

The Dissertation Committee for Jeremy Andrew Lupo certifies that this is the approved version of the following dissertation:

Towards an Improved Quasi-Realism

Committee:

John Deigh, Supervisor

Cory Juhl

Larry Laudan

Josh Dever

Ernest Sosa

Daniel Bonevac

Towards an Improved Quasi-Realism

by

Jeremy Andrew Lupo, B.A.

Dissertation

Presented to the Faculty of the Graduate School

of the University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

August 2015

Towards an Improved Quasi-Realism

by

Jeremy Andrew Lupo, Ph.D.

The University of Texas at Austin, 2015

Supervisor, John Deigh

In my dissertation, “Towards an Improved Quasi-Realism,” I lay out, evaluate, and then improve upon existing varieties of quasi-realist expressivism in metaethics, in particular with regards to solving the problem of embedding.

I outline two distinctions to classify quasi-realist theories: between fast-track and slow-track approaches, and between univocal and mimicking approaches. I argue that a successful theory will use the mimicking approach and a structure that combines fast-track and slow-track elements-- that is, it will work to justify the use of propositional form in expressing attitudes using concerns related to the wider institution that we identify as moral practice.

After historical discussion of A. J. Ayer, I. A. Richards, and C. L. Stevenson, in which I attempt to bring out overlooked nuances of how the embedding problem relates to their theories, I turn my attention to Simon Blackburn’s notion of “propositional reflection”-- the idea that a structure of our attitudes can ground an expressivist’s explanation of why we use apparent propositions to express attitudes. I argue that none of Blackburn’s candidates for the relation is adequate. Theories that separate what is expressed from what is reflected (as in Allan Gibbard’s *Thinking How to Live* view) do better, and I contribute a suggestion of my own according to which a commitment is a state of standing ready to take action, but, ultimately, the right view is one that adopts a pluralistic view of the attitudinal dynamics that underlie our

adoption of propositional form for the expressions of attitudes that constitute moral language.

Finally, I consider the quasi-realist maneuver itself, the assimilation of paradigmatically realist metaethical claims by the expressivist via minimalist or expressivist analyses of those claims. I determine that the expressivist form of quasi-realism is more promising than the minimalist, and that the quasi-realist can minimize the burden he carries from denying intuitively plausible realist claims while at the same time keeping quasi-realist expressivism clearly distinct from realism. The harsher consequences of adopting a mimicking theory can be rendered harmless by quasi-realism, resulting in the most plausible form of expressivist theory.

Table of Contents

Text.....	1
Chapter 1: Introduction: Two Problems for Expressivism.....	1
1.1 Expressivism.....	1
1.2 The Frege-Geach Problem.....	5
1.3 Dispute and Practice.....	9
Chapter 2: Slow and Fast Tracks; Mimicking and Univocal Approaches.....	15
2.1 The Slow Track.....	15
2.2 The Fast Track.....	23
2.3 Hybrids.....	27
2.4 Need for fast-track elements.....	32
2.5 Preferring mimicking expressivism.....	43
Chapter 3: Ayer.....	54
3.1 Ayer’s emotivism, version 1.....	55
3.2 Ayer’s emotivism improved.....	62
3.3 Ayer and quasi-realism.....	73
3.4 Using Ayer’s ideas within a hybrid theory.....	79
Chapter 4: Richards and Stevenson.....	86
4.1 Richards.....	86
4.2 Richards and “emotive meaning”.....	99

4.3 Stevenson.....	103
4.4 Relationship between Stevenson’s ideas and the Frege-Geach problem.....	111
4.5 What to learn from Richards and Stevenson.....	118
Chapter 5: Propositional Reflection.....	120
5.1 Propositional Reflections: The General Idea.....	121
5.2 Blackburn’s Attempts at Propositional Reflection.....	129
5.3 Another Propositional Reflection.....	138
5.4 The Pluralistic View.....	143
5.5 The pluralistic view and Richards.....	147
Chapter 6: Quasi-Realism.....	153
6.1 Some difficulties for quasi-realism.....	155
6.2 More difficulties for quasi-realism.....	162
6.3 Nuanced quasi-realism.....	170
6.4 Richards and Stevenson again.....	174
6.5 Dialectical constraints.....	180
6.6 Some signs that meta-ethics is itself expressive.....	187
6.7 Future.....	193
Bibliography.....	200

Chapter 1. Introduction: Two Problems for Expressivism

1.1 Expressivism

On the face of it, the sentences that feature in moral discourse resemble those that feature in discourse about things. Structurally, “this is wrong” seems similar to “this is red” or “this is a chair”. And we work with it similarly as well; we embed it, and thereby construct sentences such as “if this is wrong, you will burn in Hell for all eternity,” without marking any particular difference in what we are doing from when we construct “if this is a chair, its owner must be very short.” We disagree about it; one person might claim “this is wrong” and another “this is not wrong”, and thereby enter into what seems like the same sort of relationship with one another as they enter into when one claims “this is a chair” and the other “this is not a chair.” And we argue about it; one might, for instance, construct the argument “if this is murder, then it is wrong; this is murder; therefore, it is wrong” with the same procedure that we use to come up with the argument “if this is a chair, it won't break when you sit on it; it is a chair; therefore, it won't break when you sit on it.”

All of this evidence strongly suggests that “this is wrong” really is like “this is a chair”; they differ only in that the former attributes the property of being wrong, and the latter, that of being a chair. A realist holds that moral sentences ascribe (moral) properties to objects, and that there are such properties. It is not the goal of this piece to dissertation that such a view is mistaken. Rather, suppose we find, somewhere else,

reason to doubt the view, or, more moderately, to think that this view does not tell the whole story. Such a situation sometimes occurs when one's wider theory seems to conflict with moral realism-- see, for instance, A. J. Ayer's (1936 [1952]) positivism-- or in the face of arguments directed against realism itself-- a choice few examples being J. L. Mackie's (1977) queerness argument¹, Sharon Street's (2006) evolutionary argument², or the problem of moral divergence³-- often called the problem of moral disagreement, but we must rename it in order to avoid confusion with a certain older problem that will serve as a central element of this work.

If we find reason to reject realism, there are a number of views we might take up instead. A number of these carry the label of "expressivism", a term which I will now appropriate to capture the sort of theory that I am interested in examining. Let us use the term "expressivism" to denote the position that the nature of moral discourse is the expression of attitudes. This definition requires several points of clarification.

First, the word "nature" above is not a fancy way of saying "meaning," or indeed as an oblique attempt at highlighting any particular mechanism. Rather, the category of theory that I wish to highlight is one characterized not by any specific claim made by all and only such theories, but by a shared structural feature: theories in this class cite the expression of attitudes in moral discourse as the wellspring of explanatory power for dealing with meta-ethical questions: what moral discourse is for, what is at

1 See also *e.g.* Brink (1984) pp. 111-125; Horgan and Timmons (1992), pp. 221-260; Garner (1990)

2 See also *e.g.* Kahane (2011); Shaver-Landau (2012); Wielenberg (2010)

3 That is, that different groups of people often appear to possess different moral beliefs. For an overview, the introduction of Gowans (2000) is decent. Mackie (1977) is again an important resource. For more, see *e.g.* Harman and Thomson (1996); Enoch (2009); Sturgeon (1994)

stake in moral dispute, whether we need to posit the existence of moral properties with unusual motivational properties, etc. Admittedly, this sort of characterization gives us less to go on than one based on a mechanism, but I do not think that that is avoidable. Indeed, when, in later sections, we turn our attention to some of the existing theories that the definition is meant to capture, we will see that the prospects are dim for quick and dirty semantic approaches to defining expressivism; that, in fact, the most natural candidate for such a definition-- that expressivism constructs the meaning of moral sentences by appeal to the expression of attitudes-- is neither historically apt nor conducive to the proper understanding of the idea it purports to capture. So, although this style of definition may not be as satisfying as some of the alternatives, it will save us trouble in the long run. ⁴

Second, “expression” requires some discussion. Schroeder (2010) takes pains to give the expressivist a distinctive notion of expression, one where the term specifically refers to the sort of relationship found between an ordinary factual belief and the sentence that voices the content of that belief. I find myself unmotivated to follow suit. Schroeder wants terminology that will separate “expressivism” from earlier “emotivist” theories⁵, while I find the label “emotivist” to be flawed⁶-- though, perhaps, useful for

4 For definitions in the literature that share this feature, see *e.g.* Gibbard (2003); Gert (2011); Stoljar (1993) is particularly interesting (assuming we do not allow ourselves to be distracted by Stoljar's choice to use “Emotivism” rather than “expressivism”). Charlow (2014), which draws upon Rosen (1998) and Wedgwood (2007), is related and has useful discussion on this issue, though not quite an example due to some differences in detail. All of this said, I would define expressivism in this way, for practical and historical reasons, even if there were far less precedent for so doing.

5 “Emotivists” get the ordinary sense of “express”, while “expressivists” get the technical sense.

6 The problems being that, firstly, the theories most closely associated with the term-- Ayer's and Stevenson's-- are both far more sophisticated in terms of what they take to be expressed in moral discourse than is suggested by the “emotivist” label; secondly, when understood in full and in their

leading the uninitiated into the general vicinity of the ideas-- and therefore desire a general term that can capture both those theories and their descendants (once again, the historical discussions to come should bring out why this is the case). At any rate, it is also unclear if Schroeder's definition actually succeeds at capturing the theories to which the "expressivist" label is supposed to apply.⁷ If needs be, we can include specific notions of "expression" into particular theories; so long as they are near enough to the ordinary notion to justify the label, we are unlikely to cause confusion by lumping them in with one another and with theories that use the folk concept straight.⁸

most developed forms, these theories share quite a lot with their descendants, decreasing the motivation for separating them.

- 7 Blackburn, for instance, sometimes moves quite freely between "emotivist" and "expressivist" when speaking (see his interview with Darlei Dall'Agno, 2002), and in writing cites emotivism as a classic example of an expressive theory (see *e.g.* Blackburn (1984), p. 167). Anyway, Schroeder himself, in his (2008b) notes that expressivists tend to do very little to define "expression", though they should; for our purposes, though, we should not rule out *bad* expressivism from being expressivism.
- 8 A side-note on the idea of "expression": traditionally, the idea of expressing a mental state has been drawn in contrast to that of attributing the idea to oneself (as in "subjectivism"). Now there is some debate as to whether the distinction is tenable; if not, there can be no expressivism. In the literature, the dispute about this point grew from a 1998 piece by Jackson and Pettit; the point is reiterated in Jackson and Pettit (2003). The central thrust of their argument is that the process for learning to use a sentence to express a mental state will be sufficient to give that sentence the meaning of a report of that mental state, thus collapsing expressivism into subjectivism. Various replies ensued; for instance, Dreier (2004) tries to pull apart the results of the learning considerations that Pettit and Jackson are concerned about from the primary truth conditions of moral language, while Smith and Stoljar (2003) distinguish agreeing to use a word *for* something as opposed to agreeing to use it *when* something is the case to block the unappealing *for* claim that leads to subjectivism. Schroeder's own view, in his (2008b) is that the problem is solved by his preferred form of expressivism, wherein the notion of expression discussed earlier aligns the attitude with the belief in factual cases (and so there is no more danger of moral sentences collapsing into an attribution of an attitude to the speaker as there is of factual sentences collapsing into sentences about the speaker's beliefs). Personally, I find the Smith-Stoljar solution compelling, especially when it is combined with a certain observation: in social interactions, there are times when it is and when it is not appropriate for a speaker to make statements about himself or herself. Speakers who talk about themselves excessively may be regarded as self-centered or even neurotic. There sometimes emerge apparently expressivised versions of self-attributions-- for instance, the ungrammatical "hopefully" or "thankfully" at the beginning of the sentence, which differ from "I hope that" or "I feel thankful for the fact that" in having explicit reference to the speaker removed-- which appear to be usable without violating such conventions in situations where the grammatical versions would do so. A theory that allows purely expressive constructions to form, either directly by learning use-*when*, or indirectly by subtraction, has the advantage of generating a convenient explanation for the emergence of "hopefully" and the like.

Third, I find it necessary to remark on the use of “attitude” within the definition. After all, without further specification, the definition does not say terribly much; there are few mental states that the term “attitude” could not be conceivably stretched to cover. And yet the prospects seem very dim indeed for any *one* definition of attitude to capture the entire category. Perhaps the best option will be to repeat the maneuver used for “expression”: any notion of “attitude”, technical or otherwise, will do, so long as it falls within the intuitively-delineated family resemblance cluster surrounding the ordinary concept. Also, including “attitude” at all has the odd effect that hyphenated expressivisms (“norm-expressivism”, “plan-expressivism”, etc.) aren't properly expressivist unless they mean for their distinctive feature to be interchangeable with some construct of attitudes. Although this aspect is peculiar, it is not a big deal; after all, we already have the hyphenated terms to refer to such theories.

1.2 The Frege-Geach problem

If we choose expressivism as our form of anti-realism, we face the difficulty known as the *Frege-Geach Problem*. Core references for the problem are Geach's (1958), (1960), and (1965), though it was also raised by Searle (1962), and perhaps even earlier by others.⁹ As mentioned at the beginning of the previous section, moral language behaves like non-moral language. For one, it can *embed*, entering into contexts where the moral sentence is not asserted. Either side of a conditional, for instance, can

9 For instance, Schroeder (2010) attributes prefigurings of the problem to Ross (1939), ~pp.30-41; to H.B. Acton (1936); and to Hare (in his (1952 [1964]), ch. 2). The arguments are not quite the same as in Geach, but close enough to be interesting.

contain a moral term:

1.2e1 If Arnold killed Bob, he acted wrongly.

1.2e2 If murder is wrong, Arnold will burn in Hell.

Such sentences can combine with factual sentences or with ones that contain moral terms in asserted contexts to create apparently valid arguments with either a non-moral or a moral conclusion:

1.2e1 If Arnold killed Bob, he acted wrongly.

1.2e3 Arnold killed Bob.

1.2e4 Therefore, he acted wrongly.

1.2e2 If murder is wrong, Arnold will burn in Hell.

1.2e5 Murder is wrong.

1.2e6 Therefore, Arnold will burn in Hell. ¹⁰

Expressivism draws upon the expression of attitudes as its key explanatory mechanism.

¹⁰ If running the inference in this direction seems particularly odd, it may be because of a related difficulty, the “wishful thinking problem”, from Dorr (2000). It seems like an agent who accepts 1.2e2-1.2e5 thereby comes to have reason to accept 1.2e6, but if 1.2e5 is described only as an expression of an attitude, how could it function in that way? For discussion of this issue, see e.g. Budolfson (2011); Mabrito (2013); Schroeder (2011). Observe that this problem only arises for versions of expressivism in which the account of moral disputes in general is subsequent to the account given to such apparently deductive arguments in the order of explanation.

But expressing a mental state is an action, and claiming that it occurs in unasserted contexts as well would collapse the distinction between the two and lead to absurdity.

And so the expressivist has more work to do if he is to explain:

- a) The meaning of sentences such as 1.2e1 and 1.2e2.
- b) The role that “Arnold killed Bob” plays in determining the meaning of 1.2e1, and analogous cases.
- c) The close relationship between “Arnold killed Bob” as it appears in 1.2e1 and as it appears in 1.2e3, and analogous cases.
- d) The validity of 1.2e1-1.2e4, 1.2e2-1.2e6, and analogous cases.

Taken together, these difficulties comprise the *Frege-Geach Problem*, also known as the *Embedding Problem*. Some authors might prefer to divide the problem into a number of sub-problems¹¹, but, for our purposes, we will lump all of the issue's aspects together into a single challenge.

To see how this problem works, let us consider a simple version of expressivism, one on which “murder is wrong” is taken to be interchangeable with “Boo, murder!” Now consider the sentence, “If murder is wrong, then paying someone to commit murder is also wrong.” If the expression of the Boo!-attitude gives the meaning of “murder is wrong”, then “murder is wrong” cannot mean the same thing when it is embedded in sentences like our example, because asserting such a sentence does not involve expressing that attitude. Such is undesirable for multiple reasons. It is an affront to our intuitions about compositionality; it fails to explain the meaning of the

¹¹ For example, Sinnott-Armstrong (2000), divides the problem into four sub-problems.

complex sentence; it makes arguments that we think ought to be straightforwardly valid appear to be instances of a fallacy of equivocation. In order to be plausible, an expressivist theory will have to do better.

Although the difficulty is often called the “embedding problem,” it would be too simple to consider a) to represent the entirety of the problem. For the difficulty, if it is to get anywhere, ought to apply to expressivist theories with a wide range of possible accounts as to the degree of interconnectedness between metaethics and semantics. It is not only a problem for those theorists who approach metaethics as an exercise in determining the meaning of moral sentences; it is also supposed to work against theories that concern themselves entirely with making comments about moral practice, while leaving it up to the philosophers of language to determine which of the moving parts in each theory should count as “meaning” and which not. It should even work against theories that are explicitly anti-semantic, where for one reason or another the theorist thinks that ethical statements lie entirely outside the boundaries of meaning, though they nonetheless have some use.

Rather, it would seem that the center of the issue must be defined by a more general concern: not that of describing the meaning of those sentences where moral terms occur in embedded contexts, but of describing the sentences themselves-- their role and their logic. It is not that this question eliminates the meaning question, but rather that any answer to this question will also answer the meaning question, either by solving it or by showing why we needn't worry about it. And it is a question that

can be aimed even at those theories that do not use semantics as the delivery mechanism for their expressivist explanations.

In order that we might proceed in an organized and methodical manner, we will leave this problem here for now. We will pick it up again later, after we have seen enough of the lay of the land to discuss strategies for dealing with the issue at a high level of abstraction.

1.3 Dispute and practice

Another difficulty for the anti-realist that expressivists, historically, have wanted to address is the problem of moral *dispute*¹². By this I mean the problem that Moore raised¹³ for subjectivism¹⁴: if we take the statements of moral discourse to be claims about the attitudes of the speaker, then it follows that a moral claim and its negation, in the mouths of different speakers, do not contradict one another. For instance, A might disapprove of the recreational immolation of squirrels, while B might condone it; A can describe his attitude by saying, “I disapprove of squirrel-burning,” while B can describe his by saying, “I do not disapprove of squirrel-burning.” In such a case, the statements of the two do not contradict one another; they may both be true, both be false, or one be true and the other false. And yet “Squirrel-burning is wrong”

12 Perhaps better known as the problem of moral disagreement-- not to be confused with what I called the problem of moral divergence, which is sometimes also called the problem of moral disagreement, though the two are entirely different. To avoid creating confusion, I have chosen to discontinue the ambiguous “moral disagreement” for both problems.

13 See Moore (1922)

14 The view that moral utterances report the attitudes of the speaker-- so “murder is wrong” translates to “I disapprove of murder,” and so forth.

and “Squirrel-burning is not wrong” are incompatible even coming from different speakers. Thus, it seems that “I disapprove of squirrel-burning” is not an acceptable analysis of “Squirrel-burning is wrong,” and more generally, that statements of moral language cannot be analyzed as statements about the attitudes of the speaker.

While negations of moral sentences is itself an issue for expressivism¹⁵, the problem of moral dispute is not confined to negation. Where the embedding problem was about moral sentences, their meaning (perhaps), and their underlying logic, the dispute problem is about moral practice, by which I mean the set of customs and social habits involving and surrounding the discussion of moral matters. There are facts about what we *do* that any meta-ethical theory must be prepared to explain. The subjectivist position¹⁶, by holding that there is no contradiction between “murder is wrong” in the mouth of one speaker and “murder is not wrong” in the mouth of another, becomes subject to the challenge of explaining the forms of conflict that typically attend such a case: that the situation is regarded as one of disagreement, that the participants might try to resolve that disagreement through what appears to be rational argument, and that one might even be called irrational for failing to heed the consequences of such an argument.

Ayer (1936 [1952]), whose theory we will investigate in more detail in Chapter 3, recognizes that the problem of dispute raises questions for the expressivist just as it

15 Indeed, the literature currently features quite a bit of material about the “problem of negation,” attributed to Unwin (1999), (2001) and much discussed by and in connection with Schroeder (2008a), (2008b), (2008c), (2008d), (2010).

16 At least in the simple form we are considering.

does for the subjectivist. The expressivist must be able to explain, firstly, what sort of conflicts can arise between opposed expressions of attitudes, and secondly, how such conflicts can stand at the center of the discursive practices that attend moral dispute. If the expressivist seeks to appeal-- as it seems he must-- to some notion of disagreement in attitude, the second aspect of the problem becomes difficult. For disagreements in attitude *typically* call for a different sort of debate from debates about facts; when we take issue with someone's expressed attitude, we may attempt to change it, using a wide variety of persuasive methods, but we do not-- and cannot-- refute it. "Ick, broccoli!" can be reacted to in a variety of ways; it can be rewarded with a disapproving stare, met with the question, "Why don't you like broccoli?", or alternatively with "Have you even tried it?"; it can be fought against with claims that "it's not that bad," "it's good for you," or "I like it". But it cannot be disproved, shown unintelligible, or even deemed unlikely. And yet we do-- or seem to do-- all of the above in moral discourse.

The expressivist, if he is to remain in a viable position, has two strategies that he might pursue. He might attempt to explain the evidence-- that is, he might provide some sort of story to inform us as to why disputes about moral matters proceed as if we were disputing matters of fact-- or he might challenge that observation of moral discourse really does yield the datum that such discourse proceeds in the manner expected of a matter of fact. These strategies are not mutually exclusive; it is open to the expressivist to make as many observations as possible to support his position, and then cover the remainder with explanations.

Independently of the consideration of any particular expressivist theory-- which we'll leave off for now, so that we can do justice to our case studies later and not lose track of our current topic-- we can consider the strategies themselves. In particular, we should consider the extent to which the second strategy might be taken. Ayer argued-- and he would be followed in this by Stevenson¹⁷-- that moral debate was doomed to collapse into non-rational argument, and that, in fact, it would do so as soon as it became genuinely moral (rather than about non-moral facts of the issue at hand). Ayer's position in 1936 is an example of the second strategy; the idea is to show that if we really think about when rational debate proceeds and when it breaks down, we see a pattern that is exactly as we would expect for straightforward expressivism. Success in this matter, though, will hinge not on any detail of Ayer's own theory, but on dedicated arguments concerning the breakdowns (or absence of breakdowns) in moral debate.

This problem is closely tied to the Frege-Geach problem; indeed, some authors¹⁸ consider these to be the same problem. But dividing them is useful for the purpose of exploring solutions. An Ayer-style all-in second-type solution addresses the dispute worry without having to concern itself with Frege-Geach issues (at least within the bounds of the solution to the dispute problem), because such a theory labels all of the fact-like discursive practices as actually being discourse about fact. Alternatively, there seems to be room in logical space for a theory to pursue first-type strategy that is not a

17 e. g. Stevenson (1944)

18 See e.g. Björnsson (2001); Van Roojen (1996).

straightforward solution to Frege-Geach; for instance, we might have one that treats moral practice in a holistic manner, about which we will say more later.

The issue of moral dispute is part of the question of moral practice-- that is, of providing a satisfying general description of what we are doing when we think, talk, and argue about morality. We take ourselves to be engaged in a practice that involves real disagreements, ones where the participants are not merely talking past each other. An expressivist theory can attempt to show how it is that this is so, or it can deny that it is so. Either way, it will need something to say about what is going on when two people seem to disagree about a moral matter.

Here are two ways that an expressivist might answer that demand. On one hand, he might start with an answer to the embedding problem, and thereby come to understand complex moral sentences. Having done so, he will understand what is being said in moral discourse, including in cases where people appear to be arguing. And if he can explain how we can engage in rational argument using expressions of emotions, it is unclear what more we would need to satisfy the question of dispute. After all, if we can engage in rational argument about something, that seems as good a justification as any for calling the dispute genuine. On the other hand, he might start with some thoughts about the dispute problem-- that is, with an idea of what kind of practice moral dispute is part of-- and then attempt to give meanings for sentences with moral terms in embedded contexts based on the role of those sentences within the practice. These two ways are not exhaustive; hybrids are possible.

In order to make these ideas explicit, let us now turn our attention to some examples, courtesy of Simon Blackburn.

