, for instance, that we find ourselves disagreeing with a member of the Taliban (i.e. a Talib) who (i) claims that educating women is wrong, (ii) is fully informed about the factual differences between societies where women are educated and those where they are not, and (iii) believes that God disapproves of educating women *because* it is wrong (rather than believing that wrongness is constituted by God’s disapproval). Since moral expressivists explicitly deny the possibility of objectively false (or mistaken) moral judgments, assertions, and beliefs – it is hard to see how a

moral expressivist *can* consistently show that the informed Talib's claim is substantively mistaken.

As it stands, expressivists lack a viable solution to the problem. Nevertheless, it seems that they can make some progress by distinguishing *substantive* from *objective* mistakes. Since moral expressivists cannot consistently claim that moral judgments are objectively false or mistaken, the possibility of objective error is out of the question. The only remaining option would be to posit a form of substantive yet non-objective error. The trouble is that expressivists' only non-objective notion – the deflationary, or “morally engaged” sense in which an individual deems every content they disagree with “mistaken” – is not substantive. Non-objective yet substantive error *could* take the form of alleging that the Taliban's claim is false according to their own system of morals – but this option is a nonstarter if systems of morals are (as expressivists maintain) equivalent to moral beliefs (since by stipulation, the Taliban believe that their claim is true). Without the resources to carve the necessary distinction, expressivism requires modification.

Ironically, expressivists can sidestep the charge of “insidious relativism” – of lacking an account of substantive evaluative error – by borrowing techniques from genuine relativism. Genuine relativists are careful to differentiate between (a) the individuals relative to which a non-objective sentence content p is true, and (b) the individuals who believe p is true. In other words, whether p is true for an individual J (in genuine relativism) is *not* settled by whether J believes that p . Instead, whether p is true for J is settled by whether J is constituted such that p is true for them. For example, suppose that Mark is overwhelmingly disposed to enjoy sushi. In that case, ‘sushi is delicious’ is true for Mark, even if Mark sincerely believes that sentence is false.

By dint of differentiating beliefs from the relativized standards to which they are answerable, genuine relativists are able to differentiate between two forms of substantive mistake: (i) J 's believing a content which is necessarily false for everyone (in J 's world), and (ii) J 's believing a content which is false for J , but potentially true for other individuals (in J 's world). Both forms of error involve J believing a content which is false relative to their own perspective (to use Kölbel's (2002) terminology) – but whereas the first kind of error is objective, the second is non-objective. The second form of error – J 's believing a non-objective

content which is false for J – is precisely the substantive yet non-objective notion of mistake which the expressivist needs.

We are now in a position to see the kind of substantive, non-factual mistake which moral expressivists should say pernicious moral claims often involve: the assertion of a content p which is false for the speaker (in the genuine relativist's sense), in spite of the speaker's belief that p is true. The Taliban may genuinely believe that it is not good to educate women, and for all that, be substantively mistaken – without being objectively mistaken. This occurs just in case 'educating women is not good' is false for the Taliban, by dint of how they are constituted, and in spite of what they sincerely believe. The fine line the minimalist expressivist should walk, then, is to claim that what is good or bad *for* an individual depends on – is relative to – their moral constitution, rather than their moral beliefs. This position is as consistent with moral irrealist as it is with taste irrealist: just as genuine relativists can relativize the truth of taste judgments to people's palettes (rather than their beliefs) and continue to deny the possibility of objectively tasty foods, so too can one relativize the truth of moral judgments to people's moral constitutions (rather than their beliefs) and continue to deny the possibility of objectively good actions.

In review, minimalist expressivism lacks the resources to accommodate the possibility of substantive (i.e. genuine) evaluative mistakes. By definition, minimalist expressivists cannot say that evaluative statements are objectively mistaken. They *can* say that evaluative claims are "mistaken" in a deflationary, or "morally engaged" way – but this is not a substantive notion of error. Without the ability to distinguish *substantive* error from *objective* error, minimalist expressivists cannot accommodate the possibility of substantive evaluative mistakes. One way to carve the necessary distinction would be to say that evaluative statements can be false for their speaker, in spite of what they sincerely believe. Unfortunately, this approach is blocked by expressivists' view that an evaluative statement being "false for the speaker" is just a matter of what the speaker's moral beliefs are. The solution rests, then, in giving this commitment up, and accepting – in the mold of genuine relativism – that what is good or bad for an individual depends on – *is relative to* – their moral constitution, rather than their moral beliefs. Doing so permits moral expressivists to accommodate the possibility of substantively mistaken evaluative beliefs, without giving up moral irrealist. There are compelling reasons, then, to dissolve one of the only remaining distinctions between minimalist expressivism and genuine relativism. On that

note, I turn now to what seems the only important distinction remaining between the two camps: genuine relativists' reluctance to embrace the expressive role of evaluative judgments.