Chapter 2. Slow and Fast Tracks; Mimicking and Univocal Approaches

Solving the embedding problem requires describing¹⁹ sentences featuring moral terms in embedded contexts. The most obvious way to approach this task is compositionally: account for how atomic moral sentences can be combined to form complex moral sentences and an account of any complex moral sentence is in the offing. A less obvious way to approach the task is functionally: describe what moral language, in general, is *for*, and then describe the role that complex moral sentences play in the practices of reasoning that we adopt in pursuit of that goal. Each of these approaches corresponds to a different structure that an expressivist theory might adopt; the approach chosen will dictate the overall order of explanation that the theory must pursue. This chapter aims to explain these two structures and their implications for the prospects of theories that follow them.

2.1 The Slow Track

Blackburn's *Spreading the Word* (1984) view will serve as our example of the first structure. Before going into any detail, though, it is important to note that we will be greatly simplifying the views in question. In particular, we will be leaving off all quasi-realism, where by “quasi-realism” I mean the practice of using analysis, typically either minimalist or expressivist, on meta-ethical claims that are usually associated with

¹⁹ Those interested in developing a typical contemporary expressivist theory, that is, one that understands expressivism as a semantic thesis, should substitute “giving the meaning of” for “describing” here.

realism in order to render them compatible with expressivism. This omission is temporary-- we will return to quasi-realism in Chapter 6, at which point we should be better equipped to examine it properly-- and is not philosophical, but merely explanatory, in nature. The topics of this chapter are simply easier to discuss with quasi-realism excluded, even though doing so means that many of the claims here will ultimately need to be revised or even reversed. The purpose of this section is not to create a reconstruction that would be acceptable to a proponent of the original view, but to look at particular central components of the view on their own.

In *Spreading the Word*, Blackburn asks us to consider an explicitly expressive language, E_{ex} , possessed of operators $B!(x)$ and $H!(x)$, for “Boo!” and “Hooray!”, respectively²⁰. The $B!(x)$ and $H!(x)$ operators are expressive; $B!(x)$ and $H!(x)$ express disapproval and approval, respectively, towards x . As users of E_{ex} will want to express attitudes towards other attitudes, and combinations of attitudes, and belief-attitude combinations, Blackburn sees fit to introduce bars ($|$) as a way of referring to the enclosed attitude or belief rather than expressing it; thus $|B!(x)|$ and $|H!(x)|$ refer to disapproval and approval towards x , respectively. He also introduces the semicolon ($;$) as a coupling operator; judging from his examples, it seems to mean something like “having the latter upon having the former.” Blackburn calls the resulting compound a *sensibility*, and suggests that we might helpfully think of it as a function from an input mental state to an output mental state.

20 See Blackburn (1984) pp. 193-196. Note that as of his (1988 [1993]), Blackburn includes a third basic attitude-expression, “ $T!(x)$ ”, for “Tolerate!”.

E_{ex} provides us with a way of constructing expressivist analyses of moral language. For each sentence of ordinary moral language, there is (or at least, there is supposed to be) a corresponding E_{ex} sentence to show us what attitude(s) the moral sentence expresses. Per Blackburn's translations²¹, "Lying is wrong" becomes $B!(\text{lying})$. Moral conditionals are identified as expressions of *higher-order attitudes*: that is, attitudes towards sensibilities. Thus "If lying is wrong, then getting your little brother to lie is wrong" expresses the same attitude as $H!(|B!(\text{lying});|B!(\text{getting little brother to lie}))$. By deploying higher-order attitudes, the *Spreading the Word* view can give an account both of the moral conditional itself and the role that simple moral sentences play when embedded in complex moral sentences.

Note that $H!(|B!(\text{lying});|B!(\text{getting little brother to lie}))$ does not appear to be of the form $A \rightarrow B$. In general, expressivist theories take one of two approaches, which we will call "univocal" and "mimicking." Univocal expressivism seeks an account of complex moral sentences that gives those sentences the logical forms suggested by their surface forms. Mimicking expressivism does not, instead attempting to find an alternative explanation for our use of our chosen surface forms.²² Thus, a view that uses higher-

21 Ibid p. 195

22 My distinction between mimicking and univocal expressivism is similar but *not* identical to the distinction between "modest quasi-realism" and "ambitious quasi-realism" drawn by Hale (1993). Modest quasi-realism seems similar to mimicking expressivism; Hale defines the modest quasi-realist as one for whom it is the case that

he will have accomplished all he set out to do, if he manages to explain how we can respectably and intelligibly talk and think as if there were moral truths and falsehoods-- presenting our attitudinal commitments in propositional style, with all that that entails, by way of propositional embeddings-- although there are in reality (as he conceives it) no such things. (Hale 1993 p. 340)

Ambitious quasi-realism, on the other hand, claims that there really is moral truth and falsity, not just an

order attitudes to explain complex moral sentences has two choices: it can offer an account of logical connectives in general that renders $H!(|B!(lying);|B!(getting\ little\ brother\ to\ lie)|)$ an instance of $A \rightarrow B$ (and E_{ex} sentences, in general, instances of the forms suggested by their English expressions), in this case it will be a univocal theory. Or it can grant that $H!(|B!(lying);|B!(getting\ little\ brother\ to\ lie)|)$ is not an instance of $A \rightarrow B$ (and other complex moral sentences may not have the forms that they appear to have-- at least that is what the kind of mimicking theory that we are currently considering, one without any quasi-realism, would say. A theory *with* quasi-realism, such as Blackburn's own, would be highly reluctant to make such a claim-- a situation that we will not investigate more closely until Chapter 6), and seek an explanation for why we nonetheless use "if lying is wrong, then getting your brother to lie is wrong" to express $|H!(|B!(lying);|B!(getting\ little\ brother\ to\ lie)|)|$. In that case it would be a mimicking theory.

Spreading the Word presents a mimicking theory. It is worth pointing out that trying to classify existing theories as univocal or mimicking is not as straightforward of a process as it might seem. As I specified, we have been examining the mimicking approach as it might appear without any quasi-realist elements. Blackburn's views,

imitation thereof. But although the modest/ambitious distinction seems similar to the mimicking/univocal distinction, there is a significant difference due to the fact that, per the provisos earlier in this chapter, the mimicking/univocal distinction is to be reckoned before any quasi-realist elements of the theory are factored in, while the modest/ambitious distinction is not. Hence, for example, Blackburn's position is, as I am about to argue, a mimicking theory, but is ambitiously rather than modestly quasi-realist (at least by Hale's reckoning. I think Hale may be oversimplifying matters. In fact, Blackburn repeatedly demonstrates apathy [e.g. Blackburn (1984) p. 257, Blackburn (1993c) p. 367] towards the question of whether he wishes to endorse modest or ambitious quasi-realism. But this stance is enough, I think, to demonstrate the difference between the mimicking/univocal and modest/ambitious distinctions).

however, *are* quasi-realist, and so we cannot expect them to look like the mimicking theory that we have just introduced. What we see in Blackburn, though, are explanations of the type that we would expect from a mimicking theory: explanations that center on the practical utility of the use of propositional form. He writes:

E_{ex} needs to become an instrument of serious, reflective, evaluative practice, able to express concern for improvements, clashes, implications, and coherence of attitudes. Now one way of doing this is to become like ordinary English. That is, it would invent a predicate answering to the attitude, and treat commitments as if they were judgements, and then use all the natural devices for debating truth.²³

Blackburn's hope is that we might “come to appreciate why it should be natural to treat expressions of attitude as if they were similar to ordinary judgements.”²⁴ The view is not, as in a theory of the univocal type, attempting to show that expressions of attitudes *are* similar to ordinary judgments, but rather that it might be profitable for us to use propositional form for the expression of attitudes despite those expressions not having the underlying logic of propositions. This focus suggests that we should read Blackburn as providing what I have called a mimicking view.

One aspect of the view that we will return to later is the account of what, exactly, we stand to gain from the adoption of propositional form. All mimicking theories require an account of this type. For now, we will sketch the idea; we will return to it for an in-depth treatment in Chapter 5.

According to the *Spreading the Word* version of Blackburn's mimicking

23 *ibid*

24 *Ibid* p. 196

expressivism, the purpose of using propositional form is to allow us to make use of the ordinary tools of reasoning for a purpose other than that of keeping track of actual logical relationships. Consider:

2.1e1 B!(lying)

2.1e2 H!(B!(lying);B!(getting little brother to lie))

2.1e3 B!(getting little brother to lie)

One who accepts 2.1e1-e2 seems compelled to accept 2.1e3: not because there would be any logical contradiction involved in his failure to do so, but because his own sensibility instructs him thusly. Should he choose instead to reject 2.1e3, he will have failed to choose attitudes in accordance with his own sensibilities, which would be an unhappy state for him (as he would lack self-approval, or in other cases even have self-disapproval) and a confusing one for anyone who must interact with him (as we would doubtless find it more difficult to understand the actions of one who acted in opposition to his own motivations). We can keep track of the fact that there is a problem with someone who accepts 2.1e1-e2 but not 2.1e3 by using propositions rather than straightforward E_{ex} to express the same attitudes. If we express the same attitudes as:

2.1e4 Lying is wrong.

2.1e5 If lying is wrong, then getting your little brother to lie is wrong.

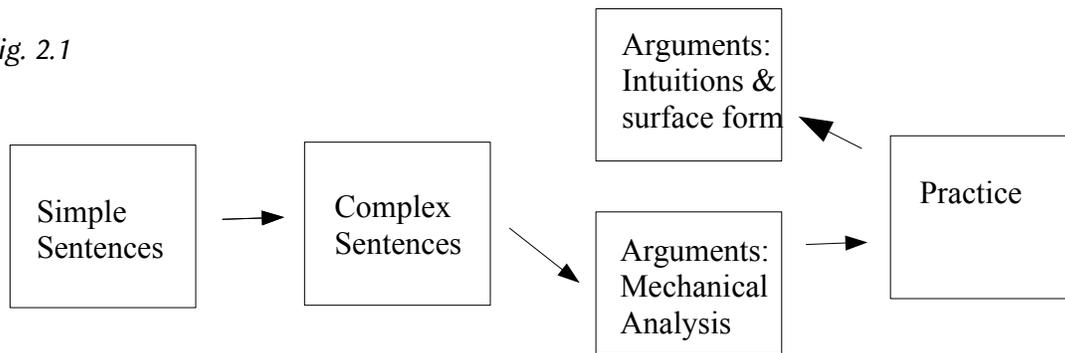
2.1e6 Getting your little brother to lie is wrong.

Then we will be able to see, using our existing devices of logical reasoning, that there is

a problem with the reasoning of someone who accepts the first two attitudes and yet rejects the third. Using E_{ex} , Blackburn can show us that when a combination of attitudes is happy or unhappy is tracked by the apparent logical relations between the propositions used to express those attitudes.

The overall order of explanation in the *Spreading the Word* theory is:

Fig. 2.1



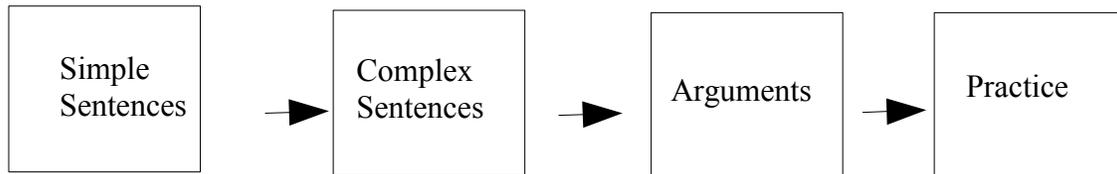
The account starts with expressivism, giving us an account of simple sentences. The idea of higher-order attitudes following the procedures of construction that are made explicit in E_{ex} then gives us a way of understanding complex moral sentences. The relationships between those sentences and between them and the simple sentences (once again, made clear by their E_{ex} analyses) govern the mechanism of moral arguments, in the way that we have just discussed. Understanding the goals and interests involved in the practice of moral discourse is simply a matter of generalizing on the previous, and the explanation both of our choice of surface form comes from the account of how the surface form serves that practice, which we have also just outlined (once again, for in-depth discussion, please see Chapter 5). As we will see in Chapter 6,

once we add quasi-realism to the theory, we also gain the ability to understand the role played by our (apparently realist) intuitions about the nature of moral argument.

At the beginning of this chapter, I described two overall approaches that an expressivist theory might take to solving the Frege-Geach problem; at the beginning of the section, I claimed that the *Spreading the Word* view was an example of the first of those approaches. Let us, after Blackburn (1988 [1993]), call the approach we have just outlined, where the theory builds upwards from simple sentences to complex sentences to arguments and then to practice, *slow-track expressivism*. We now have two dimensions which we can use to categorize the view: it is *slow-track* and it is *mimicking*.

Not all slow-track theories need take the mimicking approach. Taking the univocal approach would, after all, simplify the overall order of explanation to:

Fig 2.2



On a univocal account, moral sentences really do have the logic that they appear to have²⁵. The overall structure of any univocal slow-track theory is to start out with

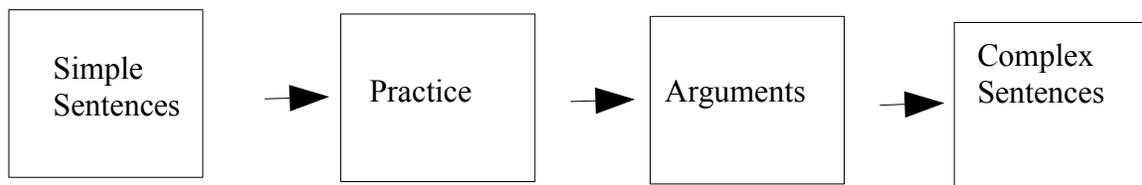
25 Although I will be focusing on mimicking forms of expressivism, it is worth noting that the univocal approach is not without its popularity. Indeed, Schroeder, in his (2010) and (2008a, b, c, d) seems to consider the univocal approach to be the single characteristic that distinguishes contemporary expressivism from early emotivist views, and as a result, the literature responding to Schroeder's treatment of Unwin (1999)'s problem of negation (see, e.g. Horgan and Timmons (2009), Sinclair (2011), Silk (2014), Skorupski (2012), Schwartz and Hom (2014)) seems to be proceeding under the assumption that the goal is to create a univocal expressivism. Mimicking expressivism, meanwhile, is not subject to the problem in the first place, as the very reasoning that mimicking expressivism uses

expressivism for simple sentences, and then to come up with an account of complex sentences *in general* (not just in the moral case). Once that has been accomplished, moral arguments need no further explanation, since they will simply be instances of their observed logical form; likewise, moral practice is just another normal part of our ordinary discursive habits. We may find ourselves concerned with moral matters in a way that is peculiar to morality, but no explanatory power needs to be drawn from that concern. Thus, the overall order of explanation is somewhat different for slow-track univocal and slow-track mimicking expressivisms. Nonetheless, the two have enough in common that they can be placed within the shared category of the slow track.

2.2 The Fast Track

Still following Blackburn (1988 [1993]), we will call the second of our two overall approaches the *fast-track* approach. The overall order of explanation for fast-track expressivism is:

Fig. 2.3



Fast-track expressivism, like the slow track, starts with a basic expressivist account of simple moral sentences as expressions of attitudes. Rather than trying to construct

against the Frege-Geach problem will, if successful, answer the negation case as well. At the very least, it seems worthwhile to be aware of both alternatives.

complex moral sentences out of these simple sentences, however, the fast track moves directly into talk about the practice of moral discourse as a whole-- its origins, its importance, etc. Included in the story will be some talk about moral disagreement and argument. The story, and in particular the account of disagreement, is then used as the explanation for the surface form. Complex sentences are said to be explained by their role in the larger practice, with special mention of their place in moral arguments. In fast-track theories, that each complex moral sentence behaves the way that it does within our practice of moral discourse is built into the analysis of that sentence.

Like the slow track, the fast track permits both univocal and mimicking approaches. I take the inferentialism of Chrisman (2008), (2010) to be an attempt at the former: if meaning *in general* is a matter of inferential role, and the expressive and descriptive (likewise the moral and nonmoral) are distinguished by the kind of roles they play, both expressive and descriptive statements will have their function within the practice of argument built in. On the other hand, a fast-track mimicking theory²⁶ embraces the more humble aim of showing that “the adoption of propositional form and style meets our need to share and discuss and dissent from attitudes or other stances”²⁷; where the slow-track form of the mimicking theory approached that aim from the ground up, attempting to form constructions out of attitudes as its way of demonstrating the suitability of propositional form, the fast-track form of the mimicking theory approaches it from the top down, attempting to show that the

26 As briefly outlined by Blackburn in the (1988 [1993]) paper where he outlines the slow-track/fast-track distinction.

27 Blackburn (1988[1993]) p. 185

practice of moral discourse, as a whole, could benefit from the use of propositional forms. It can then proceed to offer us an account of moral arguments told in terms of the function of those arguments within the wider realm of moral discourse, and of complex moral sentences in terms of their roles within moral arguments.

Although Blackburn introduces the idea of the fast-track approach, he actually favors a hybrid approach, combining elements of the slow and fast tracks. He claims that “the fast track can benefit from some of the security achieved on the slow, and the slow track can make use of some of the short cuts of the fast.”²⁸ By “security,” what Blackburn actually means is explanatory abundance. The slow-track, when used on its own, proceeds by developing an overtly expressive language, and then using it as an analysis of moral language. If this approach proves inadequate, we needn't necessarily discard the first step. There are other options for the *second* step that we might consider instead; that is, we might find some use for the expressive language other than as an analysis. Blackburn suggest that we might treat it as a “model showing why what we do is legitimate”²⁹. If we want a hybrid theory that is close to the fast-track but with additional explanatory resources, we could use the constructions of the slow track not as actual analyses of moral sentences, but rather as a way of demonstrating that attitudes can relate to one another in ways that might be worth tracking through the use of propositional form. Adding such a demonstration to what would otherwise be a fast-track theory can help make the theory's claims about practice more plausible.

28 Ibid p. 186

29 Ibid p. 197

Hybrid theories needn't follow the above model. It is also possible to build a hybrid theory that is primarily slow-track, but incorporates a fast-track element into its account of complex sentences or the attitudes that they express. Such a view might hold that the *Spreading the Word* theory, or something like it, is *almost* right, but that it omits some element which we can only characterize in terms of our expectations for arguments featuring moral sentences. It would claim that the constructions of the slow track tell us most of what we want to know about complex moral sentences, but that E_{ex} (or some counterpart thereto) and actual moral language are not quite expressing the same attitudes; the ones we express in moral language are like those in E_{ex} , but with an extra attribute that is to be understood, either in part or in whole, in terms of the inferential roles of their expressions. I think that Blackburn's view in his *Ruling Passions* (1998), which we will discuss in the next section, is most productively interpreted as a view of this type.

A third possibility for a hybrid theory would occupy a position in between these two alternatives, and thus by extension one somewhere near the middle of the slow-track/fast-track continuum. We might think that the slow track's search for a single compositional analysis that will be right across the board is futile, but that one or more slow track explanations might be right on a case by case level. That is to say, we might use slow track explanations to help us get a handle on some moral sentences in some contexts, but there is no *one* slow track explanation that is right all of the time. Rather, what justifies the similarity in surface form between compositionally different

expressions of attitudes is exactly the kind of functional consideration to which the fast track appeals. If the hybrid theory of this type is also a mimicking theory, it will hold that the binding consideration is practical in nature. It is a theory of this type that I ultimately hope to develop and defend, though it will take me the remainder of this piece to do so.

2.3 Hybrids

The view from Blackburn's *Ruling Passions* will serve as our example of a hybrid mimicking view.

As with our treatment of the *Spreading the Word* view, our presentation here will leave certain aspects for later. Propositional reflection-- the heart of the explanation that makes every Blackburnian mimicking view work-- will be covered in Chapter 5. Quasi-realism will be covered in Chapter 6. For now, here is an outline aimed at illustrating the overall structure of the theory.

Once again, the foundation of the theory is expressivism, but with a difference in what is expressed. On this view “for a subject *S* to think that *X* is good...is for *S* to value it”³⁰, and “when we assert values, we...voice our states of mind”³¹, where “Valuing something...is not to be understood as *describing* it”³². The “values” of the ruling passions view, though, are not any old attitudes; rather:

to hold a value is to have a relatively fixed attitude to some aspect of

30 Blackburn (1998) p. 50

31 *ibid*

32 *Ibid* p. 49

things, an attitude with which one identifies in the sense of being set to resist change, or set to feel pain when concerns are not met.³³

On the *Ruling Passions* view, we still express our attitudes in moral discourse, but the attitudes in question are part of a special category³⁴, rather than being the simple approvals and disapprovals expressed by the “hooray!” and “boo!” operators of *Spreading the Word*'s E_{ex} .

At first, the view of complex sentences in *Ruling Passions* seems similar to that of the older work. Blackburn claims that “By advancing disjunctions and conditionals, we avow...complex dispositional states”³⁵, where the dispositional states in question are ones to accept or reject certain combinations of beliefs and attitudes; this sounds very similar to the *Spreading the Word* view, only with “avow” substituted for “approve”. Once again, the idea is that an attitude “can be avowed, or...can be put forward without avowal, as a *topic* for discussion, or as an alternative”³⁶, which E_{ex} captured through the use of the vertical bar ($\bar{\quad}$). I suppose one might read the *Ruling Passions* view as holding that complex sentences express something like the combinations that were the objects of *Spreading the Word*'s higher-order attitudes, only without these being embedded in a higher order attitude, but as I cannot see any way of rendering this proposal intelligible, I will ignore it.³⁷

33 Ibid p. 68

34 When we return to the *Ruling Passions* view in Chapter 5, we will examine Blackburn's ideas about the distinctive nature of the attitudes expressed in moral discourse in greater detail, and will investigate exactly how the special attributes of such attitudes encourage us to express them in propositional form and work with them, in interpersonal coordination and intrapersonal thought, as if they were beliefs.

35 Ibid p. 71

36 *ibid*

37 A complicating factor here is Blackburn's own definition of “avowal”, per *ibid* p. 68, according to

Seeing the difference between *Spreading the Word* and *Ruling Passions* requires us to try and make sense out of an apparent *non sequitur*. Blackburn writes:

To avow anything of the form 'If p then q ' is to commit oneself to the combination 'Either not- p , or q ' and to be tied to that combination is to disavow the combination of p with not- q . Holding both together is therefore unintelligible. Logic is our way of codifying and keeping track of intelligible combinations of commitment.³⁸

At first glance, the reasoning here seems faulty. Why should it be unintelligible to hold a combination of attitudes that one disavows? After all, it is hardly uncommon, let alone unintelligible, to have attitudes that one does not want to have, or that one does not think one has, or that one does not wish to present oneself as having. There does not seem to be anything unintelligible or even unfamiliar about disavowing one's own combination of attitudes.

At the same time, it would be too hasty to see the claim here as a simple mistake. After all, Blackburn clearly *does* understand the idea of self-disapproval; indeed, he describes in detail (pp. 59-68) various kinds of inner conflict. He even accepts (p. 68) that one might genuinely hold a value, and yet act in ways that are at odds with that value. And yet, by extension, we ought to be able to, as Blackburn describes it, see ourselves as 'tied to a tree' (e.g. p. 71), and yet form other beliefs and attitudes in defiance of our bonds rather than in conformity to them. These ideas beg for reconciliation.

which "Avowal"...means that we express this state, make it public, or communicate it." It is something which we do with "practical states" (ibid). Unless "practical states" are themselves attitudes, though, it is not at all clear what it would be to express, make public, or communicate one.

38 Ibid p. 72

We can enact this reconciliation by starting with a smaller reconciliation.

Blackburn writes:

If we want to know...what is going on when we...put forward an attitude, we must look to the function of the indirect contexts in question. The key idea here is one of a functional structure of commitments that is isomorphic with or mirrored by the propositional structure that we use to express them.³⁹

He also says that, with regards to names for his position, “A full-dress title might be 'non-descriptive functionalism' or 'practical functionalism.’”⁴⁰ If we are to understand his position, we must find a way of incorporating a functionalist element into the account of embedding without lapsing into incoherence.

One way of accomplishing this goal is to build upon the fact that the *Ruling Passions* distinguishes between attitudes in ways that are independent from their intensity. *Ruling Passions* also contains the idea of “emotional ascent”, which we will return to in Chapter 5, but will sketch now. Attitudes, the view claims, can be distinguished not only by their intensity, but also functionally. Some attitudes are defined in part by such aspects as how unwilling we are to give them up, how we insist on others sharing them, or how we allow them to influence our actions. It is, on this view, possible to have two attitudes that are the same in valence and intensity, but which are nonetheless different attitudes, since they differ in what one has to do, or be disposed to do, in order to count as having them.

39 Ibid p. 71

40 Ibid p. 77

If some attitudes have subsequent actions built into their attribution conditions, then some combinations of practical states cannot constitute a coherent description of an agent, since their attribution conditions are incompatible. If, for instance, the attribution conditions of a certain attitude include that the attitude be held with a certain steadfastness, then an agent who had an attitude of that valence and severity but only for a short time would not have *that* attitude, but a different one-- one that, perhaps, is similar, but not a value, as Blackburn conceives of them. By the same token, an attitude, if we are to see it as a commitment and not a simple feeling, might have among its attribution conditions that we follow through with its endorsements and recommendations as to our combinations of attitudes. The E_{ex} sentence $H!(|B!(lying);|B!(getting\ little\ brother\ to\ lie)|)$, on this view, and the English “If lying is wrong, then getting your little brother to lie is wrong” do *not* express the same attitudes. The attitudes they express have the same object, valence, and intensity, but the latter imposes a functional demand that the former does not, namely, that one who accepts that lying is wrong also accept that getting your little brother to lie is wrong. We can describe, coherently, an agent as avowing $H!(|B!(lying);|B!(getting\ little\ brother\ to\ lie)|)$ and $B!(lying)$ while disavowing $B!(getting\ little\ brother\ to\ lie)$; such an agent would have an unhappy combination of attitudes, having failed to configure himself properly by the lights of his own sensibility, and this form of unhappiness is not only possible but all too frequent. Yet Blackburn seems to commit to the idea that we cannot say the same about someone who earnestly believes that lying is wrong and that if lying is

wrong, getting your little brother to lie is wrong, yet also earnestly denies that getting your little brother to lie is wrong. Such a person must have been mis-described; we must have been wrong in our understanding of his attitudes,

In integrating functional elements into its account of the attitudes expressed in moral discourse, the *Ruling Passions* view incorporates an element of the fast-track approach. Complex sentences cannot, on this view, be understood without understanding what those sentences require of us during moral argument. We can have different attitudes with the same composition in terms of simpler attitudes, but different functional roles. The difference in functional role is not captured by the composition of the attitudes in terms of simpler attitudes, but can only be understood by reading the attitude's place in practice off of the logic of the surface form of its expression. Drawing the role of complex moral sentences in arguments into the account of those sentences is a hallmark of the fast track.

2.4 Need for fast-track elements

The slow-track approach has much to recommend it. Because slow-track forms of expressivism employ a constructive procedure to understand complex moral sentences, they have a relatively easy time coming up with a satisfying account of such sentences, and of the role that the simple sentences making up the complex sentences play in determining the content of the latter. Having something to say about both simple and complex moral sentences, in turn, gives a theory a potential basis to appeal to in its

account of argument and practice. Adopting the slow-track order of explanation seems like it might give an expressivist the chance to answer the Frege-Geach problem in a direct and compelling way.

That said, there is also some reason to believe that a successful form of expressivism must include fast-track elements. Even in *Spreading the Word*, Blackburn sees reason for concern about the slow-track approach. He writes:

What I have done here is to explain how conditionals can be regarded as ways of following out implications, although it is not imperative that the commitments whose implications they trace have 'truth-conditions'. Now you might say: even if this can be done, hasn't the quasi-realist a very dreary task in front of him? For remember that the Frege point was entirely general: it could cite any unasserted context. So mightn't other arise which require separate and ingenious explanations, and the quasi-realist faced with an endless task?⁴¹

The worry that Blackburn perceives for the slow track is that the slow track only ever has accounts for those complex sentences for which it has specifically identified a construction procedure. When it comes to embeddings for which the slow-track has not identified a construction procedure, it has nothing to say about them; nor do we have any reason to assume that any such account is even possible, prior to it actually being found⁴².

By 1988 ("Attitudes and Contents"), Blackburn had found another reason to be skeptical about the slow-track approach:

I should admit not only that it threatens to look Ptolemaic but also that it seems not to correspond to any obvious cognitive processes we go through. It is not as though constructing (say) conditionals with evaluative components comes harder to us than constructing them with

41 Blackburn (1984) p. 195

42 An objection that would be repeated by Brighthouse (1990)

ordinary components, and this will need explanation.⁴³ Blackburn's worry seems to be that we, as language users, can use and work with *any* ordinary construction using evaluative language, without having to do any work to figure out how it translates into an expressive language. We use evaluative language in a way that is seamless with our use of non-evaluative language, and gladly accept that evaluative language will follow non-evaluative language in its logic. Yet no theory built on the slow track order of explanation can tell us what, if anything, could *guarantee* that any evaluative embedding will have the logic, or a parallel of the logic, of its surface form. We seem to act as if we do have such a guarantee, and the slow track cannot explain this, save perhaps by offering an exhaustive catalogue of expressive constructions, one that covers every possible form of embedding. No slow-track theory has done this, and, for the same reasons as gave rise to the problem, we have no guarantee that any such theory could do so.

Fast-track and hybrid theories need not suffer from these concerns. A fast-track theory explains the appropriateness of propositional form *in general*. For example, on the *Ruling Passions* view, the functional characterization of the attitudes involved in moral discourse is given by the propositional forms of their expressions, so the relevant attitudes will have the fact that they relate to one another and to beliefs in ways that parallel the apparent logic of their expressions in all contexts built into and thus guaranteed by their attribution conditions. On a view with this structure, we do not need to worry about whether we can come up with a construction from simple

43 Blackburn (1988 [1993]) p. 185

attitudes to serve as an adequate analysis of a given sentence where moral terms occur in an embedded context; rather, the surface grammar of the embedded context tells us what the sentence does within our practice of expressing and sharing attitudes, and we need only investigate why we would want a sentence with such a function.

On the other hand, creating a successful fast-track theory has its own challenges. Blackburn writes of the fast track that it “might seem like theft”⁴⁴: because of the fast track's order of explanation, it may seem to fall more naturally into the category of “vague gesturing” rather than an actual theory. And indeed, such a criticism could be well-founded, if the particular theory in question lacks development. The fast-track expressivist cannot be quite so quick as to claim that, as Blackburn (1988 [1993]) puts it, “[t]he adoption of propositional form and style needs our need to share and discuss and dissent from attitudes or other stances. It involves only philosophers in error, and little more need be said.”⁴⁵ Rather, a fast-track theory must devote itself to coming up with an account of moral dispute that satisfies explanatory demands from above and below: it must tell a story about an expressive practice that is independently plausible, and at the same time show why the sort of practice that it describes should be expected to incorporate or co-opt propositional form. Even once it has done so, though, it might still seem to be a bit of a cheat, insofar as it does not have neat formulae to supply as answers when asked about the meaning of complex moral sentences or the inner workings of moral arguments.

44 Ibid p. 185

45 ibid

It is hard to know how much to make of this worry. There is something emotionally appealing about the formulae of the slow track; seeing the expressivist account told in symbols rather than words, even if the meaning is the same, feels comforting (even though, in actuality, only notational work has been done). The fast-track account cannot do this, and that can create an uneasy feeling that the account has somehow been less specific or incomplete. But we should be careful to judge theories on their merits rather than their aesthetics; insofar as the feeling that a fast-track account is vague is traceable to the style with which such a theory must be described rather than its substance, it is not to be heeded.

There are, however, real difficulties in the vicinity. Schroeder (2010) gets it almost right when he criticizes inferential-commitment accounts⁴⁶, which are a species of fast-track theory, as follows:

since these theories are non-constructive, there is a natural sense in which they don't explain the semantic properties of complex sentences, as we originally set out to do, in the face of the Frege-Geach problem. To see why, compare the [cognitivist and fast-track] explanations of which someone who thinks that P and who also thinks that \sim P is undergoing the kind of clash of attitudes that we have been calling 'intrapersonal disagreement'...It is clear that the inferential-commitment theorist's explanation is no explanation at all. It merely helps itself to what it

46 See section 2.2. Now it is worth pointing out that it would be possible to create something that *looks* an awful lot like an inferential-commitment account without it being a fast-track theory. Failure in nuancing can cause an inferential-commitment expressivist account to degenerate into a slow-track approach. If “commitment” is, as per one of Schueler's (1988) worries, to be understood as a moral notion-- you ought to fulfill your commitments, or you do wrong if you fail to fulfill your commitments-- then defining a conditional as a commitment to believe the consequent upon believing the antecedent is just a roundabout way of defining it as a higher-order attitude towards combining or failing to combine the two attitudes when one has the antecedent attitude. This will not be the case for all expressivist inferential-commitment theories, though, as it depends on a particular notion of commitment which is not built into the definition of the theory-type. As for the question of what alternatives exist for “commitment”, we will return to it later/

should be trying to explain. Is there really such a mental state as $[\sim P]$?
Why does it disagree with $[P]$?⁴⁷

There are some issues with the details of the argument. When a fast-track theorist with any sophistication characterizes P and $\sim P$ as expressions of disagreeing states, it is not, as Schroeder seems to think, a specification drawn in a vacuum, but rather one that draws upon whatever account of disagreement is contained within the theory's description of the practice of moral discourse. Schroeder is looking for the explanation in the wrong place. For a fast-track theory, the explanation of the disagreement between P and $\sim P$ does not lie within anything that the theory has to say about the mental states expressed by P and $\sim P$; rather, a fast-track theory must have *already* provided an account of moral disagreement prior to offering any characterization of specific constructions like $\sim P$. If it has not done so, then it isn't a good fast-track theory in the first place; Schroeder's point would, indeed, work against it, but such a theory would fail anyway. True, the functional definitions used by the fast-track approach add far less to the explanatory story than the formulae of the slow-track approach, but one look at their respective order of operations should be sufficient to show that that is exactly as it should be.

The real trouble here, rather, is that the fast-track theory does not exist within a dialectic void; not only must its explanations succeed, but they must compete for plausibility with those of rival theories. Taking the fast track approach as a way of solving the Frege-Geach problem is of little use if doing so gives away the game to the

⁴⁷ Schroeder (2010) p. 141

realist. And in this vital regard, the fast-track theorist starts at a disadvantage, since he hasn't the realist's luxury of taking the position that the moral case works like the non-moral case, and nothing more needs to be said⁴⁸. If the fast-track theory is also univocal, it will need to motivate its general semantics with sufficient force as to render them independently plausible. Although we will return to the question of mimicking and univocal expressivism in the next section, for now, let us simply note that one must be incredibly ambitious to pursue univocal fast-track expressivism; mimicking fast-track expressivism is more humble, and thus opens itself to potential objections drawn from far fewer philosophical sub-disciplines. As for achieving mimicking expressivism via the fast track, it is not enough for the fast-track approach to have *a* mimicking story to tell; what it needs is a really *good* mimicking story. Thus, it is worth our attention to consider what kind of story would lend the most explanatory power to a mimicking account.

The central idea of a mimicking account, remember, is that the expressivist needn't achieve the surface form of moral utterances, only justify it.⁴⁹ It must, in other words, be prepared to explain why it is that, in Blackburn's words, "the adoption of propositional form and style meets our need to share and discuss and dissent from attitudes or other stances."⁵⁰ In this context, an objection from Schueler (1988) takes on the role of a constraint: if the kind of conflict or disagreement that is tracked in moral

48 A univocal theory could achieve the first half of the realist's luxury, though not the second-- it *must* say more if it is to justify its claim of univocality. But the univocal approach also carries a significant additional burden insofar as it has claims that it must justify well outside the moral sphere.

49 "Notice that, whichever track we favor, the point is to *earn* our right to propositional forms": Blackburn (1988 [1993]) p. 186

50 Ibid p. 185

argument is too far removed at the practical level from inconsistency, the explanation will fail. If there is too much difference between how we work with moral sentences when engaged in rational argument and how we would expect expressive practice to go according to the mechanics set out by a particular theory, then it becomes implausible to suggest that such an expressive practice could be expected to co-opt propositional form and the accompanying logic. Thus if, for instance, a theory says that we are just expressing attitudes and more attitudes (albeit higher-order ones), and that the logic that we apply to those expressions of attitudes is actually tracking whether combinations of attitudes are happy or unhappy (that is, self-approving or self-disapproving), it faces a tough question in whether such an ill-fitting expression could really be explained away by some purported practical benefit.

Observe that the (slow-track) *Spreading the Word* view, Schueler's original target, indeed seems to struggle on this front. Some further attitude completing or failing to complete a combination of attitudes of which you approve or disapprove might be an important relationship to keep track of, but the strategy of keeping track of it by expressing attitudes with propositions and higher-order attitudes with conditionals seems to be overkill. Schueler writes:

[W]ithout some account of how there is something untoward or mistaken in a mere clash of attitudes..Blackburn cannot claim even to have shown that, on the antirealism theory he is advocating, a person who asserts the premises and denies the conclusion of the argument...has made any mistake at all, let alone any sort of logical mistake in the usual sense.⁵¹
The problem at hand with tracking unhappy combinations of attitudes using

51 Schueler (1988) p. 500

propositional forms is not that one cannot do so-- not that there is reason to doubt that there can be an isomorphism between the two (though, as we will see in Chapter 5, there is reason to suppose that the two *do* come apart, as well as, as argued earlier, *no* reason to suppose that the isomorphism can be maintained across all possible embeddings). Rather, the issue is that the propositional form leads us to thinking in certain ways about moral argument: for instance, that one who accepts the premises and yet denies the conclusion of a moral argument that we typically consider valid has thereby made a mistake in his reasoning; that such a person is to be condemned as irrational; or that, having been led to accept the premises of such an argument, it is incumbent upon us to accept the conclusion (or reverse our acceptance of one of the premises) regardless of how we might feel about that particular moral belief. Of course, a defender of Blackburn's 1984 position would claim that we are not misled by any of these intuitions; he might claim that we ought to regard one with an unhappy combination of attitudes as mistaken in his reasoning and irrational, and that regarding such people in that way is part of an important practice of interpersonal and intrapersonal coordination. But this reply is not convincing; while having a combination of attitudes of which you yourself disapprove might indeed cause you some grief, it hardly seems worth thinking of as a logical mistake. Indeed, Schueler points out that this particular form of inner conflict is neither uncommon nor particularly "untoward or fishy"⁵². To track such a thing with apparent logical inconsistency would be to treat it with a seriousness that is wholly out of proportion to

52 *ibid*

its actual importance.

For a mimicking theory to be viable, it must do better than this; it must point to some conflict worth tracking with logical contradictions. When it comes to determining what would make a conflict worth tracking with logical contradictions, it seems unlikely that there is anything like a concrete rule to which we might appeal. Theories of the sort under discussion live or die based on the plausibility of their explanations, and while there are some ways that a theory could easily lose its plausibility (contradicting itself, for instance), if a theory is at least reasonable, plausibility or implausibility is harder to make definite than proof or disproof. Nonetheless, we do, at least, tend to have some inkling of when an expressivist theory gains or loses out on the basis of its account of the foundation of moral “inconsistency.” If a conflict strikes us as the sort of thing that one could be in without a serious flaw in one's thinking, then it is unsuitable as the foundation of a fast-track theory, because it would not be worth tracking with logical contradictions. On the other hand, a conflict that *is* a serious flaw in thinking-- and it will be up to individual theories to find a conflict and explain *how* it is a serious flaw *in thinking*, without collapsing into cognitivism-- could play the part. Only if a suitable conflict can be found will a mimicking theory be able to succeed.

In Chapter 5, when we turn our attention to the subject of propositional reflection, we will return to the issue of making our use of propositional forms to express attitudes make sense. For now, however, let us settle for noting that the fast

track seems to have an advantage over the slow track in answering Schueler's objection⁵³. A slow track mimicking theory must start by constructing complex attitudes, and then find a relationship between its complex and simple attitudes that is worth tracking with apparent logical disagreement. A fast track mimicking theory, on the other hand, can attempt to incorporate a more appropriate relationship into its account of the relevant attitudes or of the practice of moral discourse more generally. The *Ruling Passions* theory, for instance, cites conflicting attribution conditions of attitudes, which would indeed seem to involve a mistake in thinking if not kept straight. While we cannot rule out the slow-track approach for this reason alone, in the absence of any slow-track theory that can show us how such a theory might come up with a satisfactory relationship, it seems reasonable to be pessimistic about the slow track's chances for success in this regard.

The arguments of this section attempted to motivate adopting the fast track approach, or a hybrid theory with fast-track elements, over the slow track. One of the arguments was that the fast track permitted the development of better mimicking theories than the slow track, since the fast track is able to appeal to relationships that it makes more sense to track with propositional form than any of the compositional relationships available to the slow track. For this advantage to matter, there must be some reason to think that mimicking theories are correct. The next section attempts to demonstrate that this is so.

53 Historically, this stands to reason, since Blackburn (1988 [1993]) introduces the fast track in large part as an answer to Schueler's critique of his earlier view.

2.5 Preferring mimicking expressivism

Mimicking theories, compared to univocal theories, have more work to do. They must answer the question of how the mimickry they propose is occurring is appropriate, and defend against the challenge that their account amounts to an (untenable) error theory. Nonetheless, we have reasons to think that mimicking theories are more promising. First, as we have already mentioned, univocal theories are incredibly ambitious; they make remarkable claims in the philosophy of language which they will have to be prepared to defend. Mimicking theories are far more humble; they can, if so desired, let our general theory of language go whichever way it pleases, and then, later in the order of explanation, provide their account of why it is nonetheless useful for us to express our attitudes as we do. Second, mimicking theories can use quasi-realism to adopt many of the advantages of univocal quasi-realism. Third, when equipped with the kind of quasi-realism that I will seek to defend in Chapter 6, mimicking expressivism can retain the ability to resort to outright denials of realist meta-ethical claims, which might be convenient for solving various problems.

We have already covered the first of these three points, so let us move directly into the second. This is the idea that the costs of mimicking expressivism can be mitigated through the use of quasi-realism. For example, several authors, including Van Roojen (1996), Schroeder (2010), Miller (2003), and Zangwill (1992) criticize the *Spreading the Word* view for conflating pragmatic (or “Moorean”) inconsistency with

actual inconsistency⁵⁴. Such a criticism represents a failure to distinguish between mimicking expressivism, which Blackburn is pursuing, and univocal expressivism, which is Schroeder's own favored strategy. Nonetheless, one might think that this weakness in the foundation of the objection does not wholly collapse it. For a mimicking theory makes sense insofar as the mechanisms that it proposes would be practical; if those mechanisms would lead us into error, given that we generally regard being in error as bad, we would expect ourselves to resist them. Likewise, we might wonder about the persistence of the mechanisms. There might very well be good reason for us to adopt a mimicking practice, but we must marvel at the duration of our pre-Blackburnian innocence-- that is, that we have managed to mislead ourselves for so long in order to achieve practical benefits.

While I find myself with little inclination to give weight to these concerns (I see no reason to suppose that the human capacity for self-deception has such meager

54 Schueler (1988)'s objection follows similar lines, but, as we have already discussed, Schueler raises the problem in a way that demonstrates awareness of the actual nature of Blackburn's theory. Hale (1986) is sometimes read as a member of this group, but also seems to understand that what Blackburn needs for "inconsistency" is far less demanding than is thought by those who misunderstand the univocal-mimicking distinction.

Wright (1988b) also maintains an atypical understanding of the problem, according to which the higher-order attitudes view fails to keep proper count of the number of mistakes made by someone who does not follow a moral *modus ponens*: there is the moral mistake, that is, of not following through with the principle expressed by the conditional, and at the same time a second, distinct *logical* mistake. Like Hale and Schueler, Wright recognizes that Blackburn's aim (as it would be understood in a hypothetical version of his theory *sans* quasi-realism) is not to satisfy our logical intuitions, but merely to vindicate them; if, however, the higher-order attitudes account mis-counts problems, we will be left wondering how to incorporate our idea of the proper number of problems into the mimicking story. There is, thus, a real problem for mimicking expressivism in Wright's objection. I am not sure how severe a problem it is, because I do not find the idea of a second mistake terribly compelling (would it not be more plausible to say that there is *one* mistake, that of reasoning badly, regardless of whether the maltreated argument descriptive or expressive in nature?), but, just in case there is a genuine problem here (or in the vicinity), note that the arguments to come in this section about the usefulness of quasi-realism to a mimicking theory should solve it.

bounds, or indeed any bounds whatsoever), I suppose others⁵⁵ could be forgiven for finding them more troubling. Although the mimicking expressivist will do his best to find strong reasons for our use of propositional form, our intuitions about the forms of complex moral sentences and the arguments into which they enter weigh heavily for the other side. The mimicking theory *does* try to incorporate those intuitions; our seamless treatment of language is, for the mimicking expressivist, an important part of our practice of coordinating attitudes. But the mimicking explanation is far more complex, and thus less desirable as a piece of theory, than the explanation that those intuitions are simply correct.

Quasi-realism is a tool by which the expressivist comes to adopt meta-ethical claims from realism. It can be of use to the mimicking expressivist in that it allows mimicking expressivism to bring itself into conformity with our linguistic intuitions while both a) maintaining its account of the practicality of adopting propositional form for the expression of attitudes, and b) remaining harmless to our overall theory of language. The quasi-realist mimicking expressivist thus keeps the advantages of mimicking expressivism while also gaining the right, otherwise only available to univocal expressivism, to claim to have actually solved the Frege-Geach problem rather than merely finding a less-unpleasant-than-usual way of biting the bullet.

A mimicking expressivist who is also quasi-realist would, in all likelihood, no longer call himself a mimicking expressivist⁵⁶. This is so because, once equipped with

55 Harcourt (2005) is an example.

56 Blackburn, who holds a position of this type, seems to find it difficult to know what to call himself; he seems to like the term “projectivist” in his earlier works, only to reject it in *Ruling Passions*.

quasi-realism, the mimicking expressivist will no longer claim that complex moral sentence or moral arguments fail to have the logical form that they appear to have.⁵⁷ Rather, he will claim that they *do* in fact have such forms-- only, this further claim is itself to be understood in some way that an actual realist, or, for that matter, a univocal expressivist, might find surprising.⁵⁸ Exactly what that way is will require sustained examination, which will occur in Chapter 6. For the time being, let us assume that a successful method for the assimilation of realist claims will be found. Given that assumption, a mimicking expressivist would not have a theory on which any mimicking occurred. The practical account of our assumption of propositional form for expressing attitudes would remain the same, but the outcome of the account would not be a reason for expressing attitudes with a surface form that does not reflect their logical form. The mimicking story tells us how our expressions came about, but does not, according to the quasi-realist, disqualify those expressions from being genuine instances of their surface forms. Rather, it would be a reason for the surface form, one which might be combined with further reasons (generated by quasi-realism) to justify

“Mimicking expressivism” suffers from the same problem that Blackburn points out for “projectivism”-- it suggests that there is some kind of fakery going on, which is not what the quasi-realist holds (Blackburn 1998 p. 77). I am not sure whether appropriateness or consistency should take priority in naming, but, for now, let us go with consistency (and thus continue to use “mimicking expressivism”-- if a jellyfish can be called “jellyfish” without being a fish, then mimicking expressivism can be called “mimicking expressivism” even when it no longer counts its mechanism as mere mimicking.

57 In Hale’s (1993) terminology, such an expressivist ceases to be “modest” and instead becomes “ambitious.”

58 Assuming only a single application of quasi-realism. The quasi-realist procedure is, in fact, recursive; the quasi-realist might wish to claim that not only do complex moral sentences and moral arguments have their apparent logical forms, but that they have those forms in the usual way-- where what it takes to have those forms in the usual way is then to be subjected to analysis, assuming the quasi-realist does not wish to add yet *another* layer of quasi-realism.

the claim that those expressions really *do* have their apparent forms.