4.4.2 Genuine relativism's need for expressivism

Genuine relativists say remarkably little about what taste judgments are *typically used for*. What are people doing when they say that rollercoasters are fun, or that sushi is tasty? What are they trying to communicate or achieve? Contextualists and expressivists give straight answers. Contextualists say that taste judgments are primarily used to *report* or *describe* facts – facts about people's mental states, or about the contents of idealized standards. Expressivists reject this account, claiming instead that taste judgments are merely used to *express* mental states.²² There is a clean divide between the two groups: whereas contextualists advance a descriptivist account of the functional role of taste judgments, expressivists defend a form of nondescriptivism and emphasize their expressive function.

What do genuine relativists say? They say that the contents of taste judgments are context-invariant, truth-evaluative, and non-objective; they claim that taste judgments express the same content in all contexts, that these contents are true relative to some perspectives (and false relative to others), but neither true nor false *tout court*. On the basis of these points, genuine relativists reject both the realist view that taste judgments describe mind-independent facts, and the contextualist view that taste judgments describe mind-dependent facts (such as mental states and social standards). But if taste judgments neither describe mind-independent facts, nor report mind-dependent facts, then they are not used to report facts of any kind. Genuine relativism's anti-realism and anti-contextualism therefore suggests some form of *non*-descriptivism. This much is clear – but this thesis remains purely negative. If taste judgments are not used to describe any kind of fact, then what function(s) *do* they serve? Genuine relativists owe us a positive account.

²² Some caution is necessary here, as describing a mental state is not (in general) mutually exclusive with expressing a mental state. Many declarative sentences – for example, 'Jones likes ice cream' – report a mental state (viz. Jones's enjoyment of ice cream) *and* express a mental state (viz. belief that Jones likes ice cream). Expressivists suppose that all declarative sentences express mental states. According to taste expressivists, the distinctive feature of taste judgments is that they are *not* used to describe facts. Hence, taste expressivism is virtually a negative thesis: it is the view that taste judgments are *merely* used to express mental states, since they are *not* used to describe facts.

Unfortunately, genuine relativists do not explicitly say what the communicative function(s) of taste assertions is (or are); they simply do not offer a robust positive account to go hand-in-hand with their implicit nondescriptivism. The most that we can gather from their semantic program is that taste judgments are typically used to assert propositions one believes are true in one's own perspective – but what does that come to, in the end? We have just seen that genuine relativism must, of necessity, be a form of nondescriptivism. But if, in the act of asserting a proposition a speaker believes true in their own perspective, they are neither *describing* their taste beliefs nor *reporting* what is true in their unique perspective – what *are* they doing? At this stage, the remarkable simplicity of minimalist expressivism's answer is (or ought to be) tempting: to assert that *p* is to *express* (rather than *report*) one's belief that *p*. Expressivists point out that sincere assertions express (i.e. reflect) something about their speaker: the mental state that the speaker needs to have in order for their assertion to *be* sincere. The minimalist expressivists' key insight is to recognize that since every judgment can be sincerely asserted, every judgment expresses (i.e. gives voice to) some mental state; even if a judgment is primarily used to describe (or report) a fact, it is also used to *express* belief in that fact. It follows that if a judgment *p* is *not* used to describe any fact, that *p* is *merely* used to express a mental state.

The minimalist expressivists' answer – that taste judgments are merely used to express mental states – follows from two theses: (i) a controversial negative thesis (nondescriptivism) and (ii) an uncontroversial positive thesis (that assertions express the mental states that make them sincere). Genuine relativists should adopt the expressivists' answer, not only because genuine relativists already (implicitly) accept nondescriptivism about taste – but also because there is little reason to deny that assertions express (i.e. give voice to) the mental states that make them sincere. In lieu of providing a new and principled explanation of what speakers are doing when they assert taste judgments, the path of least resistance is to acknowledge that the function of taste assertions is to express (rather than report) taste beliefs.

In what remains of this subsection, I will show that the expressivists' answer can be teased from genuine relativists' views about the conditions in which taste assertions (and beliefs) are justified. As I will attempt to show, genuine relativism suggests that to assert a taste content *p*

is (typically) to express one's belief that p – but without explicitly reporting this belief, as contextualists generally suppose. As Lasersohn (2005) writes,

[W]e typically *assert* [taste contents] from an autocentric perspective. That is, in making an assertion, we regard it as in some sense *justified* iff it is true relative to that context [of assessment] in which we ourselves serve as judge. (p.670)

In other words, the assertion of a taste judgment is (typically) *justified* just in case the judgment is true in (or relative to) its speaker's perspective. So, to *believe* that one's own taste assertion is justified is (typically) to *believe* that it is true in one's own perspective. Since all assertions believed by their speaker to be epistemically justified are sincere, it follows that an assertion of a taste content p is (typically) *sincere* if its speaker *believes* that p is true in their own perspective – even if this belief turns out to be false. Now, assertions express (or give voice to) something about their speaker – assertions express (or give voice to) the mental state(s) that would make them sincere.²³ ²⁴ It follows that taste assertions that p (typically) *express* one's belief that p is true in one's own perspective.