Since mimicking expressivism with quasi-realism says that we are not in error about the forms of complex moral sentences or arguments, it is not *inherently* subject to the objection from error. It may be subject to a lesser version of the objection targeting the quasi-realism itself-- that is, that quasi-realism claims that we are surprisingly in error about the content of many philosophical statements. But such an error would be far less offensive, as it deals with a more exotic species of sentence, which lessens the error, since the error is both less surprising and less often made. Additional applications of quasi-realism⁵⁹ can further lessen the error, until it is completely inoffensive.

This defense is, for the time being, largely speculative; it depends on what is (for now) an assumption that a quasi-realist maneuver capable of delivering the stated advantages can be found. I say “largely” because there is a way that the defense can help the mimicking expressivist, albeit only slightly, even without the arguments of Chapter 6. This is simply to note that the mimicking expressivist need not be nearly so brutal in the formulation of his view as I have been. Even without quasi-realism, he needn’t pursue contrast with univocal expressivism through outright denials. Rather, the mimicking expressivist can simply provide his account of the usefulness of propositional form and leave it at that. It is up to the objector to show that that account puts us in error, and that “us” is more than just a few philosophers. This attitude towards the issue is exactly that taken by Blackburn when he originally (1988)

⁵⁹ See previous footnote.

introduces the fast-track approach. We do not need a slow-track analysis, Blackburn's argument goes, to show that what we do is legitimate, because what error there is involved in adopting a mimicking-type theory is minimal and in the main confined to philosophers. We *could* invoke slow-track-style explanations to try and bolster the legitimacy of the practice-- hence hybrid theories-- but these need be no more than secondary explanations, that is, models, or perhaps speculative analyses specific to a particular instance (towards the end of Chapter 5, we will have a chance to investigate this kind of theory in more detail). There need not be anything more because the overall claims of the fast-track mimicking theory are undemanding (except, perhaps, for philosophers). Ordinary folk do not, need not, and perhaps even *should* not distinguish between the kind of mimicking the mimicking expressivist posits and actuality.

Let us turn now to the third point in favor of mimicking expressivism. Equipped with the proper kind of quasi-realism-- such as the kind that I attempt to develop in Chapter 6-- mimicking expressivism possesses the ability to let us look behind the mimicry if necessary. That is to say, mimicking expressivism is free to acknowledge differences between moral and non-moral language, so long as those differences are apparent to us; such differences can only reinforce the impression that moral language is a creature of a different sort that is merely imitating descriptive language. As Blackburn puts it in his (1993c):

We can *earn* our right to use the logic and the propositional forms that

we do, from the anti-realist starting point. Seen like this, a quasi-realist may or may not regard the forms in question as perfectly propositional. He could be receptive to signs, embedded in normal discourse, that we do not always think and speak as if the sentences in question were just like others.⁶⁰

A univocal expressivist⁶¹, on the other hand, has no such freedom. If the univocal expressivist shows that moral language has the forms that it seems to have, any cases that suggest moral language is actually something different will then require an explanation, one other than expressivism itself. The mimicking theory retains the ability to appeal to expressivism in explaining cases that are awkward for realism. The univocal theory cannot appeal to expressivism to explain any such differences because it argues that expressions of attitudes are no different from any other proposition; they have the same forms, and thus should be expected to behave in the same ways. The mimicking theory *can* do so, simply because a mimickry needn't be perfect⁶².

Are there any such cases? Blackburn does not provide an example in his (1993c), and we cannot dwell too long on this topic without going well outside of the intended scope of this piece. For now, let us content ourselves with a mere suspicion. One might suspect, especially if one takes an interest in the so-called “moral twin earth” problem

60 Blackburn (1993c) p. 366

61 Or, for that matter, a mimicking expressivist with a different kind of quasi-realism from the one I endorse.

62 I should point out that this statement is not in conflict with my earlier remarks about needing a very compelling mimicking story. Mimicking expressivism is about explaining why we do what we do; insofar as a purported mimicking practice would lead us into confusion, we should be skeptical about such a practice. In the kind of case that I am now discussing, we are not led into error, but *see* the difference between the moral and the non-moral.

of Horgan and Timmons (1992)⁶³ (itself an elaboration on an idea from Hare⁶⁴) that what it takes to disagree with a moral statement is not quite the same as what it takes to disagree with a descriptive statement. Simple expressivism, such as that of Ayer (1936) and especially Stevenson (1944), can explain this difference as a result of the fact that moral disagreement itself is something different from ordinary descriptive disagreement; mimicking expressivism can do likewise. But, for the univocal expressivist, moral disagreement is “genuine” disagreement of the type observed in descriptive discourse⁶⁵, so he will need to carry out additional theorizing if he is to explain the difference⁶⁶.

Whether or not there is a genuine reason to favor mimicking expressivism in this

63 See also Blackburn (1991 [1993]). For discussion see *e.g.* Merli (2002), (2008); Sayre-McCord (1997); Gert (2006).

64 *e.g.* Hare (1952 [1964])

65 In fact, this is regarded as a desideratum by Schroeder (2008a, b, c, d), (2010).

66 I should point out that the univocal theorist struggles with explaining disagreement in general. Schroeder, our example of a univocal expressivist, argues in his (2008b) that genuine disagreement requires one and the same inconsistency-transmitting attitude being present towards inconsistent contents: for example, “there is a cat in the room” and “there is not a cat in the room” express the same attitude (belief) towards inconsistent contents (cat in room, not cat in room). In the moral case, the attitude is one of “being for”, which seems identical with Blackburn’s $H!(x)$. Yet, to handle the full range of moral attitudes, Schroeder must introduce substantive proposals for one to be for when saying that something is good or bad, or right or wrong (he suggests “blaming for” in the negative case; presumably, some other substantive proposal is necessary in the positive case, to explain what one is for when one denies that something is good-- perhaps such a person would be for not praising for the act in question); this move is problematic because it seems that I can be for or against something without thereby being for one particular proposal for what to do about it.

If we were to flesh out the argument for mimicking expressivism being developed here, the dialectic would start with the Moral Twin Earth problem to motivate the search for an alternative to descriptive inconsistency for the account of moral disagreement. The obvious alternative is a clash of different attitudes, which is how the emotivists explained moral disagreement. The univocal expressivist closes this option off to himself because it is dissimilar to descriptive disagreement (even though this was why one might have considered it as a solution to the problem in the first place), which renders it incompatible with his overall strategy. Yet, he lacks a compelling alternative; furthermore, even if a viable alternative could be found, it would lack the elegant simplicity of the emotivist solution. The solution is to reject univocal expressivism in favor of mimicking expressivism, which can and will use the same solution as the emotivist.

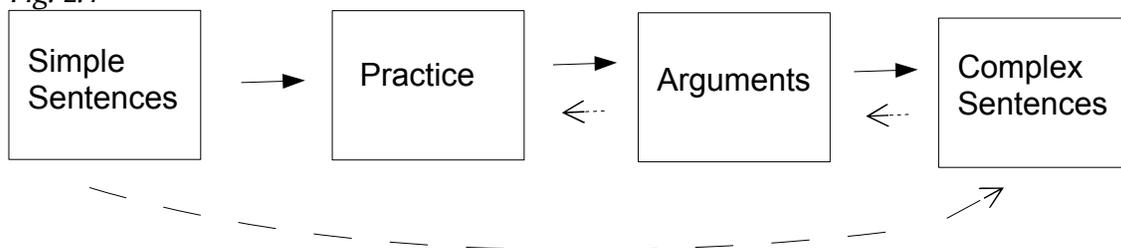
vicinity, the fact remains that mimicking expressivism, since it has the ability to drop the act, has a tool in its toolkit which univocal expressivism lacks. I suspect that this tool will prove useful. If it does not, then the mimicking theory loses nothing for having it, and is still preferable for other reasons, in particular its modesty with regards to our theory of language.

One interesting consequence of the mimicking theory's approach to disagreement is that it brings the mimicking theory closer to the old emotivists than we might have suspected. Although the theories of Ayer and Stevenson are now quite out of date, they might nonetheless be of use to us because of this similarity, while post-Geach work that attempts a univocal answer to Geach's problem is less likely to be useful. Naturally, Ayer and Stevenson's theories will not do on their own, but they can be helpful in a different way.

Earlier, I introduced the idea of a hybrid theory, one which follows Blackburn's preference for a combination of slow-track and fast-track elements. Such a theory, unlike a pure fast-track theory, has things to say about complex moral sentences and their relations to their parts besides what it determines from their relative places in our practice of reasoning. Unlike a *Spreading the Word*-style, slow-track theory, the hybrid theory stops well short of trying to build an entire analysis from simple attitudes and rules for assembling them into more complex ones. Instead, it builds on a smaller scale. It uses constructions to illustrate the ways that attitudes can relate to one another, and to give us at least some insight (though not necessarily a complete understanding) into

the nature of at least some (not necessarily all) of the complex attitudes involved in morality. These insights are to be completed and bound together by fast-track elements. By understanding our overall aims in using propositional form for expressing at least certain kinds of attitudes, and then understanding the role of complex moral sentences in our pursuit of those aims, we can more precisely and universally account for the behavior of those sentences. If we use dashed arrows to represent the secondary explanations provided by the slow-track elements, a hybrid theory's order of explanation can be pictured as:

Fig. 2.4



The theories of old emotivism, despite their incompleteness, can contribute to a theory with the above structure. We know that attitudes can relate to one another in a number of different ways; in fact, we have already seen multiple options simply by looking at different iterations of Blackburn's view. It would be tempting, especially, I think, for philosophers, to base our understanding of the interactions relevant to the practice of moral argument on a single relationship that attitudes can bear towards one another. Towards the end of Chapter 5, I will argue that doing so is a mistake. Rather, it is important that we understand the incredible variety of relationships that can hold

between attitudes, a variety that we use propositional form to help us manage. A theory that did not work on its own, but which makes use of accurate observations about our attitudes, can add to our list of interactions. At the same time, it can contribute to our list of secondary explanations; it can serve as a model, and as an occasional source of insight into particular cases.

In the next two chapters, therefore, we will turn our attention to Ayer and Stevenson.

Chapter 3. Ayer

I would like to begin my discussion of Ayer with an examination of a historical oddity. Looking back upon his view of ethics in his 1983 Whidden lectures, A. J. Ayer took a moment to consider the problem of embedding-- also known as the “Frege-Geach problem”. We know that Ayer was not unwilling to modify or even reject the views that he expounded in *Language, Truth, and Logic*; indeed, in a mid-1970s interview with Bryan Magee, he commented of the *LTL* view that its “most important defect was that nearly all of it was false,”⁶⁷ and in a 1989 interview with Ted Honderich, he goes as far as to call *LTL*'s philosophy of mind “incoherent”⁶⁸. But the Ayer of 1983 was largely unrepentant on the matter of ethics, a position that he would continue to hold; for that matter, in the 1989 interview, Ayer declares outright: “As for morals, I hold my old view”⁶⁹ As Ayer died two months after that interview, we can say without too much reservation that he kept his view on ethics (though perhaps, as we will see, with certain refinements upon the original formulation) to the end. On the subject of embedding, Ayer had this to say:

I think that this objection is less formidable than it may appear at first sight. Once it is admitted, as in the case of 'this is good', that an emotive expression can be put into an assertoric form, there seems no sufficient reason why the disguise should not be adapted to cover one's emotive

67 Magee 1979 p. 131

68 Honderich 1989, in Griffiths (ed.) 1991 p. 216

69 Honderich 1989, in Griffiths (ed.) 1991 p. 215

reaction to hypothetical instances. As for the deductive inference, it has, indeed, to be assumed that our attitudes are consistent, but thereafter the inference is simply secured by the transitivity of the comparative form.⁷⁰

We will, in due time, attempt to draw out the nature of the reply being given here. For now, though, let us focus not on the substance of the reply, but on the overall attitude to the problem expressed therein. We might take as some further indication of Ayer's appraisal of the problem the comment directly following his discussion, where he labels the objection that he fails to capture ordinary usage as “more serious” than the worry from embedding.⁷¹

We might wonder at Ayer's attitude here. For Geach's statements of the problem-- especially in Geach (1960) and Geach (1965); Geach (1958) takes Hare's prescriptivism as its direct target-- seems like it is aimed directly at him. Yet it seems uncharitable, or at least overly hasty given Ayer's willingness to discard his earlier views, to write off his reaction to the embedding problem as nothing more than a display of stubbornness. So we might wonder: is there something in Ayer's view that might give it some resistance to the attacks of the embedding problem? Or is there some attempt at such that fails, but might nonetheless underlie Ayer's confidence?

3.1 Ayer's emotivism, version 1

LTL gives us a starting point for developing Ayer's theory. Ayer explains the function of ethical statements as follows:

70 Ayer 1984 pp. 29-30

71 Ayer 1984 p. 30

Moral judgements do not say anything. They are pure expressions of feeling and as such do not come under the category of truth and falsehood.⁷²

More specifically:

If now I...say, 'Stealing money is wrong,' I produce a sentence which has no factual meaning-- that is, expresses no proposition which can be either true or false. It is as if I had written 'Stealing Money!!'-- where the shape and thickness of the exclamation marks show, by a suitable convention, that a special sort of moral disapproval is the feeling which is being expressed.⁷³

We might wonder what to make of this “special sort of disapproval.” Indeed, an uncharitable reading of the idea could easily predestine our reconstruction to commit suicide by circularity, as would be the case if we identified the moral attitudes by their role in moral discourse while simultaneously defining the function of such discourse as the expression of just those attitudes.⁷⁴ Even if we avoid circularity, it is hard to see how to pick out the appropriate feelings⁷⁵.

But I am not sure that this notion of a distinctive moral attitude need be anything more than a descriptive shortcut on Ayer's part. He does say that the notion he is after is (in part) “the different feelings [that ethical terms] are ordinarily taken to express”⁷⁶, and that we are talking about the expression of “certain feelings in the speaker”⁷⁷ in a particular situation. Seeing as Ayer's view posits the expression of a variety of different attitudes, the objection above seems to have a problem with

72 Ayer 1936 [1952] p. 108

73 Ibid p. 107

74 See Miller 2003 p. 44

75 See Wright 1988 pp. 11-12, and *ibid* p. 45. Alston 1968 also considers these issues.

76 Ayer 1936 [1952] p. 108

77 Ibid p. 107

number; it would be odd to think that Ayer would not only posit distinct moral feelings, but attribute to each occurrence of them an arbitrary and unspecific multiplicity. And the role of these emotions in moral discourse is not meant to be our only handle on them; Ayer's reference to "a peculiar tone of horror"⁷⁸ tells us that these are, at least to some degree, familiar.⁷⁹

A more charitable interpretation would go as follows. In moral discourse, we express only ordinary feelings, but not one at a time; rather, we give voice to bundles consisting of many different feelings, the co-occurrences of which in certain combinations and ratios is associated with morality. And so a "moral disapproval" is not some unanalysable, fundamentally distinctive feeling, but something like ordinary disapproval tinged with outrage and flavored with horror, with an undercurrent of discomfort (to give one of a great many possibilities, no one of which need be the whole story). Of course, we are still left wondering how to define these emotions, but

78 *ibid*

79 I should also point out that, whatever conclusion we may come to about the nature of emotions in Ayer's emotivism, the theory is not so simple as to rely *entirely* on them. Here we should take a moment to address the criticism of Ayer by Hare mentioned in Chapter 1. Hare alleges that Ayer disrespects the difference in logic between telling someone to perform an action and merely getting him to do it (see Hare (1952) p. 13, and for support his "Freedom of the Will", in Hampshire et al. (1951)). This criticism would be well-founded if Ayer claimed that the prescriptive function of moral terms was derived from the emotive function: that is, if we read Ayer's claim (*LTL* p. 108) of moral sentences having "the effect of commands" with the emphasis on effect, and do not read much into the structure of the following sentence-- which claims that a moral sentence "may be regarded as **both** the expression of...feeling...**and** as the expression of the command" (*ibid*, bold added). But in light of Ayer's remarks immediately following, that does not seem like the right reading. Rather, Ayer holds that the emotive and prescriptive are separate dimensions in the determination of the so-called "meaning" of ethical terms (more on *that* issue in a moment). He claims that, for instance, "it is your duty to tell the truth" and "you ought to tell the truth", sentences both of which one might categorize (following Ayer, *ibid*) as expressing a pro-attitude towards truthfulness, are to be differentiated from one another insofar as they issue commands with a different degree of forcefulness. And so, in Ayer's theory, the commanding effect of moral sentences is not to be described in terms of getting people to do things by expressing your emotions or attitudes, but as a dedicated function of moral language in addition to the expression of attitudes.

at least they have the advantage of familiarity.

Now that we have at least commented on the emotional component of the emotive theory, we must turn our attention to the question of what, exactly, the theory is meant to *be*. We must resolve the appearance of tension between, on one hand, the passage quoted above claiming that “moral judgements do not say anything” (and its allies, e.g. that the use of such terms “adds nothing to the literal meaning of [a] sentence.”⁸⁰, that “ethical concepts are pseudo-concepts and therefore unanalysable”⁸¹, that “sentences which simply express moral judgements do not say anything”⁸²), and on the other hand, such comments as:

In fact we may define the meaning of the various ethical words in terms both of the different feelings they are ordinarily taken to express, and also the different responses which they are calculated to provoke.⁸³

(Though note Ayer's scare quotes in the sentence immediately preceding:

And thus the “meaning” of the word “good,” in its ethical usage, is differentiated from that of the word “duty” or the word “ought”.⁸⁴

Though they are dropped in the next sentence, it is interesting that they occur here.)

In other words, we must answer the question of whether or not Ayer's emotivism is meant as a semantic theory-- if he is proposing an analysis whereby ethical terms are to be treated as having an “emotive meaning”, separate from yet equal to descriptive meaning, or if the proposals of emotivism are to be understood in

80 *ibid*

81 *Ibid* p. 112

82 *Ibid* p. 108

83 *Ibid* p. 108

84 *ibid*

some other way. Fortunately, *LTL* provides us with the necessary resources to answer this issue. The first consideration to bear in mind is the role of the emotive theory within the greater project of *LTL*; it exists to answer a certain challenge to Ayer's verificationism:

There is still one objection to be met before we can claim to have justified our view that all synthetic propositions are empirical hypotheses. This objection is based on the common supposition that our speculative knowledge is of two distinct kinds-- that which relates to questions of empirical fact, and that which relates to questions of value. It will be said that "statements of value" are genuine synthetic propositions, but that they cannot with any show of justice be represented as hypotheses, which are used to predict the course of our sensations; and, accordingly, that the existence of ethics and aesthetics as branches of speculative knowledge presents an insuperable objection to our radical empiricist thesis.⁸⁵

Ayer's tactic for responding to this challenge is to attempt to provide a theory of ethics-- the emotive theory-- that is compatible with the meaninglessness of ethical statements, as is suggested by the criterion of verification, without thereby being eliminativist. Any such theory would benefit from the ability to explain ordinary attributions of meaningfulness. Ayer thus clarifies in his introduction to the second edition:

In putting forward the principle of verification as a criterion of meaning, I do not overlook the fact that the word "meaning" is commonly used in a variety of senses, and I do not wish to deny that in some of these sense a statement may properly be said to be meaningful even though it is neither analytic nor empirically verifiable. I should, however, claim that there was at least one proper use of the word "meaning" in which it would be incorrect to say that a statement was meaningful unless it satisfied the principle of verification; and I have, perhaps tendentiously, used the expression "literal meaning" to distinguish this use from the others...It is

85 Ibid p. 102

indeed open to anyone to adopt a different criterion of meaning and so to produce an alternate definition which may very well correspond to one of the ways in which the word “meaning” is commonly used...Nevertheless, I think that, unless it satisfied the principle of verification, [a statement] would not be capable of being understood in the sense in which either scientific hypotheses or common-sense statements are habitually understood.⁸⁶

I have reproduced the rather significant block of text above in order that we might examine Ayer's ideas as to what counts as a “meaning”. For a number of positions are available as to whether or not we take Ayer's emotivism to be a theory of meaning. We could choose to divide the theory into a negative semantic component and positive pragmatic component⁸⁷, or we could see the project as attempting to highlight a certain special kind of meaning⁸⁸; we could even view Ayer's statements on the results of his positivism as essentially disposable, and thereby open the way to subsuming his view to that of Stevenson⁸⁹. In light of the passage above, though, I find it most reasonable to take the first of these as substantially correct. Resolving the apparent tension in Ayer's claims requires taking Ayer's terminology of “literal meaning” at face value, and then writing off the “meaning” we attribute to moral claims as something else-- not

86 Ibid pp. 15-16

87 e. g. Stoljar 1993. For period (pre-Geach) responses that treat Ayer as following Stoljar's model, see e.g. McGill (1954), p. 70; Schuster (1953), p. 655. In fact, according to Harrison (1950), there existed an entire category of objection to emotivism based on the idea that, if emotivism held ethical propositions to be meaningless, it could not respect the commonsense differentiation between those that are gibberish and those that are not. Harrison's solution to the problem on behalf of emotivism is to allow for other forms of meaning besides the communication of beliefs, or, in other words, to take the alternative reading of Ayer's theory to the one I am advocating. The solution, though, seems to me to be a bit of an overreaction. All that we really need to defuse the worry is some factor to sort out the likes of “it is wrong to dissolve babies in acid” from “it is wrong to tuna”; there is then a *further* question about whether whatever difference in use we might pick out deserves the status of *meaning*, with all the attendant rights and responsibilities.

88 e. g. Schroeder 2010

89 e. g. Wilks 2002

something illegitimate, perhaps, but something that can be safely pushed to the pragmatics⁹⁰.

What we see in the earliest version of Ayer's theory, then, is expressivism in a fairly straightforward form, drawing on a simple but not wholly un-nuanced notion of an attitude, and deploying that notion to form a theory about ethics that avoids entanglement in semantics⁹¹. To bolster this frame, Ayer adds an account of moral dispute, which I alluded to in section 1.3 but now have the liberty to explain more fully. In its original form, it is an example of the approach to moral dispute that argues that the data, properly viewed, does not show anything that ought to make even a very simple expressivist theory worried. Ayer claims⁹² that we never really engage in rational debate about morality, but only about the facts of contested cases; debate always breaks down at the moral level. Where there exists a genuine difference in attitudes, rather than a disagreement about the facts of a case, argument gives way to condemnation and conflict. These claims are meant to be borne out by observation; reflection on actual moral disputes, rather than on the kind of inference that is used to

90 Perhaps we can better understand the idea of non-literal attributions of “meaning” if we look at some ordinary examples, for instance, instances where we give a “meaning” of a term in order to explain a particular example of its usage, or to help someone else use the term, but do so in a way that does not constitute a definition of a term. For example, consider the following dialogue:

A: “Let’s go to the movies!”

B: “Meh.”

A: “What does ‘meh’ mean?”

B: “It means, ‘I’ve been sitting in front of a computer screen working for the last forty hours straight, and spending even more hours sat staring at screens really isn’t my idea of fun at the moment.’”

In this dialogue, B’s second line is not a definition of ‘meh’. Most likely, it is not even what he meant by ‘meh’; rather, his first line merely expressed his unenthusiasm, and the second provides an explanation for that lack of interest. So his second line does not provide, in any literal sense, the meaning of his first. This example, then, shows a non-literal use of “meaning”.

91 And thus also any version of the Frege-Geach problem that is wholly semantic in nature.

92 See Ayer (1936 [1952]) pp. 116-117

demonstrate the supposed logic of moral language, is supposed to convince us that Ayer has it right when it comes to moral dispute.

Ayer's strategy is an "all-in" approach. What I meant by that is that it relies entirely on the claim that we are always arguing about matters of fact when we apply the method of rational debate in moral disputes. Such an approach has its advantages; with no account for why we would use fact-like practices in an emotive realm, Ayer's theory gains in parsimony, and does not suffer any plausibility penalties from an account that is less than fully comfortable. But dealing in absolutes leads to a worst-case scenario vis-a-vis counterexamples; one case of what Ayer claims to be impossible would sink the entire account of dispute. Given that Ayer's theory is supposed to be backed up by observation, with all theorizing further along in the order of explanation than the claim about the nonexistence of moral dispute (and thus no possibility of recourse to principles to rule out the possibility of counterexamples), this weakness is quite dangerous.

3.2 Ayer's emotivism improved

Accordingly, in his introduction to the second edition of *LTL*, Ayer modifies his approach to the problem of dispute. One modification, made without explanation⁹³, is to allow that once debate has "broken down", it needn't degenerate into abuse, but can continue through emotive persuasion. Ayer says very little about this change, and I will follow suit, since there seems to be no reason to exclude the possibility in question.

93 Ibid p. 22

Another change is to allow that arguments about the facts of moral cases needn't presuppose a common moral (emotive) position⁹⁴; it might also be possible to persuade one's interlocutor into a change in attitude through discourse about facts⁹⁵.

By far the largest change, though, is Ayer's addition of a mechanism whereby moral language can participate in disputes about the relevant facts for moral issues. The way this idea is supposed to work is that we have a certain body of sentences, which are literally meaningless, that use moral language. Sometimes, these meaningless sentences are used simply to express feelings. "Lying is wrong", for instance, is a meaningless sentence that may be used to express disapproval of lying--as if one had said the word "lying", as Ayer puts it, "in a peculiar tone of horror, or written it with the addition of some special exclamation marks."⁹⁶ But there are other cases where

what seems to be an ethical judgement is very often a factual classification of an action as belonging to some class of actions by which a certain moral attitude on the part of the speaker is habitually aroused.⁹⁷

In his discussion in the introduction to the second edition of *Language, Truth, and Logic*, Ayer gives the examples of one committed to utilitarianism using moral language as a vessel for talk of happiness, or a religious person using it for talk of rules laid down by authority, but the factual uses needn't be ones that refer to a *theory*. Rather, the idea is that our actual attitudes are not so great in number or fine in grain as we may assume,

94 Ibid. Contrast with *ibid* p. 111

95 *ibid*

96 *Ibid* p. 107

97 *Ibid* p. 21

or as Ayer puts it:

the common objects of moral approval or disapproval are not particular actions so much as classes of actions; by which I mean that if an action is labeled right or wrong, or good or bad, as the case may be, it is because it is thought to be an action of a certain type.⁹⁸

We might, I suppose, see these wider-ranging attitudes endorsed by Ayer as filling the role of “moral principles”, though Ayer does not give us enough to go on to know just how wide-ranging the attitudes are meant to be (and thus how fitting the label of “principles”). It is these that we voice when we use ethical language to express attitudes. When we argue about ethics, if we are proceeding rationally, we are not arguing over the principles themselves, but rather about whether or not a shared principle is relevant to a certain case or cases, or in other words, whether or not some pre-existing attitude shared by the debaters applies applies to the case(s). Ayer's theory already included the idea of winning a moral debate through debate about facts that determine whether one's wide-ranging attitudes apply in a particular case:

In short, we find that argument is possible on moral questions only if some system of values is presupposed. If our opponent concurs with us in expressing moral disapproval of all actions of a given type *t*, then we may get him to condemn a particular action *A*, by bringing forward arguments to show that *A* is of type *t*.⁹⁹

It is this kind of debate that we are actually engaged in when we form logical arguments using what looks like moral language. The idea that is new in the second edition of *LTL* is that, when we do so, we use the same meaningless sentences as we do

98 *ibid*

99 *Ibid* p. 111

to express certain attitudes-- that is, as we use when we express our moral principles. Although the words are the same, what they are used for is quite different¹⁰⁰.

The effect of this change on Ayer's theory is to add another reply to the dispute problem. If much of moral discourse is actually factual rather than emotive in nature, then we would expect such discourse to proceed in the manner expected of discussion about facts, even if morality *itself* is essentially emotive. According to this account, you can have what almost anyone would call a genuine moral dispute, with moral principles being applied and normative conclusions generated. Basing the expressivism on wide-ranging attitudes gives the theory the ability to explain arguments with normative conclusions: if we express attitudes¹⁰¹ that speak to more than just the individual object at which they are immediately directed, and then argue that a certain case falls within the category over which the attitude ranges, we thereby apply the expression of the attitude to case.

Does Ayer have the basic ingredients for a solution to the dispute argument? Here we might gain some understanding by examining, and contrasting, a different expressivist theory-- the earliest version of Simon Blackburn's theory, that of his 1973 [1993] paper "Moral Realism". Blackburn starts with a component that should be familiar to us from our examination of Ayer: the idea that, when using moral language,

100 We should beware not to think of the difference in terms of a rigid separation of uses. There is, after all, nothing to stop us from expressing emotions during a moral argument, as that is something that a great deal of language, whatever it else it might be for, can do. Likewise, we might sometimes use ethical terms in their factual capacity even when not making an argument. Rather, the difference is between the typical uses outside of and within arguments.

101 The use of "attitudes" here is indicative of another change in Ayer's theory over time. In interview (Honderich 1989, in Griffiths (ed.) 1991 p. 215), Ayer admits to moving away from the use of emotions in his theory, and to "attitudes".

we sometimes express attitudes, and other times discuss those attitudes:

The question, for the anti-realist, is that of how his account of the hypothetical is to cohere with his account of straightforward assertion in such a way that '*P*, and if *P*, then *Q*' entails *Q*, where *P* is a moral proposition. The theory tells us that anybody asserting '*P*, and if *P*, then *Q*' where *P* attributes worth to a thing *expresses* his attitude to that thing, and *asserts* that that attitude involves a further attitude or belief. There is, when that has been done, a logical inconsistency in not holding the further attitude or belief.¹⁰²

Blackburn's objective here is to try and reclaim ordinary argumentative practices for the expressivist. The idea is that an expressive premise can be stretched to include a conclusion by a second premise that makes a statement about the first; if the expressive premise does include or involve the conclusion, then it would indeed be a contradiction to fail to accept that conclusion. And so Blackburn hopes to give deductive arguments the force that we expect.

Unfortunately, as Blackburn realizes, there's no way that the expressivist is going to get away with this notion of involvement:

It may be thought at this point that too much weight is put upon the word 'involves.' For can an attitude involve a belief (or a belief an attitude in such a way? We recognize that people can hold the moral attitude without the belief and *visa versa*: what other involvement is there?¹⁰³

The first problem for the maneuver is that the view of involvement needed for this maneuver is untenable on a straightforward notion of the attitudes expressed in moral discourse. Even *without* considering arguments that combine the normative and the

102 Ibid p. 127

103 Ibid. Naturally, the fact that Blackburn does not endorse the attitude-involvement account has not stopped critics from seeing his (1973) as endorsing it; see, e. g. Schroeder (2010), whose criticisms (pp. 112-115) of the involvement account mirror Blackburn's own.

descriptive, the account has trouble. For example, consider the following argument:

3.2e1 If pushing people off of bridges in New York is wrong, then pushing people off of cliffs in Canada is wrong.

3.2e2 Pushing people off of bridges in New York is wrong.

3.2e3 Therefore, pushing people off of cliffs in Canada is wrong.

The account that we are currently considering holds that 3.2e1 says that the attitude expressed in 3.2e2 (disapproval of pushing people off of bridges in New York) includes the attitude expressed in 3.2e3 (disapproval of pushing people off of cliffs in Canada). But clearly that is *not* the meaning of 3.2e1; it seems utterly absurd to think, for instance, that one could not believe in 3.2e1 without committing to the idea that one's feelings about the use of Canadian geographic features as gravitic murder instruments are nothing more than a subset of one's feelings about the use of New York transportation conduits for the same purpose.

Things get even worse for the involvement account when we start to consider arguments that mix the descriptive with the normative. For example:

3.2e4 If Bob is allergic to fish, I should not give him a bagel with lox.

3.2e5 Bob is allergic to fish

3.2e6 Therefore, I should not give him a bagel with lox.

Now the problem is that I can accept 3.2e4 and yet deny exactly what the involvement account would translate it into, namely, that my belief in Bob's fish allergy includes disapproval of giving him a bagel. I might, while accepting the entire argument, think

that there are some situations under which I would come to reject 3.2e4 and 3.2e6 without rejecting 3.2e5. I realize, for instance, that I might come to believe that Bob deserves death by anaphylaxis, which could then combine with the idea that Bob is allergic to fish to yield the conclusion that I ought to give him lox. Yet, if my belief in Bob's allergy involves the disapproval of giving him lox, then I could not cease to disapprove of giving him lox while at the same time maintaining my belief in his allergy. If I really thought, in accepting 3.2e4, that 3.2e5 involves 3.2e6, then I should not think it possible that I could come to deny 3.2e6 but not 3.2e5. It seems impossible, on the involvement view, to *in fact* accept the argument 3.2e4-6, while at the same time acknowledging the *possibility* of coming to deny 3.2e6 without denying 3.2e5. And this seems wrong, because cases where one who accepts the argument might come to deny 3.2e6 without denying 3.2e5 are not difficult to imagine.

Blackburn's own response to this problem is to move away from the idea that complex moral sentences make statements about attitudes, and to the position that such sentences are expressions of attitudes towards involvement relationships between beliefs and attitudes:

To say that the attitude of approval to courage involves the belief that courage is a quality by which a man must act to achieve happiness is, as I said, to express a moral standard, to make a moral claim. The subtle thing is that the subject of the claim is now the attitude of approval to courage, and it is said that this *ought* to involve the belief that courage involves happiness.¹⁰⁴

The idea of higher-order attitudes-- attitudes about attitudes and combinations

thereof-- has already been touched upon in Chapter 2, and we'll hear more about the

104 *ibid*

theories that Blackburn bases on them in Chapter 5. For now, though, a quick note about this specific version of the idea: “ought to involve” seems unnecessarily strong. The notion of involvement might seem useful if you want a way to make a statement about attitudes and their relations to beliefs and one another, but if you are merely judging the states that you have expressed, and not making any claim about them, having the judgement be about involvement relations inherits potential pitfalls from the involvement account for no apparent benefit over simply judging combinations of attitudes.

Ayer's position differs from Blackburn's in that Ayer has the idea of wide-ranging attitudes. As a result, he needn't (and doesn't) claim that we make statements that are directly about our own mental states in the process of moral debate. Application of wide-ranging attitudes is not an analogue to the inclusion of one attitude in another, because demonstrating that a case falls within the set of situations towards which an attitude holds can be done entirely with arguments about the facts of the case.

Ayer does not give us enough text to even attempt a more specific reconstruction. Nonetheless, it is possible to comment about the position at a very general level, and this is enough to formulate a problem. We can either reconstruct Ayer as offering a slow-track theory and try and build an account of dispute out of a translation procedure for moral sentences, or we can reconstruct him as offering a fast-track theory. The former approach runs into difficulties that feels reminiscent of the

problems with Blackburn's involvement theory. First, if we believe that there is more than one moral principle-- more than one fundamental attitude involved in morality-- we run into the problem of drawing conditionals across attitudes. For example:

3.2e7 If telling a harmless lie is wrong, then dynamiting a hamster is wrong. If our disapproval of harmless lies and our disapproval of dynamiting hamsters are manifestations of different attitudes, then it is not at all apparent how we could explain 3.2e7 in terms of the application of a wider attitude. Second, it seems that we ought to handle the likes of 3.2e4 by making them general statements-- that is, statements about how, if a certain condition holds, some attitude or other applies (without specification of which, beyond the valence specified by the conditional itself). Yet, if there is no specification of the particular attitude involved, there is nothing to say whether it is or is not the same attitude expressed in the expressive premise.

A fast-track theory built out of the components that Ayer leaves us fares no better. One difficulty with this tactic is that Ayer's theory draws upon ordinary factual dispute as its explanation of moral dispute. Yet, we know how to form sentences to express the moving parts in a factual dispute in a direct and explicit manner, in such a way as to render unsatisfying the approach of defining complex moral sentences by their role in moral practice. In other words, if Ayer is right about the complex sentences in moral dispute being disguised statements of facts, they should be translateable into undisguised statements of facts. But if we could do that, we would not be using the fast-track approach in the first place. In other words, our reconstructed-Ayer seems to

have committed himself to the slow-track approach. But even if that approach works, he hasn't given us enough pieces to reconstruct a theory that succeeds in executing the approach.

Perhaps reconstructed-Ayer can get around this problem. He might, if he has abandoned many of the commitments of the actual Ayer, claim that we engage in a practice wherein we exchange facts in order to adjust our application of wide-ranging attitudes, but that this practice sometimes involves sentences that do not permit of further analysis beyond their role in it. Of course, this is not something that the actual Ayer claimed or would claim, but it is an interesting variation of the theory, so we will consider it anyway. The move here would be to say that a conditional, for instance, isn't necessarily voicing a specific relationship of derivation or manifestation between attitudes, but rather constitutes a kind of generic move with a wide variety of possible underliers drawn from all of the possible factual relationships that might make it so that in a case, an attitude applies, if that or another attitude applies in that or another case. Such a view would be a genuinely fast-track position, but one that ultimately leaves us searching for a better alternative. The kind of relations that it is based on--factual relations governing the application of attitudes-- is not adequate to the dialectic needs of the fast-track. Due to the fast-track's order of explanation, it relies on the inherent plausibility of whatever account of dispute it provides. Now the real Ayer *does* argue for the position that there are no reasoned moral disputes that are not actually factual disputes, but it is only a "look and see" kind of argument-- upon reflection, we

are supposed to realize that the point at which moral dispute breaks down to become nothing more than name-calling is exactly the point at which we finish arguing about non-moral facts¹⁰⁵. There is certainly *something* to Ayer's observation; careful reflection does reveal that it is surprisingly difficult to find realistic examples of moral disputes that on the one hand have the potential for rational argument to be fruitful and yet on the other hand cannot be explained as disputes about facts. But that is insufficient for present purposes. We would need far more than that the radical idea that there are no genuine rational moral disputes seems less crazy upon reflection than it does at first; we would need more even than that the idea is not a stretch at all, which goes beyond what the kind of argument Ayer offers can achieve. We would need Ayer's theory about what is really going on in moral dispute to seem so plausible, so *right* upon reflection that it becomes a net source rather than a net recipient of explanatory power. But the kind of argument that Ayer gives us, while interesting, is not capable of going that far.

Ayer's theory, then, serves to illustrate a point about the kind of solution to the dispute problem that one needs in order to take the fast track. I think, though, that there is a deeper lesson that we can learn from Ayer. Ayer's theory is very simple, even when we account for the fact that Ayer writes in such a condensed style that explanations of his theory tend to be longer than the original statement. It is also modular; because of the simplicity of the theory, the account of dispute is separable from the rest of the theory, and indeed most of the account was not even present in the original version of *LTL*. But, of course, this raises questions. If the theory is

105 Ayer (1936 [1952]) p. 111

fundamentally modular, what if we put on a different account of dispute? What other modules might the theory acquire to help its overall plausibility?

3.3 Ayer and quasi-realism

The thesis of this section is that Ayer's theory, in the form to which it had evolved by the end of Ayer's life in June 1989, deserves the label of "quasi-realist" (strictly speaking, I should say that I aim to find some subtle foreshadowing of the quasi-realist view in the complete Ayer, but the insane version of the thesis is more interesting, and thus more likely to lead to good results, so we'll go with that). "Quasi-realism" is sometimes used as a blanket term for the sort of view expounded by Simon Blackburn and Alan Gibbard, which distinguishes itself from other forms of expression by a kind of mimicry of the realist. As Blackburn (1993) writes:

I dramatized the question of whether a non-descriptive story can indeed understand the use we make of [embedded] contexts by inventing the figure of the quasi-realist, or someone who 'starting from an anti-realist position finds himself progressively able to mimic the thoughts and practices of realism.' Quasi-realism is not really another 'ism' in the sense of a position or an ideology in the same space as realism or anti-realism; it represents more of an attitude of exploration of the reality of the boundaries that those 'isms' demand.¹⁰⁶

In chapter 2, though, I used "quasi-realism" not in the general sense wherein it serves as a synonym for expressivism, but in reference to a certain kind of theoretical maneuver, namely that such theories take the basic expressivist strategy of analyzing moral sentences as expressing attitudes, and then apply to apparently meta-ethical

¹⁰⁶ Blackburn (1993) p. 4.

statements about things like the “existence of objective values” or “moral properties” an analysis that places them within the same practice. Using this component in a theory embodies an “attitude of exploration”, per Blackburn's quote, insofar as it allows expressivist theories to incorporate apparently realist claims into themselves without (hopefully; we'll come back to this in a minute) thereby becoming realist. Here, and when we return to quasi-realism in Chapter 6, I will continue to use quasi-realism in the more specific sense, the one in which it refers to the specific theoretical maneuver described.

We can find passages in Ayer which seem to employ a maneuver of this type:

To say that anything is right if someone thinks so is unobjectionable if it means no more than that anyone is entitled to use the word 'right' to refer to something of which he morally approves. But this is not the way in which it is ordinarily taken. It is ordinarily taken as the enunciation of a moral principle...it may appeal to some; it does not, in fact, to me.¹⁰⁷

For a subjectivist, egoism would be a coherent theory. There would be no inconsistency in advocating as a principle of conduct that each man should aim at the realization of his own interest. But this argument could be countered by our simply equating the belief in so-called objective goodness with the principle that, in default of special reasons, no one person's interests should count for more than anyone else's. It would then be left for us to decide whether this principle was acceptable.¹⁰⁸

This strategy is used to help the expressivist in preserving, rather than having to

107 Ayer (1954) p. 248. Compare this to Blackburn's response to the objection from unacceptable mind-dependence (Blackburn 1984, p. 218)

Suppose someone said 'if we had different sentiments, it would be right to kick dogs', what could he be up to? Apparently, he endorses a certain sensibility: one which lets information about what people feel dictate its attitude to kicking dogs. But nice people do not endorse such a sensibility. What makes it wrong to kick dogs is the cruelty or pain to the animal.

108 Ayer 1984 p. 34

challenge, much of ordinary practice, e.g.:

it is by no means improper to refer to ethical utterances as statements; when someone characterizes an action by the use of an ethical predicate, it is quite good usage to say that he is thereby describing it; when someone wishes to assent to an ethical verdict, it is perfectly legitimate for him to say that it is true, or that it is a fact¹⁰⁹

Such a strategy raises the question of just how far the expressivist wants to take it. For in each instance he has a few options: 1) he can analyze a claim and thereby render it harmless; 2) he can deny the claim; or 3) he can say that the claim accepts multiple readings, ones internal and external to the practice. There are dangers lurking on all sides here, especially for (3), but I doubt we'll have time to get into them in any depth. For now, let's pretend, contrary to fact, that we're confident that, one way or another, we can keep all three options alive. Then we might ask how often we ought to apply each of them.

Let's introduce a little bit of terminology: call **boldness** the frequency with which a theory chooses option (2) (we might have to weight this by significance, but let's let that concern rest for now).

Boldness might be important. Why? Because, without any **boldness**, the quasi-realist gains burden with regards to distinguishing his position from realism. If he sticks to (1) at all times, his theory will consist of the same claims as realism, and in addition, a number of further claims about the analysis of the first claims. If the quasi-realist prefers to utilize (3), he will have all of the same claims as if he had only used

109 Ayer 1954 pp. 231-232

(1), plus further claims as to the additional readings of the first set of claims, and yet more claims about the infelicity of whatever those additional readings claim is being said. Both of these cases have the quasi-realist saying everything that the realist says, so it is unclear how the two are in disagreement. Worse still, any argument that gives us reason to think the shared claims are false could trouble the quasi-realist as well.

The jury is still out on this issue (it's called the “problem of creeping minimalism”, courtesy of Dreier (2004b)).¹¹⁰ For our part, we will return to the question in Chapter 6. For now, let us simply note that some strategic application of **boldness** would stop us from having to worry about it in the first place. And if simplicity is regarded as a theoretical virtue, **boldness** can help us achieve it.

Ayer, as one might suspect, is heavy on **boldness**. Examples include the straightforward denial of objective values, e.g. Ayer 1967 p. 217, Ayer 1984 p. 33, and the denial of genuine logical debate of basic values: e.g. Ayer 1936[1952] p. 112. To increase his **boldness**, he takes the position that, while he might be forced into (3) in order to make the quasi-realist maneuver, he does so only out of a kind of compassion for ordinary use, while holding that we really ought to be stricter and then apply (2). For example:

Very often, what [a philosopher] is doing, although he may not know it, is to recommend a new way of speaking, not just for amusement, but because he thinks that the old, the socially correct, way of speaking is logically misleading, or that his own proposal brings out certain points more clearly. Thus, in the present instance, it is no doubt correct to say that the moralist does make statements, and, what is more, statements

¹¹⁰ See also e.g. Asay (2013), Dunaway (2010), and Chrisman (2008).

of fact...But when one considers how these ethical statements are actually used, it may be found that they function so differently from other types of statement that it is advisable to put them into a separate category altogether¹¹¹

If we consider Ayer's reply to the embedding problem, he seems to be taking a slow-track approach-- indeed, the idea of "emotive reactions to hypothetical instances", if we take the more plausible reading (attitudes towards hypotheticals) sounds a lot like Blackburn's higher order attitudes approach. But the interesting thing is that Ayer's **boldness** means that he can *only* take the slow track. Why? Consult the above quote. According to Ayer, many of our ordinary linguistic practices with regards to morality are misleading. We don't have to take his word for it, either; his **boldness** speaks for itself on this count. Yet, at the same time, Ayer's theory is of the mimicking sort; his comment, quoted earlier, about propositional form being a "disguise" attests to that. Since this claim is itself a part of his **boldness**, we cannot allow our reconstruction to diverge on that matter, at least not while the result of his **boldness** level is exactly what we are considering. On the mimicking approach, explanatory force is generated from the fittingness of the surface forms. Ayer denies that there is any such fit; in fact, much of his writing on the matter stresses the difference between moral and non-moral cases. so he loses his technique for explaining the presence of the surface form. He also doesn't believe in genuine moral arguments. But without moral arguments, he loses the main tool for constructing a functional account of complex moral sentences. So it seems like the fast-track order of explanation is closed off to Ayer; at Ayer's level

111 Ayer 1954 pp. 232-233

of **boldness**, any advantages to the use of propositional form in general would be more than outweighed by the massive error into which it has (apparently) led us.

As a fast-tracker, then, Ayer would get nowhere. But we've already realized that we have reason to worry about the slow-track approach in general, and about Ayer's attempts at following it in particular. We can resolve the historical puzzle from the beginning of this chapter by saying that Ayer's confidence was at least partially justified-- his expression of it coming before the problems with the slow track were known-- but his theory cannot stand as it is, nor even with our attempts at highly charitable reconstruction.

For our purposes, there is a further matter of interest. In the previous section, we discussed the modularity of Ayer's theory. We see that once again with the question of his **boldness**. Some **boldness** is important, but an excess of it is fatal. If we look at the selections above, we see that his **boldness** is determined piecemeal; he decreases **boldness** by adding quasi-realist aspects to his theory, and increases it again with aggressive clarification of those aspects. Since Ayer's theory allows for the easy adoption of instances of the quasi-realist maneuver, there seems to be no reason why the theory could not be modified to have whatever **boldness** level was necessary. In theory, the possibility exists for an Ayer-type theory, armed with a different theory of disagreement and a proper level of **boldness**, to attempt the fast-track approach. Whether that theory succeeds or fails will depend on its details, and we have no particular reason to suspect that it will succeed, or, for that matter, that it will even

maintain enough of the original theory over the course of its modifications to be seen as any more closely related to Ayer's view than to any other expressivist theory. But, in the interests of precision, we must point out that the concerns that we have considered do not conclusively rule out any possible variant on Ayer's theory.

3.4 Using Ayer's ideas within a hybrid theory

While the involvement and application accounts do not work as overall pictures of moral dispute, none of the arguments directed against them are so extreme as to say that the phenomena referred to cannot happen. Nowhere have we, or the opponents of Ayer or Blackburn, claimed that it is impossible for an attitude to involve another attitude, or for people to have wide-ranging attitudes that manifest as a number of simpler attitudes. And so, while these phenomena cannot occupy the central position in whatever theory we come up with, there is nothing to stop them from playing a secondary role. We cannot use them to ground a general explanation of moral dispute, but they may, nonetheless, be useful in explaining specific cases. Since our plan is to try and construct a hybridized fast-track theory, supplemental explanations for individual cases can be useful. Besides, if Ayer's theory is meant to be based, first and foremost, on observations about moral discourse, it would be nice if whatever we came up with could do justice to the same observations.