Moreover, genuine relativists suppose that there is (typically) a strong connection between believing that p and believing that p is true in one's own perspective. As Lasersohn (2005) writes,

[I]n the typical case we adopt an autocentric perspective[:] to believe a [non-objective] sentence content should normally involve believing that it is true relative to ourselves. For example, if John believes that the Giant Dipper is fun, ... he should also believe that the Giant Dipper is fun for himself.

This is *not* to say that *The Giant Dipper is fun* and *The Giant Dipper is fun for John* express the same content relative to John. It is merely to claim that if John, adopting an autocentric

²³ For example, to assert 'Larry sent me the wrong damn package' is to express (or give voice to) both (i) one's belief that Larry sent the wrong package, and (ii) one's anger that Larry sent the wrong package.

²⁴ Although I invoke what is essentially the expressivists' positive thesis here, I do so only because it is virtually incoherent to deny. A common truism about assertion is that to assert that p is to present p as true. But to present p as true is – among other things – to present oneself as *believing* that p . (Even in cases of deception, audiences falsely conclude that a speaker believes that p because the speaker presents themselves as believing that p .) Hence, for a speaker to (sincerely) assert that p is – among other things – to (sincerely) present oneself as believing that p ; to make an assertion is to express (or give voice to) a belief with the same content. More generically, to assert that p is to present oneself as having whichever mental state(s) suffice for an assertion of p to be sincere.

perspective, believes the content of one of these sentences, he should believe the content of the other one as well. (pp.675-676)

Genuine relativism therefore suggests the following three points:²⁵

(Autocentric) Belief Expression (BE): To assert a taste content p is (typically) to express one's belief that p is true in one's own perspective.

(Autocentric) Belief Assessment (BA): To believe a taste content p is true in one's own perspective (typically) involves believing p (and vice versa).

Anti-contextualism (AC): To assert a taste content p is not to *report* that one believes p .

The first two of these points – **(BE)** and **(BA)** – jointly entail that to assert a taste content p is to express one's belief that p .²⁶ Together with **(AC)**, then, genuine relativism suggests that to assert a taste content p is to express one's belief that p , but *not* to report (or describe) that belief. This, of course, is identical with the expressivists' account. In conclusion, genuine relativists can (and should) join expressivists in saying that the functional role of taste judgments is merely to *express* – rather than *report* – taste beliefs.

4.4.3 Taste nondescriptivism

In Sections 4.1 - 4.3, I showed that minimalist expressivism and genuine relativism are remarkably similar and largely compatible views – that they are two sides of the same coin. In the previous two subsections, I argued that the outstanding differences between the two views reflect shortcomings for both, the remedies to which dissolve what little remains at issue between them. In effect, minimalist expressivism and genuine relativism resolve (or should resolve) to the same thing: a form of evaluative nondescriptivism replete with (i) an irrealist metaphysics, (ii) an

²⁵ Kölbel (2002) expresses equivalent conditions of justification for taste assertions and beliefs in his relativized mistake criterion: if a person believes (or asserts a sentence expressing) a proposition which is false according to the perspective that they possess (at that time), then they have made a mistake (or are at fault). Insofar as individuals aim *not* to make mistakes, it follows that their assertions and beliefs aim at what is true in their own perspective. Hence, *sincere* assertions express what speakers *believe* is true in their own perspective.

²⁶ Note that we may drop 'typically.' This is because whichever perspective a speaker "adopts" in assertion – say, S – then **(BE)** to assert p will be to express one's belief that p is true for S . But by virtue of "adopting" the perspective S , **(BA)** to believe that p is true for S will involve believing that p . Hence, whichever perspective one "adopts" in assertion, to assert a taste content p will be to express one's belief that p .

implicitly relativist truth-conditional semantics, (iii) a notion of substantive yet non-objective error, and (iv) an expressive account of the functional role of evaluative assertions.

If I am right that minimalist expressivism and genuine relativism are aligned in this way, then the time has come for the contours of the literature on the semantics of taste judgments to change. Rather than identify three genuine alternatives (contextualism, genuine relativism, and expressivism), we ought to say that there are just two: taste *descriptivism* and taste *nondescriptivism*. In that light, my arguments against realism and contextualism in *Chapters 2* and *3* show that taste descriptivism is fraught with difficulties, and that taste nondescriptivism presents the solution to these difficulties. What I take this to show, in the end, is that there is at least one discourse – taste discourse – for which nondescriptivism is not just plausible, but probable.