In our discussion of hybrid theories in Chapter 2, we touched upon the idea of rescuing slow-track elements for use as a "model", and discussed some of the ways that

parts of a slow-track theory could contribute to a fast-track theory. Blackburn's theory of higher-order attitudes and their place in reasoning, for instance, may have seemed inadequate as a core account of moral reasoning, but it was able to make itself useful to a fast-track as a possibility proof and supplemental explanation. We can deal with Ayer's insights into the surprising abundance of questions of fact in moral dispute and the relatively small number of attitudes necessary to cover a wide variety of cases by giving these insights the same status. The kind of dynamic that Ayer gestures at to reinforce his claim that there are no genuine reasoned moral disputes might not succeed at that job, and might not be capable of playing the lead role in a more sophisticated expressivism, but it can be of use in several lesser capacities. It can help us understand how some moral debates are likely to go, and what sort of considerations might be productive; it can show that an emotive practice needn't be devoid of factual reasoning. None of these functions require the theory to generalize to the extent that we would require for an overall account of morality. It suffices that the dynamic occurs *sometimes*, that the interplay of general attitudes and facts is somehow involved in morality-- denying that seems implausible.

We can find many examples of this dynamic in everyday life, after all. One might, for instance, have a general attitude of being for the underdog in sporting competitions¹¹². But it can be surprisingly difficult to work out which team is the underdog, especially if one has only a small degree of familiarity with the sport in question. What are we to make of such a case? It seems false that one would not favor

112 Where, incidentally, the "Hooray!" and "Boo!" model is quite apt.

either team; one favors the underdog, which implies that one favors a team. The basic attitude of favoring the underdog is not dispelled by one's ignorance in the case at hand, because we have other handles on it; the attitude can still be brought to mind by associations that do not concern the particular case. Nor is it merely that one *wishes* to support the underdog, although that attitude may hold as well¹¹³, since not all ways of manifesting support for the underdog require knowing which team is which. For example, in motor racing, inclement weather might increase the chance of an upset; thus, one who favored the underdog might hope for rain.¹¹⁴

More generally, it is clearly possible to have an attitude without realizing all of its potential manifestations. Now it is important to realize that unknown manifestations are not reducible to predictions of future attitudes. For example, suppose that I like pizza. We can then distinguish two separate claims about a particular pizza that I have not yet tried: that I *do* like it, because it is pizza and therefore subject to my general attitude towards pizza, and that I *will* like it when I try it. One might argue that an attitude cannot apply in unrecognized cases by claiming

113 That is, deploying higher-order attitudes seems unlikely to lead to a superior explanation of this case.

114 Some other examples in the mold of the underdog case: I might disapprove of plans that will place me in a dangerous situation, but not know which of several alternatives will do so; I might dislike rabid animals, and wonder if a certain squirrel is rabid; I might like Tuesdays, but not know what day of the week it is. Here, for simplicity's sake, I am focusing on one type of case, where there is a general attitude and uncertainty as to its applicability. I have chosen this kind of case because it is the sort of case where thought and argument are likely to become involved; the uncertainty demands a resolution, and the resolution will reveal the extension of the attitude. That said, this is not the only kind of case that we might look at to investigate the interplay between general attitudes and cases. Simple subsets are another; for instance, if I dislike dust, it follows that I dislike the 154th dust-mote to be swept up in a particular cleaning session, even though I do not have any thoughts about, or take particular notice of, that specific mote. These are perhaps clearer at making the point that our attitudes are coarse-grained rather than fine-grained, but present fewer interesting details by which possible alternatives may be drawn out.

that all apparent cases of an existing attitude applying in an unrecognized case are actually cases of predicting that one will come to have a new attitude towards the case, of the same valence as the general attitude, when the case is encountered. The position is motivated by the apparent oddness of Ayer's dynamic; is it actually possible to prove to someone not that they *will* have an attitude towards something that they have yet to encounter or realize, but that they *do* already have an attitude towards it?

This argument does not work. The apparent plausibility of the story it tells comes from failing to recognize the true implications of the view. For instance, we might not feel terribly troubled by the idea that I can't actually like a brand of pizza that I have not tried, even if I like pizza in general. But as the scale drops, our troubles increase. Am I unable to like a brand of pizza, but only to like the pizzas of that brand that I have had, and predict that I will like the ones I have not? Am I unable to like a whole pizza, but only to like the slices that I have eaten, and predict that I will like the ones I have not? Am I unable to like a whole slice of pizza, but only the bites of the slice that I have already eaten, and predict that I will like the ones I have not? As the regress continues, the argument comes to demand that our attitudes are finer and finer-grained, until it comes into direct conflict with how our attitudes seem to be.

We can also, with a bit of reflection, explain the feeling of discomfort that motivated the argument. Attitudes are closely tied to feelings; indeed, as mentioned earlier, Ayer's own theory employed feelings in its early version, and later moved away from them in favor of attitudes. When we imagine a case of the sort that we have been

talking about, we picture ourselves confronting the unknown thing, and not having any particular feelings about it (yet). But, although the two are closely related, attitudes are not feelings. At this point, we do not have an account of what attitudes *are*; we will return to this issue in the next chapter, though even then, the range of possibilities is such that it would be very premature to claim finality for what we come up with. That said, since attitudes are familiar to us, we can say a bit about them without the need for an overarching account. For one, attitudes are something that one has, even when one is not experiencing them. My attitude towards pizza is my attitude towards pizza, even on days when I have not even considered consuming anything other than coffee. Furthermore, it is possible to be mistaken about one's own attitudes. I might, for instance, tell myself that I am okay with lettuce as a pizza topping; after all, lettuce itself is fine, and most foods can work on pizza-- but, every time I have actually encountered such a pizza, it has not been to my preference. Sooner or later, I will just have to admit to myself that I do not like salad pizza, but in the meantime, it seems quite possible for me to possess the attitude B!(salad pizza) without being aware of the fact. This potential for ignorance works both ways; it could well turn out to be that what I dislike is not, in fact, salad pizza, but some subset of salad pizzas, ones sharing a common feature in their composition, that happens to include the salad pizzas I have thus far experienced.¹¹⁵

¹¹⁵ In this regard, the theory does quite well at tracking our intuitions. It seems natural to say that I might think that I dislike salad pizza, but be proven wrong; a future experience might demonstrate that what I dislike is not, in fact, salad pizza, but only a certain kind of salad pizza. Getting the right answer here without attributing massive error to anyone requires the possibility of general attitudes that could include a future case, so that it can be proved that one does not have the attitude by its failure to manifest in that case.

If we can have attitudes even when we are not experiencing them, and we can be wrong about our attitudes, then the absence of a feeling, leading to a belief in the absence of an attitude, cannot be relied upon to tell us that the attitude really is absent. And so I might have the general attitude H!(underdog), and this attitude might apply to one of the teams at the event, without me feeling compelled to actually “hooray” for one team or the other (in the case where I do not know which team is the underdog). The absence of the feeling is not indicative of the absence of the attitude, especially if the inability of the attitude to manifest itself in feeling is easily explained by the absence of the information required to direct it. Rather, the best explanation of such cases is that the attitude is present even though we do not know which object(s) it encompasses.

What these considerations show us is that the simple idea of coarse-grained attitudes seems able to stand on its own. As such, it can serve as a source of explanatory power. Now, we know from our earlier examinations that this power is limited, at least insofar as it applies to the moral case; we cannot use coarse-grained attitudes as the foundation of an expressivist theory. The dynamic, in short, exists. If attitudes often work in this sort of way, then an expressivist would do well to recognize the possibility of the attitudes-- or at least *some* of the attitudes-- involved in morality working in this way. In fact, we can strengthen this claim; there seems to be absolutely nothing to *stop* an attitude, just because it is of whatever kind is involved in moral judgments, from taking a scope that allows for question on specific cases. It would,

given all of this, be quite awkward to claim that moral attitudes do not follow the pattern; at worst, it could even jeopardize the familiarity of the claimant's notion of attitudes. For our purposes-- that is, for the sake of using this part of reconstructed-Ayer's theory as a secondary component in a fast-track expressivism-- this degree of resolution is sufficient.

Chapter 4. Richards and Stevenson

In *Ethics and Language* (1944), C. L. Stevenson presents us with a view that bears both close similarities and striking differences to Ayer's emotivism. Both views place a heavy explanatory emphasis on the role of moral sentences as expressions of our attitudes. Moreover, like Ayer, Stevenson believes that rationality applies only to matters of fact, writing:

*If any ethical dispute is rooted in disagreement in belief, it may be settled by reasoning and inquiry to whatever extent the beliefs may be so settled. But if any ethical dispute is -not rooted in disagreement in belief, then no reasoned solution of any sort is possible.*¹¹⁶

On the other hand, where Ayer ascribes no meaning to ethical sentences, Stevenson gives them multiple kinds of meanings, including “emotive meaning” and “descriptive meaning”¹¹⁷. And he has rather different ideas about moral disputes (at least, those that are not merely persuasive or rhetorical, a category that both he and Ayer take to encompass a significant proportion of moral arguments). It is to Stevenson's theory that we will now turn our attention.

4.1 Richards

Before we discuss Stevenson's ideas about the meaning of moral sentences, we

¹¹⁶ Stevenson (1944) p. 138

¹¹⁷ See *ibid* p. 37 and thereafter. I should like to flag that I find the phrase “emotive meaning” useful and will continue to use it even after our discussion of Stevenson concludes. If calling it “meaning” outside of the historical context in which it was named offends your semantic sensibilities, feel free to substitute “impact” or “effect”. In particular, though, I happen to prefer the emotive meaning of “emotive meaning”, and in particular that it resists the superstition-driven urge, discussed later in this chapter, to demote its subject to second-class function status.

should settle the question of what he means to be talking about when he talks of “meaning”. For he is not working within the popular model of words having a single meaning (even in a particular context) that we can use to come up with the analysis of a sentence in which they occur. In fact, he cautions us, using a phrase from I. A. Richards, to beware the “One And Only One True Meaning Superstition”¹¹⁸. For the interpretation of this phrase, we turn to Richards himself¹¹⁹, where the phrase occurs in a discussion of a “context theorem of meaning”¹²⁰. From the name of the theory, it would be easy to interpret it as a claim only about variations in meaning across different contexts, but that interpretation would miss the point. To understand the theory, we must first understand what is meant by a “context.” Ogden and Richards

118 Stevenson (1944) p. 84, after Richards (1936) p. 39. Stevenson takes much from Richards, utilizing not only the latter’s collaboration with C. K. Ogden, *The Meaning of Meaning* (1923), which is noteworthy for containing an early form of emotivism, but also many of Richards’ solo works, including *Principles of Literary Criticism* (1924), *The Philosophy of Rhetoric* (1936), *Interpretation in Teaching* (1938), and *Mencius on the Mind* (1932). Stevenson also published on Richards, contributing to Brower (1973)’s collection *I. A. Richards: Essays in His Honor*. Note that, in Richards, the “One and Only One True Meaning Superstition” is also sometimes called the “Proper Meaning Superstition” (e.g. Richards (1936) p. 11) or the “Proper Usage Superstition” (e.g. *ibid* p.75).

119 I should like to flag that in the discussion to follow, I will be moving between Richards’ various works fairly freely. Some who have studied Richards (for example, Max Black, as represented in his (1948)) would instead prefer to divide Richards into “early” and “later” periods, with the line somewhere around 1934-1935. To some extent this is no doubt justified. In a later footnote, for instance, I remark upon the gradual erasure of the fourfold division of Total Meaning, which seems to have largely gone by the board by about 1933, though the individual parts of the division recur, though with a somewhat different status, in 1934’s *Coleridge on Imagination*. And Black is quite right to consider the problematic idea of a “referent” as belonging primarily to the early Richards; for evidence of this, one need do no more than compare the theory of definition on p. 113 of *The Meaning of Meaning*, wherein the important thing in the process of definition is “to find the referent”, to the theory of multiple definition described below. I think, though, Black underestimates the cohesiveness of Richards’ corpus, and in particular, the degree to which the theory of *The Meaning of Meaning* remains in the background of Richards’ later works. In particular, the idea of context (in a technical sense) and its relation to meaning persists though *The Philosophy of Rhetoric* (see Lecture II of the same), and even to 1938’s *Interpretation in Teaching*, which goes so far as to refer the reader back to *The Meaning of Meaning* for further details (Richards 1938 p. vii). We will note other recurring themes as well. And so I consider it no great crime to reconstruct Richards from bits drawn from works of different ages, so long as we are also careful to note changes where they occur.

120 The phrase is from Richards (1936) p. 38; Ogden and Richards (1923 [1948]) lays out the theory.

start with the idea of “recurrent clumps of events”¹²¹: occurrences that regularly present together¹²². These Ogden and Richards divide into external contexts-- clumps of events out in the world-- and psychological contexts-- clumps of mental events.¹²³

The next component is the claim that “what a word means is the missing part of the contexts from which it draws its delegated efficacy.”¹²⁴ “Delegated efficacy” is Richards’ shorthand for the idea that a word acquires a meaning by taking the place of some of the items that constitute a particular context. A context can thus come to obtain through the use of a word, even if some or all of its members are absent. Here, too, we must be careful not to oversimplify. It would be easy to confuse this idea with a theory where a word’s meaning is a particular referent, which it then supplies to the situation to complete a context. But a look at the Ogden-Richards triangle of reference¹²⁵ shows us that this reading would be a misunderstanding; according to that diagram, there exists no direct relationship between a symbol and its referent. The two are only related through the mediation of “thought or reference”¹²⁶, which must therefore be prior to the symbol-referent relation (“stands for”¹²⁷). Understanding the idea properly requires making sense of Richards’ position that “metaphor is the omnipresent principle of language”¹²⁸, which he explains thusly:

121 Ogden and Richards (1923 [1948]) p. 56.

122 The official definition, from *ibid* p. 58: “A context is a set of entities (things or events) related in a certain way; these entities have each a character such that other sets of entities occur having the same character and related by the same relation; and these occur uniformly.”

123 *ibid*

124 Richards (1936) p. 35

125 Ogden and Richards (1923 [1948]) p. 11

126 *ibid*

127 *ibid*

128 Richards (1936) p. 92

The view that metaphor is omnipresent in speech can be recommended theoretically. If you recall what I tried to say in my Second Lecture about the context theorem of meaning; about meaning as the delegated efficacy of signs by which they bring together into new unities the abstracts, or aspects, which are the missing parts of their various contexts, you will recollect some insistence that a word is normally a substitute for (or means) not one discrete past impression but a combination of general aspects. Now that is itself a summary account of the principle of metaphor. In its simplest formulation, when we use a metaphor we have two thoughts of different things active together and supported by a single word, or phrase, whose meaning is a resultant of their interaction.¹²⁹

This passage gives us a better idea of what Richards is getting at, especially if we stay true to the triangle of reference's assertion that the relation between a symbol and a thought is causal in nature.¹³⁰ If, for instance, a thought of a dog is a thought caused by a dog, and a thought caused by a dog is a thought caused by certain recurring bundles of external and internal sources, and the word "dog" functions to complete those bundles, then the word "dog" can bring about the thought of a dog. The word "dog" refers to a dog insofar as it brings about the same thought as a dog, which it does by substituting for the missing parts of those contexts that constitute the experience of a dog. Observe that "dog" supplies what is necessary to complete the appropriate contexts-- not the dog itself.

We can make sense of the comments about metaphor by positing that meanings are over-generated in this view. The completion of some contexts may involve the completion of others in turn. In cases of metaphor, a word or phrase simultaneously completes multiple contexts, including some which may be relatively

¹²⁹ Ibid p. 93

¹³⁰ Ogden and Richards (1923 [1948]) p. 11. In fact, they go so far as to say that "to say 'I am thinking of A' is the same thing as to say 'My thought is being caused by A'" (ibid, p. 55).

rare to see completed by the word; in such a circumstance, the word would have multiple meanings. Richards claims that this phenomenon is ubiquitous.

The frequency of multiple meanings is increased by various dynamics. For one, there is nothing in the idea of a context to prevent words from participating in contexts in the same way as anything else-- that is, not as a substitute, but as a regular part of a bundle of recurring stimuli. Words might also complete contexts in this capacity, and thereby fulfill the criteria for possessing additional meanings. In fact, Ogden and Richards suggest that a word could even enter into its own contexts, such that substituting any other word would change the contexts and thus the meaning-- thus does their theory explain why some people become quite resistant to change in their verbiage¹³¹. In addition, words can contribute to contexts in indirect ways, an idea that Richards terms the "interanimation of words"¹³². For example, when multiple words sound alike, and have similar meanings-- an occurrence called a "morpheme"¹³³-- the existence of the cluster can influence the effect of similar-sounding words.¹³⁴ Occasionally, words can influence the effect of similar-sounding words even when there is no morpheme¹³⁵. Sound is not the only factor, either; Richards claims that words can also influence those that are related to them in meaning¹³⁶. When these factors are combined, the result is a theory in which, as Richards says, "As the movement of my hand uses nearly the whole skeletal system of the muscles and is

131 Ibid p. 216

132 Ibid pp. 47-66

133 Ibid p. 59

134 See Ibid pp. 59-62

135 See ibid p. 63

136 See ibid p. 64

supported by them, so a phrase may take its power from an immense system of supporting uses of other words in other contexts.”¹³⁷

Richards is not merely aiming to point out variation in the effects of words. Rather, he intends to make a point on the subject of meaning itself. One premise of his argument is that interanimation dynamics, and the other factors that influence the effect of words, are not constant-- the effect is not a fixed sum of the various factors-- but that, rather, any of the various factors can be brought out to varying degrees in individual instances. Indeed, the ability to control these variations, Richards claims, and to derive great effectiveness from the interanimation of words, is a hallmark of a skilled writer¹³⁸. Another is that the magnitude of interanimation effects is at least the equal of any other mode of language-operation.¹³⁹ The first claim functions to remove one possibility for identifying a single meaning with a word, and to bolster the interanimation account by demonstration of its explanatory power. The second claim functions to remove a possible reason for discounting interanimation effects from the theory of meaning.

To these claims, we should add a pair of error theories intended to help explain away the “one and only one true meaning superstition.” First off is “the magical theory of names”¹⁴⁰-- the idea that a study of customs from around the world reveals the frequent attribution of magical powers to language, accompanied superstitions about

137 Ibid p. 65. For more on Richards’ theory of interanimation, see also his *Interpretation in Teaching* (1938), chapter 14.

138 Ibid p. 75

139 Ibid p. 69

140 Ibid p. 71; the idea is the subject of Chapter II, “The Power of Words”, of Ogden and Richards (1923 [1948]).

the power of names to confer a supernatural dominance over the object, person, or even deity named. Ogden and Richards think that the impressive utility of language invariably gives rise to superstitions about words¹⁴¹; those who engage in the study of language-- whether they be philosophers, rhetoricians, linguists, psychologists, logicians, or what have you-- are not immune to the temptation. They diagnose the desire to create, in theorizing, a direct connection between words and things (that is, to fill in the leg of the Triangle of Reference that is, in their own diagram, conspicuous by its absence) as a result of a lingering attraction to word magic¹⁴².

The other error theory employed by Richards in *Philosophy of Rhetoric* concerns the "Rule of Club Spirit"¹⁴³: the idea that we are often driven to separate meanings into *correct* and *incorrect*, *proper* and *improper*-- we might as well add *strict* and *loose*, *literal* and *nonliteral*-- by nothing more than factors wherein the making of such distinctions is tied to participation in, and motivated by the desire for inclusion in, social groups. In short, the Rule of Club Spirit is that the social benefits of staying strictly on the right side of a distinction belonging to a certain class are often sufficient to serve as the sole reason in favor of making the distinction for individuals. For groups, demanding such adherence is likewise pragmatically motivated, at least at the origin (though societal changes may have since shifted the weight of practicality away from the practice); if the distinctions are taught as part of education, for instance, they create a very easy

141 Ogden and Richards (1923 [1948]) p. 24

142 Ibid p. 47; see also Richards (1936) p. 71.

143 Richards (1936) p. 79

way of distinguishing between the educated and uneducated¹⁴⁴. In all of this, the distinctions themselves are simply arbitrary; there need neither be any independent fact of the matter, nor any present practical benefit, for a distinction of this kind to persist. Richards illustrates his point with examples based on pronunciation, but insists that the same point applies in the case of meaning¹⁴⁵.

The points about word magic and Club Spirit are not proofs. That is, Richards does not prove that either or both of these are the source of the Proper Meaning Superstition. But although Richards is often forceful in his assertions about these phenomena, a careful look at the end of Chapter II of *The Meaning of Meaning* will show the order of explanation in a different light:

The contextual theory of Signs to which, then, we first proceed, will be found to throw light on the primitive idea that Words and Things are related by some magic bond; for it is actually through their occurrence together with things, their linkage with them in a 'context' that Symbols come to play that important part in our life which has rendered them not only a legitimate object of wonder but the source of all our power over the external world.¹⁴⁶

This passage suggests a more modest role for the error theory. Instead of the error theory wiping out the Proper Meaning Superstition, and the context theorem of meaning moving in to fill the vacuum, the idea rather is to start with the context theorem. Ultimately, the context theorem and the account of meaning-superstitions can be mutually reinforcing; the context theorem provides an alternative account of language, while the meaning-superstitions simultaneously help the theorem explain

144 Ibid p. 79-80

145 Ibid p. 78

146 Ogden and Richards (1923 [1948]) p. 47

contrary intuitions about the nature and role of meaning and grow in their own plausibility from the context theorem removing the explanatory benefits of a friendlier position towards Proper Meaning. In particular, the error theories can step in once the context theorem has made its claims about language in general; they can be deployed to calm any urges to seek something additional for the theory of meaning.

The result of all this, according to Richards:

Preeminently what the theorem would discourage, is our habit of behaving as though, if a passage means one thing it cannot *at the same time* mean another and an incompatible thing. Freud taught us that a dream may mean a dozen different things; he has persuaded us that some symbols are, as he says, 'over-determined' and mean many different selections from among their causes. This theorem goes further, and regards all discourse -outside the technicalities of science -as over-determined, as having *multiplicity of meaning*.¹⁴⁷

The phrase “multiplicity of meaning” is interesting for what it tells us about the

147 Richards (1936) pp. 38-39; italics added. The passage here is interesting in part because of what it tells us about how we are to apply Richards' theory at levels other than that of individual words. At certain points in *Philosophy of Rhetoric*, Richards writes as if the Proper Meaning Superstition was a problem to do with the analysis of words in particular, and that a proper meaning could be found at a higher level-- that of the sentence, or of the entire utterance-- and then leveraged into an account of the (singular) meaning of each word in that particular instance. For passages that seem to suggest this sort of reading, see *e.g.* *ibid* p. 66, p. 50. I do not think, though, that this reading is the one that we should prefer, in light of his attention here to entire passages and bodies of discourse. For Richards also writes: “Often the whole utterance in which the co-operating meanings of the component words hang on one another is not itself stable in meaning. It utters not one meaning but a *movement* among meanings.” (*ibid* p. 48). And then: “in the strictest prose the meanings of the separate words theoretically stay put and thought passes from one to another of them. At the other end of the scale the whole meaning of the sentence shifts, and with it any meanings we may try to ascribe to the individual words. In the extreme case it will go on moving as long as we bring fresh wits to study it.” (*ibid* p. 48). Richards argues that, as a result of the Proper Meaning Superstition and the psycho-social factors promoting it, we tend to see the strict case, which is in fact very rare and largely confined to the domain of the hard sciences, as far more common than it actually is (*ibid* p. 120), and to express a preference for using the strict case as a model when doing so is, in fact, inappropriate (*ibid* p. 73). Accordingly, even if there *is* a strict end of the spectrum on which sentence-level meaning is fixed, there is no counterexample to the multiplicity of meaning to be found there, for such cases are neither representative of language as a whole nor to be given an exalted status in our data-set on account of its simplicity.

relationship between the various powers and uses that Richard ascribes to words and meaning. If we are to understand Richards-- at least to the extent needed to understand Stevenson, which is our goal-- we need to understand how to collect his remarks on the complexity and diversity of meaning into a theory on the meaning of meaning. Our first step in that regard will be to take stock.

Richards' target, particularly in *The Philosophy of Rhetoric*, is the following view: that words have definite meanings, that these meanings are in turn the determinants of sentence-level meaning¹⁴⁸, and that the title of "meaning" should be reserved for those powers that are fixed and definite in the case of words, or compositionally determined in the case of sentences. His strategy against this target, in very broad terms, consists of two parts. The first part of the strategy is to point out the great variety and flexibility of the powers of language. The context theorem of language and the idea of the interanimation of words are what drive this half of the approach. The second part is to attack the idea of taking one particular slice of the powers of a bit of language and giving the powers within that slice a special status. This half can itself be divided into halves: one set of arguments connected to the context theory of meaning, which oppose the idea of a neutral or empty context, and dispute the frequency of rigid language, and another group of arguments which are meant to call into question the intuitions that there must or should be a special slice. Into the latter category fall the

148 "The traditional Usage Doctrine...treated language on the bad analogy of a mosaic, and conceived composition and interpretation as if they were a putting together or taking apart of pieces with a fixed shape and color, whereas, in fact, the interanimations of the meaning of words is at least as great as in any other mode of mental performance." (Richards (1936) p. 69)

arguments concerning “word magic” and “Club Spirit.”

Without some further exploration, though, we could go wrong in how we put together these pieces into a theory. After all, if we take the whole mess of powers that a word could ever have, it would not have multiplicity of meaning; even summing up all of the powers possessed in a specific context would be inconsistent with Richards' claim of simultaneous incompatible meanings. It's not that each individual power is a meaning, either, for that would be inconsistent with the incompatibility part of the same claim. It seems, then, that we will have to go beyond the obvious options.

Help is forthcoming by way of the theory of definitions. While Richards does provide us with some hints of this in *The Philosophy of Rhetoric*¹⁴⁹, for further elaboration, we will also have to draw upon his 1933 “Multiple Definitions” and his 1932 *Mencius on the Mind: Experiments in Multiple Definition*. First, we should see what we get from *The Philosophy of Rhetoric*, where he writes of the following sense of the word “usage”: “Some specific power which, in a limited range of situations and with a limited type of verbal context the world normally exerts.’ (This is often called a *use* or *sense* and is what the Dictionary attempts to record in its definitions, by giving other words, phrases and sentences with the same specific power.)”¹⁵⁰. Richards claims that most words, including those used in philosophy, admit of many such senses, and that attempts to diversify our vocabulary in order to have words with only one sense are likely to be futile¹⁵¹. A natural question would be how many, or, in other words, how

149 See in particular his “Note” at the bottom of p. 65

150 *ibid*

151 Richards (1933) p. 31

fine-grained a “sense” is. The *Philosophy of Rhetoric* comment suggests maximally so, but this is misleading. Better is this remark from “Multiple Definitions”:

It is important to keep in view this fact that [in listing different senses of words] we are not here putting on paper something which is given to us, so much as making a machine-- a machine for controlling thought which will let us do some things and keep us from doing other things. It is a good machine if it is of use to us; any changes which will make it of more use to us will make it better. They are not able to be tested in any way other than this.¹⁵²

This passage suggests a) that a sense can capture more than one power of a word, and b) that the nature of the boundaries of the sense is practical; the inclusion of some powers into a sense and the exclusion of others is done in such a way as to be useful, rather than to conform to any preexisting fact. Additional support for this reading is found a bit further down:

The attempt to make a machine like this is, in fact, a way (and the best way) to the discovery of how our minds do their work. But, as we will see, our minds do their work in a number of different ways. They put the chief divisions, upon which all the others are dependent, in a number of different places for different purposes. So a number of different machines, different “philosophies,” different “logics,” are possible and necessary.¹⁵³

What Richards often presents us with, when dealing with philosophical questions, is a long list, usually numbered, of different senses. “Multiple Definitions” attempts this project for central terms in philosophy¹⁵⁴; *The Meaning of Meaning* gives us lists for “beauty”¹⁵⁵ and “meaning”¹⁵⁶; *Mencius on the Mind* gives us “beautiful”¹⁵⁷; “knowledge”¹⁵⁸;

152 Ibid p. 39

153 ibid

154 See ibid for a list

155 Ogden and Richards (1923 [1948]) pp. 142-143

156 Ibid pp. 186-187

157 Richards (1932) pp. 100-104

158 Ibid pp. 105-109

“truth”¹⁵⁹; “order”¹⁶⁰; and “principle”¹⁶¹¹⁶². The passages I have quoted above tell us what the lists in a multiple definition *are*; they are slices cut out of the sum of the powers of a word, cut not upon the lines of any natural division, but as we see fit to cut for our purposes, whatever they may be.

The interpretation above is consistent with the remark from *Philosophy of Rhetoric* that led to its generation. Multiplicity of meaning is achieved, since the causal powers of a word can be divided in many different ways; simultaneous incompatibility is achieved by locating opposed causal powers, which may cancel one another out in the current sign-field¹⁶³, and presenting a meaning that captures only the one, and another meaning that captures only the other.

159 Ibid pp. 111-115

160 Ibid pp. 118-119

161 Ibid p. 122

162 *Mencius on the Mind* retains the fourfold division of function or meaning from *Practical Criticism* (1929), wherein the “Total Meaning” (the capitalization is Richards’) of a passage can be divided into “sense” (a state of affairs being pointed to, a proposition, or thoughts to be caused in the listener), “feeling” (emotions expressed), “tone” (variations of wording in response to the speaker’s attitude towards the listener), and “intention” (why the speaker is speaking) (see Richards 1929[1930] pp. 180-183). For convenience, *Mencius* combines feeling, tone, and intention together as “gesture” (Richards 1932 p. 98); both “sense” and “gesture”, on the *Mencius* view, are important to the project of multiple definition, and so we often see (e.g. for beauty) lists divided into sections for “senses” and “gestures”. This division is not present in “Multiple Definitions”, most likely because it is itself working with a different, more inclusive sense of “sense” (not 8.1, “A general property of a thought by which what the thought is about is fixed”, but rather 8.3, “A use of sense which is nearer to the one we are making is that in which *persons who are wise are said to have sense*, that is, to have *good sense*. Good sense is, at least in part, a power to keep our thoughts, the senses of our words, in the right places.” (Richards 1933 pp. 47-48)); on the operant sense of “sense”, it seems like “gesture” ought to be included.

Interestingly, in *Coleridge on Imagination*, “sense”, “feeling”, and “tone” recur (p. 88) as parts of (no longer capitalized) total meaning. But by then they have acquired an interesting status, for Richards claims that “Meanings may be said to have any parts which, for our purposes, we find useful as instruments in comparing them” (pp. 87-88), and gives the aforementioned parts as examples. Accordingly, whether or not we divide a list of multiple definitions into “sense” (8.1) and “gesture” categories will depend on whether or not that axis of division is profitable.

163 A.k.a. “settings”, per Richards (1938) p. ix; in other words, “contexts” in the non-technical sense.

4.2 Richards and “emotive meaning”

Now that we have an understanding of the place of meaning in Richards' theory, we will turn our attention to the more specific question of his theory of *emotive* meaning. We have already seen a little of this; of the senses (8.3)¹⁶⁴ or meanings of a word, some may be characterized in terms of sense (8.1), and others in terms of gestures, which, as mentioned, include expressions of feeling¹⁶⁵. We have also seen various arguments that Richards directs against the idea of choosing one particular meaning as *the* meaning of the word. To the latter set, we should now add arguments to the effect that meanings associated with emotive uses of language have just as much right to the name “meaning” as those that are symbolic.

First, gesture, though sometimes derivative of sense, is not always so. Richards gives the examples of “Hooray!”¹⁶⁶ and “Damn!”¹⁶⁷; these words are wholly emotive in character. Certain phrases, for instance “Very pleased to meet you!”¹⁶⁸, Richards claims, can function the same way; while he does not elaborate, the idea seems to be that, although the latter may look like a report of an emotion rather than an expression of one, a bit of introspection will reveal the contrary. If we are to be able to provide meanings for such words and phrases, we must be prepared to draw from their gestures, since sense is unavailable.

164 For readers who missed the use of such numbering in a previous footnote, it's derived from Richards' (1933) list of senses for philosophical terms, including for “sense” itself; the number indicates which sense is being used. Please consult aforementioned footnote for elaboration.

165 *The Meaning of Meaning* uses “symbolic” and “emotive” (p. 149) for the sides of this distinction.

166 *Ibid*

167 Richards (1932) p.99

168 *ibid*

Second, Richards argues¹⁶⁹ that understanding sense is not necessarily prior to feeling. When reading poetry, for instance, we often start to receive feeling from the poet before resolving any concrete ideas; the poem may well begin to work on us with its first words, before it has given us enough to constitute any proposition. I suppose that in extreme cases, such as nonsense poetry, we might never be able to put together any sense worth the name. Of course, feeling *can* be subsequent to or at least influenced by sense, but it need not be.

Third, Richards argues¹⁷⁰ that in some cases, although one *could* assign a sense to a passage, doing so results in nonsense. In such cases, claiming the passage to have the sense as its meaning would presumably fail to satisfy the purpose of meaning, and so our best course of action is to disregard the sense and consider the passage at a wholly emotive level. Statements such as “Poetry is a spirit”¹⁷¹ or “Man is a worm”¹⁷² fall into this category: we could describe them as having such meanings as ascribing the property of being a spirit to poetry, or of ascribing the property of being a worm to Man, but doing so would (quite plausibly) be nothing more than misunderstanding. Now in such cases, the gestures may be derived, either wholly or in part, from the sense that the passage might otherwise have had; but this need no more imply that the passage has the sense in the present use than the fact that “bank” can mean the side of a river means that the “bank” in “I went to the bank to withdraw some quarters” carries

169 See *Coleridge on Imagination* p. 89

170 This point is presented as distinct from the first in *Mencius on the Mind* (p. 99), though combined with it in *The Meaning of Meaning* (p. 149). I see no reason to think the difference is particularly substantial; I present them separately in my own treatment merely for expository reasons.

171 Ogden and Richards (1923 [1948]) p. 149

172 *ibid*

that meaning.

This third argument returns us to the subject of Richards' account of metaphor, which is closely linked with his account of emotive meaning. Metaphors, in Richards' theory, often (though not always) operate by means of the likeness between two things¹⁷³, which he terms "tenor" and "vehicle"¹⁷⁴. Richards claims that we should not identify the meaning of a metaphor with (a statement about) its tenor; that is, that a metaphor is, or at any rate can be, more than a fancy way of making a statement that we could just as easily have made straightforwardly¹⁷⁵. His argument for this claim is that we observe other sorts of cases. For example, we encounter cases where the focus of a metaphor is not really on the tenor save, perhaps, to preserve topicality, or otherwise justify a statement about something that, were it not used as a metaphor, would be irrelevant¹⁷⁶. We also, Richards argues, encounter cases where we simply cannot seem to come up with any satisfactory non-metaphorical translation of the metaphor.¹⁷⁷ In such cases, there may be gestures made by the metaphor; being

173 Richards (1936) p. 118.

174 Ibid. Now "precise" definitions-- though of course Richards would argue that the sort of definition that the phrase calls to mind would really be anything of the sort-- for the terms "tenor" and "vehicle" are elusive. Richards argues (in *Interpretation and Teaching*, pp. 120-121) that common candidates, like "subject" and "expression", or "meaning" and "metaphor", will all add confusion rather than clarity, and (in *Philosophy of Rhetoric*, p. 96) that we in fact have *no* adequate terminology for distinguishing the two. If the absence of synonyms, simultaneous with the absence of explicable rules of use, for these terms is the bad news, then the good news is that the terms seem workable enough without anything of the sort.

175 See Richards (1936), p. 100: "the vehicle is not normally a mere embellishment of a tenor which is otherwise unchanged by it but that vehicle and tenor in co-operation give a meaning of more varied powers than can be ascribed to either. "

176 Richards (1936) p. 100

177 Richards (1932) p. 114: "Sometimes the metaphorical expression is a convenience only; it is not difficult to substitute another non-metaphorical expression in its place. But sometimes the metaphor is irreducible-- no literal expression seems to be a satisfactory substitute."

(perhaps) expressive rather than declarative in nature, these could account for the difficulty of finding a satisfactory translation.¹⁷⁸

Because of the above difficulties, Richards thinks that the idea of understanding metaphors by de-metaphorizing them-- of putting them into non-metaphorical language-- is itself a mistake. Richards argues that we may feel that, in de-metaphorizing a metaphor, we are gaining clarity, but that feeling is a mistake, one driven by an emotional bias in favor of methods that look and feel scientific (even if they are not, in fact, using the scientific method)-- this is "scientism"¹⁷⁹. In fact, all that we are doing in such a process is mangling our subject matter and confusing ourselves. And at the same time, our feelings against an alternative approach, wherein we would allow metaphorical language to remain metaphorical, are influenced negatively by a parallel bias against metaphor. When we look at metaphor without bias, we see that it is neither particularly difficult to understand¹⁸⁰ nor any more prone to disruptive variation than non-metaphorical language¹⁸¹. We also see that there is plenty of scope for useful work on metaphor that does not rely on de-metaphorizing it¹⁸². What all of this amounts to, for Richards, is an argument on the level of meta-theory; the conclusion of the argument is that the success or failure of a theory of metaphor does not hinge on the theory's ability to translate metaphors out of metaphor. If anything, we ought to be suspicious of theories that attempt to do so.

178 *ibid*

179 Richards (1948), p. 151

180 *Ibid* p. 150

181 *Ibid* p. 151

182 *Ibid* p. 150

Richards, in replying to Black's (1948) criticisms of emotive meaning, claims that, while there are non-emotive uses of metaphor, and non-metaphorical expressions of emotion, the two often co-occur, to such an extent that the study of emotive meaning is inextricable from the study of metaphor.¹⁸³ Richards asks us to consider, given such, what form it would be reasonable to expect a theory of emotive meaning to take. If emotive meaning is often associated with metaphor, then any theory of emotive meaning will often be dealing with metaphorical cases, and when doing so, it will have to follow the guidelines for a theory of metaphor that we have just finished discussing. At the very least, our expectations for a theory of emotive language will have to be informed by what we ought (and ought not) to expect from a theory of metaphorical language.

4.3 Stevenson

Stevenson, who is heavily influenced by Richards, presents a theory the structure of which answers to the concerns we have been discussing. He gives us an account of ethics that is comprised of two different “patterns of analysis,” each of which is “a 'pattern' for making definitions”¹⁸⁴; these definitions, in turn, are not final or all-encompassing, and differ from one another. If we understand Stevenson as inheriting a substantial amount of meta-philosophy from Richards, we can make sense of this style of analysis in a way that does not attribute to it any self-contradiction or back-tracking;

183 Ibid p. 152

184 Stevenson (1944) p. 89

we can avoid having to think of Stevenson's theory as roundabout, indirect, or unsuccessful at achieving greater specificity. Rather, the theory is exactly what it is trying to be: an attempt at multiple definition for a term that Stevenson finds to be shot through with emotive meaning.

Before we go further, though, I would like to highlight a few of the *differences* between Stevenson's theory and Richards'. Where the Ogden-Richards theory is causal; Stevenson's is dispositional; rather than the actual causal powers of a term in the current setting, the fundamental component is the dispositional powers of the term¹⁸⁵, which it possesses across all settings. More precisely:

The meaning of a sign...is...a dispositional property of the sign where the response, varying with varying attendant circumstances, consists of psychological processes in a hearer, and where the stimulus is his hearing the sign.¹⁸⁶

This formulation maintains both the psychologism of the Ogden-Richards theory and its grounding in the effects of a sign. But it does offer a few advantages:

First, Stevenson's theory allows for meaning to remain constant. In fact, he saw constancy as necessary for a theory of meaning¹⁸⁷. Using dispositions, which remain whether they are realized or realizeable in the current setting, allows him to achieve this desideratum.

185 I'm being slightly more specific than necessary when assigning the disposition to the term. In fact, Stevenson thinks that it is completely arbitrary whether we choose to assign the dispositions to the terms themselves or to people; he declares a choice in favor of attributing the disposition to the term, on the basis that we find it more natural to talk about the meanings of the terms, but also that, strictly speaking, the other choice would have been equally correct. See *ibid* pp. 55-56.

186 *Ibid* p. 54

187 *Ibid* p. 43

Second, Stevenson's theory allows him to cite meaning all the way down to the level of individual words. Richards was actively opposed to the idea of meaning at the word level; he saw our tendency to seek meaning at that level as an unfortunate side effect of modern English writers to separate words with spaces¹⁸⁸, and of the lingering effects of Word Magic. Stevenson has no such hostility, and is perfectly willing to talk about the meaning of a word¹⁸⁹.

Third, Stevenson's theory achieves an elegant homogeneity, in that dispositions on different scales are themselves dispositions; a cluster of dispositions taken together themselves constitute a disposition, and a subset of the input-output relationships that constitute a disposition will itself constitute a disposition. And so Stevenson can talk readily at the lower level of the dispositions that make up the meanings of a sign, and still say that the meaning of a sign is itself *a* disposition, whereas Richards' theory can only be described in terms of sets, subsets, "slices", and other operations upon either individual causal relationships or the whole mass of causal relationships, which is a bit more awkward.

Of course, we should not make this distinction between Stevenson and Richards into more than it is. After all, reassigning the label of "meaning" from causal powers to the dispositions that undergird them is a terminological maneuver; it adds the presence of dispositions to the theory's list of

188 See e.g. Richards (1936) p. 47

189 See e.g. Stevenson (1944) p. 67

commitments, but the overall substantive change involved in this particular modification is limited. We should notice, in particular, that Richards' idea of "meaning" as to a degree arbitrary, where we are free to pick and choose from among the causes of a sign in response to practical concerns, is compatible with the dispositional account, though Stevenson chooses to be slightly less flexible in his own account of meaning. Stevenson distinguishes what is meant from what is merely suggested or implied by way of conventional rules-- meaning picking out those dispositions, or the parts of the disposition, that are bound and propagated by established rules¹⁹⁰, but that, for actual language, the rules are often absent, provisional, temporary, leaving us with a great deal of choice in how to pick out the meaning¹⁹¹. Indeed, it is often an analyst's own creation of rules that distinguishes meaning from suggestion for him.¹⁹² Compare the following remark from Stevenson to Richards' comment from "Multiple Definition", quoted above, on the idea of making a machine: "We must be sensitive, then, to the fact that ethical terms are not predestined to abide by any one set of rules, and that analysis cannot "discover" the "real" sense."¹⁹³ Despite the differences that I have highlighted, it seems that Stevenson's theory takes not only its overall structure and methodology from Richards, but also a significant amount of substance.

190 Ibid p. 86

191 ibid

192 Ibid pp. 86-87

193 Ibid p. 87

Stevenson's theory also deploys the notion of a disposition in its account of attitudes. For Richards, an attitude is “some special direction, bias, or accentuation of interest towards [a thing], some personal flavour or coloring of feeling”¹⁹⁴; for Stevenson, it is “purposes, aspirations, wants, preferences, desires, and so on”¹⁹⁵, all, ultimately, captured under the heading of disposition.¹⁹⁶

Attitudes, for Stevenson, are made up of, and thus are, dispositions.

Stevenson provides us with two “patterns of analysis”-- general ideas that are each intended to suggest a number of meanings for moral terms, though not forming any sort of exhaustive formula for capturing such meanings. The first pattern of analysis is derived by combining a subjectivist element with a coordinative element. This pattern is based on, but not identical to, the phrase “I approve of this; do so as well”¹⁹⁷: the idea of the first pattern is to replace the imperative in the phrase with an emotive element that maintains a sense of universality (that is, that it is the sort of thing about which conflicts should be resolved).¹⁹⁸ Alternative paraphrasings are “Oh

194 Richards (1929[1930]) p. 181

195 Stevenson (1944) p. 8

196 e.g. p. 224 of Stevenson (1963): “[A]n attitude is of a dispositional nature, involving a variety of responses that may progressively attend a variety of stimuli. Among the responses, in particular, there are various differences in *feeling* between admiration, desire, respect, etc.; and as I need scarcely add, these differences have their behavioral correlates.”

197 Stevenson (1944) p. 81

198 Stevenson decides (almost immediately after introducing it) that the imperative component is inadequate because what he really wants to capture is not a command, but an emotive *effect*. He tells us that the descriptive components “do not evidence the *contagion* of warmly expressed approval-- the interaction of attitudes that makes each man’s favorable evaluation strengthen and invigorate the other’s.” (p. 22). His subsequent remark that “This latter effect is highly characteristic of an articulate ethical agreement” (*ibid*) seems to suggest that the prior statement describes an observed datum of moral practice for theory to capture.

that you might approve of this as I do!”¹⁹⁹ and “I approve of this, how fine it is!”²⁰⁰; these, he claims, may feel a bit closer to correct, but are still incapable of being anything other than approximations.²⁰¹ Stevenson claims that this particular emotive meaning is not exactly captured by terms outside the moral realm, and so cannot be defined by means of a synonym; he writes, “The term ‘good’ is indefinable...if a definition is expected to preserve its customary emotive meaning. It has *no* exact emotive equivalent.”²⁰²

The second pattern of analysis is laid out by Stevenson as follows:

“This is good” has the meaning of “This has qualities or relations X, Y, Z...,” except that “good has as well as laudatory emotive meaning which permits it to express the speaker’s approval, and tends to evoke the approval of the hearer.”²⁰³

A second-pattern definition, then, combines an attribution of properties with a coordinating positive emotive meaning. Stevenson tells us that there must be boundaries as to what properties can be involved, though he declines to specify where they lie. He also alerts us to the fact that examples of the second pattern are likely to be “persuasive definitions”-- definitions with some emotive force of their own that comes to bear when they are invoked.

We might profitably linger for a moment²⁰⁴ on the idea of a persuasive

199 Ibid p. 81

200 ibid

201 ibid

202 Ibid p. 82

203 Ibid p. 207

204 Since it is only for a moment, we will not delve deeply into the technical workings of persuasive definitions. For readers who are interested in such things, a few contemporary papers on the issue that one might consult are Macagno and Walton(2008); Burgess-Jackson (1995); Walton (2001).

definition, if for no other reason that that the idea serves as an illuminating illustration of how Stevenson's theory works. To define the phrase more precisely:

A “persuasive definition” is one which gives a new conceptual meaning to a familiar word without substantially changing its emotive meaning, and which is used with the conscious or unconscious purpose of changing, by this means, the direction of people's interests.²⁰⁵

Recall what we mentioned earlier: that for Stevenson, definition is not (or at least not always) a matter of discovering rules that already exist, but as often as not is part of the creation of such rules. We might do so-- might take the senses of a term and divide them up through the establishment of rules-- for the sake of improving communication and our own clarity of thought; such is the motivation given by Richards for his project in “Multiple Definition.”²⁰⁶ But if the process of establishing such rules is executed on an interpersonal level-- if we have the power to establish, or work towards establishing, rules for more than just ourselves-- then other purposes become possible.

According to Stevenson's model of a persuasive definition, there often occur situations where a term, over time, comes to possess an increasing amount of emotive meaning, and at the same time suffers a diminution of its descriptive meaning.²⁰⁷ Once enough of the descriptive meaning has fallen away, and the descriptive content of the term become sufficiently vague, it becomes possible for individuals to change or add to the descriptive content without seeming to create a homophone (thanks to the

205 Stevenson (1938), p. 32.

206 Richards (1933) p. 31

207 Stevenson (1938 [1967]) pp. 33-34

vagueness), enabling the new term to inherit the emotive meaning of the old.²⁰⁸

I would like to make a couple of points about persuasive definitions. The first is that when Stevenson claims that second-pattern definitions are persuasive definitions, he is not thereby presenting an argument against them; they are not defective on account of being persuasive. This mistake would be easy to make, especially in light of the negative emotive meaning that “persuasive” often possesses in philosophical contexts. But if we employ Richards, we can construct the idea of a persuasive definition in a way that does not accuse such definitions of an improper ambiguity. Stevenson's picture portrays persuasive definition as involving an existing term being given a new definition, while at the same time maintaining its own identity in order to keep an emotive meaning. Yet we can portray the same substance in a more emotively neutral way if we think of it in terms of interanimation: if we create a new term by giving a new descriptive meaning to an old word, and decline to employ the cutting-off techniques of making explicit our stipulation, then our new term may well wind up strongly interanimated by the old. Of course, such a maneuver-- effectively a joining of contexts through the reuse of a sign-- may be confusing, if we are not on the lookout for it. But such potential for confusion is, on Richards' view, par for the course.

The second point that I should like to make about persuasive definitions is that, on Stevenson's view, our philosophy (including domains that we might think of as “meta-philosophy”) is shot through with persuasive definitions, and thereby with emotive questions that we must be careful not to confuse with factual questions. For

208 Ibid

instance, “meaning,” Stevenson argues²⁰⁹, possesses a positive emotive meaning; we praise forms of discourse by calling them meaningful, and condemn them in calling them nonsense. When creating a theory of meaning, we propose a persuasive definition (not, in itself, a bad thing, per my first point, above); and so will have to consider whether the resulting emotive meaning is desirable. The interesting thing about this point is that it illuminates the operation of Richards' myths within philosophy. The lingering appeal of Word Magic, for instance, can influence outcomes in the theory of meaning by causing us to accept certain persuasive definitions and reject others, where if we had been properly cured of these last vestiges of superstition, we would have done otherwise.

4.4 Relationship between Stevenson's ideas and the Frege-Geach problem

Neither of the patterns of analysis contains any explicit statement for how to deal with embedded contexts, but we can at least come up with some suggestions for how the theory might deal with them. Stevenson *does* give us a notion of compositionality (albeit a loose one), wherein the meaning of complex sentences is yielded by holistically combining the dispositions (according to Stevenson, descriptive meaning, like emotive meaning, is roughly or approximately characterizeable as a matter of dispositions, in this case “the disposition of a sign to affect cognition”²¹⁰ that constitute the meanings of the parts. He also tells us that words in sentences do

209 Ibid p. 42

210 Stevenson (1944) p. 67

maintain their own meanings:

Each word has an independent meaning in the sense that if it is replaced by certain others in any context, there will be a typical sort of difference in the meaning of the context; but the precise way in which the word's meaning is realized will depend on the other words that accompany it.²¹¹

From these ideas, we can begin to reconstruct how embedded contexts ought to work on Stevenson's theory. It would be all too easy to be lured by the forms of the patterns of analysis into thinking of the meanings of moral terms as straightforward conjunctions of a descriptive content with an emotive expression and/or an imperative. It would also be all too easy, given the amount of attention and even praise that both Stevenson and Richards devote to linguistic ambiguity, to start characterizing meanings disjunctively. Neither of these options, though, captures what Stevenson is trying to construct. Rather, the idea of a disposition-- which is also used to categorize attitudes themselves-- tells us how to put together the various aspects of a meaning. Stevenson wishes us to consider meanings as dispositions to have a psychological effect, either on "cognition" or on attitudes (note that, given that attitudes themselves are defined in terms of dispositions, and cognition may be so as well, meanings will be "second-order dispositions"²¹²). A disposition can be possessed even when it is not realized; if a glass can have a disposition to break when dropped even when it is not dropped, so too can a term be disposed to have a certain psychological effect even if it occurs in a context where it does not have that particular effect. We might use that possibility to

211 *ibid*

212 *ibid*

describe unasserted contexts: the dispositions of a term to cause the effects that we associate with its assertion are not active in unasserted contexts, but they are still present. At the same time, whatever dispositions might be active in unasserted contexts (perhaps further cognitive dispositions, including ones that affect how we treat the entirety of the sentence in which the unasserted term occurs) will still belong to the word in asserted contexts. Thus, we can say that there is no change in meaning between asserted and unasserted contexts.

What we ought to be asking, at this juncture, is how much we really ought to be expecting here. I argued above that the change from Richards to Stevenson that enables Stevenson to talk about meaning at the level of individual words, though not entirely without substantive import, is in large part terminological; in other words, there has been relatively little change from a theory that saw the idea of compositionality, or at least the kind of compositionality that we might capture by an image of building-blocks, as a myth to be overturned. A rule by which we might take “the meaning” of each part of a complex sentence, and composite them into “the meaning” of the whole, might well be too much; at the very least, any such rule, it seems²¹³, will need either a number of provisions, possibly infinite, that modify it in individual cases, or else can remain constant, but at the expense of requiring many individual words to contain, within their dispositions, a number of aspects, possibly infinite, that activate when the word enters into combination with particular other words, and change the resulting meaning. In either case, the required amount of

213 Following Stevenson, *ibid*

flexibility will make it difficult, if not impossible, to distinguish there being a rule from there not being one.

Let us turn our attention, then, to Stevenson's account of arguments, on which matter he gives us a few suggestions to work out. Let us begin our reconstruction with the idea that we might use the descriptive meaning when arguing:

In general, ethical statements, like all others that have at least *some* descriptive meaning, are amenable to the usual applications of formal logic. Care must be taken, of course, that verbally seeming contradictions are not merely apparent, due to a change in sense of the particularly ambiguous ethical terms; and further care must be taken to avoid emotive repercussions of otherwise innocent tautologies. Otherwise, this aspect of ethical methodology brings with it no special problems.²¹⁴

Despite the fact that Stevenson has an account on which the emotive meaning of a term can remain unchanged even in embedded contexts, he nonetheless chooses to claim that formal logic applies at the descriptive level. The emotive level is ruled out because, while he thinks minimal truth *is* applicable in emotive contexts, he also claims that minimalism “is not a sense that we shall want to be related to the terms ‘valid’ and ‘invalid’”²¹⁵, because validity grounded on minimal truth would simply be too “different from the more impersonal study in which students of validity in science and logic-- so at least they usually insist-- are engaged.”²¹⁶ So while we might deploy minimalism to explain our propensity for applying the labels of “true” and “false” in emotive contexts, it will not feature in the account of arguments.

214 Ibid p. 116

215 Ibid p. 169

216 Ibid p. 171.

It is worth pausing for a moment to identify the differences between this aspect of Stevenson's account and Ayer's, and also to investigate what is meant by Stevenson's warnings above. On the first count, the obvious difference is that Stevenson situates his argument on the semantic level, and Ayer the pragmatic-- and indeed, we might see some resonance between Stevenson's distinction between the said and the merely suggested on the basis of rules, and the fact that Ayer's theory quickly degenerates into a mass of unanswered, and perhaps unanswerable, questions if rule-bound procedures for determining what is expressed are insisted upon at all stages.

Likewise, Stevenson is aiming for standard validity (albeit only applied within a certain sense of the terms being used), whereas Ayer is not. The two are similar, though, in their basic dependence on factual argument alone to explain the use of logic in ethics, and therein lies a difficulty: in Ayer's theory, arguments were translated as needed from the moral surface form to the factual argument underlying it, whereas in Stevenson's theory, descriptive meaning is held fast by individual terms.

Stevenson's warnings about "emotive repercussions" seem to hint at, at the least, a feeling of unease in the general vicinity of the problem, but awareness is of little help. Taking his suggestion of "caution" as a call to simply rule out certain kinds of arguments will not solve the problem, because we seem to have no way of capturing all and only the right arguments to rule out, nor any explanation of what such a ruling out would entail. Stevenson suggests, by this warning, that we ought somehow to take

emotive meaning into account in argument, but we are not provided with a way of doing so.

To be fair to Stevenson, his attention with regards to moral arguments is focused elsewhere: on arguments where attitudes themselves are contested, through persuasive methods (especial attention is given to the providing of *reasons*). These other methods are taken to comprise the core of ethical discourse, and we can still engage in logical argument on related matters (about, say, the applicability of a reason), so perhaps Stevenson does not lose much in failing to fully answer the argument problem.

To further shore up any potential weakness, Stevenson calls our attention to certain pragmatic aspects of discussion about “validity” with regards to methods of argument. He has the insight that, when compared to that which is “valid,” other methods of argument pale; they become “merely”-methods (“merely persuasive,” “merely rhetorical,” and so on)-- and so seem inappropriate for ethics. Stevenson puts it as follows:

A dismissal of validity, even in this partial way, risks opening the way to certain misunderstandings in the course of guarding against others. The validity of a method stands out as the most conspicuous ground for choosing it; hence when certain methods, or aspects of them, are denied any connection with validity, one may feel that no ground for choice between them remains. Or if such a ground is recognized, it may seem to involve only a crude, forensic success. So long as one's opponent is impressed (a hasty critic may suppose), one method is as good as another; for the whole purport of ethics is to sway attitudes. Where Plato and Kant sought eternal principles of reason, are there merely the empty rules of rhetoric? After this one is likely to envisage disillusionment and chaos, and the many other disturbing “implications” which objective

theorists so habitually attribute to their opponents.²¹⁷

We can take from this argument the moral that, when it comes to evaluating how implausible it is to deny the validity of moral arguments, and how plausible it is to cite other forms of dispute in the place of valid logical arguments, our gut may lead us astray. We must be wary of reacting so strongly to the loss of validity that we fail to explore the potential of non-rational methods as explanations for arguments in select realms of discourse (such as ethics). Such caution will not make ethical arguments any less invalid, but it might go a ways towards letting us discover that that's not as bad as we might think. We must also be on guard for a related mistake: that, upon hearing that a certain theory does not return the result that moral arguments are valid, we gain the image of moral discourse (on that theory) being rendered impotent or superficial, and in reaction to that image reject the theory. But the image is born of the pragmatics of validity; calling ethical arguments "invalid" *feels* like calling them *bad* arguments, and so we picture a theory with an invalidity claim as discounting all of moral discourse. "Invalid", where validity is expected, carries a strong negative emotive meaning (or at least suggestion), and we dispute its applicability to ethical terms much as we wind up disputing because of the emotive components of ethical terms.

This insight is not restricted to Stevenson's theory in terms of its application. It could, perhaps with a few changes in how it is described, be applied by any theory that moves the majority of moral arguments beyond the sphere of logic-- or to any theory that relies on a non-literal account of moral arguments, such as Ayer's. It is not, in

²¹⁷ Ibid p. 156

itself, sufficient to solve any of our problems, but it does, at least, make the relevant positions more plausible, by providing an error theory for a certain kind of supposed implausibility.

4.5 What to learn from Richards and Stevenson

Richards and Stevenson do not have an answer to the Frege-Geach problem. Stevenson, as we saw, did offer some suggestions about compositionality, but nothing like a full-fledged theory. Perhaps this is to be expected. I wonder if they would not simply argue that Geach's worry is based on an incorrect view of language. Such an argument would claim, first of all, that the entire idea of words having a single, independent meaning, which they then contribute in some way to a sentence when they occur in embedded contexts, is simply mistaken, a holdover of *Word Magic* that is shown to be such by the atomic homogeneity of the Causal (or, for Stevenson, Dispositional) Theory of Reference. Nor is it the case that Geach's "Frege Point" is required in order to explain the validity of arguments; the study of logic through "symbolic procedure", for Richards in particular, is only an indirect approach, and not the exclusive approach, to what is fundamentally a matter of structures of contexts²¹⁸,

218 On this point, see Ogden and Richards (1923 [1948]) p. 68. I should point out that I am not being wholly faithful to Richards here; for Richards writes, in Richards(1934 [1962]): "But the peculiar reference of thoughts to the things-they-are-of gives them modes of interaction with one another which are lacking in the case of feelings. And this interplay is studied as their logical compatibility or incompatibility, and other relations." (p. 89), and so might not be congenial to dissolving the puzzle in this way. Presumably Stevenson would have to posit a greater complexity of interaction between attitudes than Richards was comfortable with; given that one of Stevenson's criticism of Richards is that the latter pays insufficient attention to relations between attitudes, particularly agreement and disagreement (see Stevenson (1944) pp. 10-11), this would not be much of a burden.

and so if we must sometimes discard it, we ought not find ourselves overly troubled.

In this argument, explanatory power is drawn holistically from Richards' arguments; it succeeds if Richards' theory, as a unit, is right, and fails otherwise. In the context of that debate, the Geach concern represents a slight liability for Richards' side: the other side, one might argue, has a clearer and simpler view of arguments, even without taking onboard any desire to accommodate emotive meaning. Since the overall theory of meaning endorsed by Richards and, in somewhat modified form, by Stevenson is no longer fashionable, we might regard this reply as similarly unsatisfactory. But, though it does not satisfy us philosophically, its availability might nonetheless help us understand Stevenson's choice of focus.

Although Richards and Stevenson do not have a solution to our central problem, they can help us in other ways. First of all, even if we do not follow Stevenson in defining attitudes in terms of dispositions, we can nonetheless add relations between the dispositions associated with attitudes to our understanding of the complexity with which attitudes might relate to one another. Secondly, we can make use of the observations of Richards and Stevenson even if we do not endorse the theories that they derive from those observations. Both of the next two chapters will draw upon such observations.

Chapter 5. Propositional Reflection

The mimicking expressivist seeks the freedom (in how he understands the conditional and other sentences with moral terms in embedded contexts) that can be obtained when he denies²¹⁹ that moral arguments have the forms that they seem to have. Any theory with such a denial, though, owes us an explanation of what is going on when such arguments are used.

Blackburn (1984), (1993), (1998) correctly argues that the right form of explanation is based on the idea of an isomorphism between the logical relations of the surface forms and some other relation between the attitudes expressed. Valid moral arguments, on such a view, compel one who accepts the premises to accept the conclusion, not on pain of irrationality, but on some other pain instead²²⁰. On a good theory of the type that we are considering, this alternative pain will be one that, first of all, is important, and second, matches what we are led to expect by the surface forms. Such a theory will have a far easier time arguing that it makes sense (from a practical perspective) for some of our expressions of attitudes to take propositional form, since, on such a theory, our reasoning does not go astray in the moral realm. We as philosophers might be tempted to go astray as to the nature of the enterprise with which our reasoning is engaged when we reason in the moral realm. But ordinary

219 Here we are once again operating under the proviso specified for chapter 2, that is, with quasi-realism excluded. As a result, positions will be attributed to Blackburn that will, no doubt, seem highly un-Blackburnian in character, and which are unlikely to be regarded as acceptable by Blackburn himself.

220 I expect that Blackburn would find what I am claiming here to be highly objectionable-- on which count please see footnote #1.

reasoning can proceed without concern.

My goal is to show how these desiderata can be met. Although Blackburn (1998) gives us an account that comes close, and a variation upon Blackburn's general theme comes closer, the best option takes a different approach. The overall theme of the best approach is that, while it would make for a very neat theory if there were only one relationship between the attitudes that we express in propositional form driving the adoption of that form, it is ultimately neither productive nor reasonable to expect any such thing. Rather, the matching of the logic of the surface forms is achieved piecemeal by an assortment of different relationships between the attitudes, each of which need be nothing more than locally isomorphic to the logic of the surface forms within the specific instances where our needs motivate us to track that relationship.

5.1 Propositional Reflections: The General Idea

If an expressivist chooses to use the mimicking strategy, he will claim that the assumption of propositional form by expressions of attitudes occurs for practical reasons, and not (barring quasi-realist accommodations) because expressions of attitudes somehow fit into that form. The latter alternative requires a direct solution to Geach's (1958), (1960), (1965) famous worry about the non-embeddability of attitude expressions, and so we can hardly blame the expressivist for looking for an alternative. Blackburn (1984) suggests the idea of taking advantage of propositional form to help make the coordination of attitudes orderly and serious. This tactic is all well and good;

after all, putting something in propositional form-- declaring "This. Is. X!"-- feels like putting your foot down on the matter in a way that other forms of expression do not. We might suppose that, having been conditioned by the logic of ordinary descriptive language to have a great aversion to the compresence of X and not-X, we can exploit that aversion for rhetorical effect even when we are doing something other than trying to get our facts straight.

But practical concerns weigh against one another. The assumption of propositional form suggests the applicability of the logic associated with the form. If the underlying logic of moral sentences is something different, then it would seem that that form is misleading-- and indeed *has* misled us, given that we regularly reason with moral sentences as if they were ordinary descriptive sentences. The use of propositional form might well be useful as a rhetorical tool for expressing an attitude as serious, and with unwillingness to tolerate the alternative, but if such a use came at the cost of causing all of us to go astray in our reasoning very often, it would not make sense for us to have adopted it as a regular feature of our language.

To address this objection, we must posit that using propositional form as a tool for the expression of attitudes is not done *merely* for the sake of expressing seriousness or unwillingness to tolerate the opposite. Rather, we use such form when the attitudes in question also bear some relation to one another that we want to use the propositional forms to help us track. The idea is that we can appropriate our logical skills for other purposes by using propositional form. Thus, far from being deceptive,

the fact that propositional form brings with it the tendency to think in accordance with certain patterns is simply another part of how propositional form can be of use to us.

But why should we suppose that doing so actually *works*? This point can be illustrated by an example, for the sake of which we will employ a rather broken variant of Blackburn's higher-order-attitudes picture, once again from his (1984). We will suppose, for the moment, that the purpose of our enterprise is to track the happiness of sets of commitments. Now the way this broken version of Blackburn's theory works is that, unlike in the real version, there are no particular rules as to which attitude gets which expression, or even which order of attitude gets which kind of expression. Thus, in this picture, "fish are purple" expresses approval of Kona coffee; "if fish are purple, then the sky is freedom" expresses approval of Sumatra coffee; and "the sky is freedom" expresses both disapproval of Kona coffee and approval of murdering short people.

Now imagine someone trying to use the "instrument of serious, reflective, evaluative practice"²²¹ that I have just described. I submit that it would not go very well for him, especially if he happened to be partial to both Kona and Sumatra coffees. Propositional reflection-- a phrase from Blackburn that refers to the idea that we use propositional form to help us track relationships between attitudes that are isomorphic with the apparent logical relations between the expressions of those attitudes²²²--

221 Blackburn (1984) p. 195

222 A note on terminology: the phrase "propositional reflection" can also be used to refer to the sentence used to express an attitude in propositional form within a system of expression-attitude isomorphism as

allows us to explain why this form of expressivism does not work nearly as well as Blackburn's. In the parody theory, there is a mismatch between when there exists (if you do not mind a play on the emotive sense of validity) a happy combination of propositions, and when there exists a happy combination of attitudes. In a better version of the theory, such as, if all goes well for him, Blackburn's own, there would be a happy combination of propositions exactly when there is a happy combination of attitudes.

A thought experiment of a different sort aims to clarify the same idea. Imagine a world in which there exist three lights, named A, B, and C. B is far and away the most important; we care greatly (let's say there is some ritual which is considered blasphemous unless B is lit when it is performed) about whether or not B is lit. Unfortunately, B is far away, and so we cannot see whether it is lit or not. To solve that problem, we built A, to serve as a kind of B-indicator. Unfortunately, our wizardry was not very good at the time, and so A suffers from various problems. For one, it suffers from an abundance of false negatives (though, thankfully, not false positives); so many so, in fact, that when A is not lit, we should consider that B may be either lit or not lit. And A's troubles do not end there. From a mechanical perspective, it is error-prone; periodically, it falls out of alignment, causing it to return an abundance of false positives. In addition, when A is out of alignment, the rate of false negatives increases to 100%: that is, it will always be off if B is on (and so the only time when *both* lights will be on is when B is on and A is in alignment). We can't fix it, but eventually, it will

described in the preceding definition. The phrase occurs throughout Blackburn's work.

wiggle itself back into the proper alignment, and will once again be free of false positives (though it will still have false negatives; when A is working properly, it is never on when B is not on, though it may be off when B is on).

One day, a clever wizard figured out that, although A could not be fixed, it was possible to detect when it was in alignment and when it was out of alignment. In order to display this information, he constructed a third light, which he named C after a departed relative, which shines when (and only when) A is detected to be in its proper alignment. C was, indeed, well-designed; had it been installed as the designer recommended, it would have been a perfect indicator of when A was in alignment and when it was not. Unfortunately, there was something of a mix-up in the installation of C, and its wires became crossed with those of A. So long as A was on, there was no problem; but when A was turned off, the extra power would flow into C with enough energy that C would remain lit regardless of the state of A.

Nevertheless, we were, in general, quite happy with C. To be sure, when A was off, we would still have no idea whether or not B was on; but when A and C were both on, we could know that B was on (and thus perform the aforementioned ritual safely). And so an official declaration of “that’ll do” was made by the Wizards’ Council, and for a while, everything was fine.

Over time, however, the clever wizard who created C fell into despair, because people were not calling C by its proper name. Most people, already confronted with the need to keep track of two different important lights with single-letter names,

balked at the idea of a third such light. Thus, in order to have a different kind of name for C, they began to call it “if A, then B.” This name, they found, not only helped them to keep track of the light and what it did; it also made it possible, once the lights received labels proclaiming their names, for them to get new light-watchmen to read the lights properly without any actual instruction as to how they worked, so long as the light-watchman had passed a course in elementary logic. As wizards willing to spend time teaching mere light-watchmen were in short supply, this development was most welcome.

In this story, there exists an interesting isomorphism between the formulae after which the lights are named and the illuminations of the lights themselves. We can see the isomorphism on the following tables:

Fig. 5.1 Lights:

<u>A</u>	<u>B</u>	<u>C, a.k.a. “A->B”</u>
O	O	O
O	-	-
-	O	O
-	-	O

Fig. 5.2 Formulae:

<u>A</u>	<u>B</u>	<u>A->B</u>
T	T	T
T	F	F
F	T	T
F	F	T

The renaming of C to “A->B”, in the story, enables the light-watchmen to make use of the isomorphism. A novice light-watchman who came to the job already familiar with the bottom table

would at the very least have a leg up in mastering the top one. How the lights are named is relevant here; had the lights had names of a different sort, the novice would probably not be so invited to attempt the application of the lower table.

Is this story plausible? Yes, there is a bit of a “jump” that the novice must make. But, at the same time, it is not an enormous one; it should not be an obstacle to any novice save, perhaps, one who is dense, painfully literal-minded, or willfully obtuse. In fact, if the lights were named as I have described, but were able to combine their states in different ways, so as to break the similarity between the two tables, confusion would ensue. It would then be *more* difficult for a novice to learn to read the lights than if they had been called, say, “horseradish”, “freedom”, and “7032”-- that is, the effect of the names is strong enough that an unexpected failure of isomorphism could lead us astray.

The doctrine of propositional reflection holds that a similar relationship exists

between the patterns of apparent logical relations between the expressions of attitudes and some sort of patterns-- it will be up to the theory to tell us what-- in the attitudes themselves, as between the patterns of apparent logical relations between the names of the lights and the patterns of illumination in the lights themselves. If this doctrine is correct, it explains, in part, *why* the adoption of propositional form for the expression of our attitudes can be useful to us. At the same time, the feeling that our expressivist theory would have us be doing something somehow illegitimate may be somewhat calmed by the idea that we are not just using propositional form for rhetorical effect. Rhetorical effect, mind, may still be a significant or even the main motivation behind the assumption of propositional form, but, as the example of a broken higher-order attitudes theory above shows, rhetorical effect without propositional reflection could easily lead to disaster.

One potential confusion should be addressed at this point, using the example of the lights. One who was troubled by the worries related to “confusion” or “deception” that we have considered might think that mimicking expressivism requires that we have some false beliefs about our own language-- that we have to believe that certain pieces of language are descriptive when they are, in fact, emotive. In accord with what I argued earlier, we could apply quasi-realism here. But the expressivist also has another line of reply to explore. For he has never actually claimed that such false beliefs are necessary. The case of the lights has, built into it, many different reasonable options for the “meaning” of A being lit, or of C being lit; I see no reason why a novice

who follows any of these would not be able to “get it”. Likewise, he could just as easily take C's name-card to be giving him a raw direction as to be making or suggesting any fact about lights or the patterns, significances, or purposes thereof.

5.2 Blackburn's Attempts at Propositional Reflection

Let us turn, then, to the place of propositional reflection in Blackburn's various theories. Our objective will be to figure out what, in each theory, is supposed to be on the attitude side of the isomorphism relation-- that is, what patterns are supposed to match up with the interrelations between the expressions. In both the “Moral Realism” (1973 [1993]) theory (if we take a sophisticated reading of the theory, on which his claims about the expression of facts about attitudes are themselves to be understood as made in the spirit of quasi-realism) and the *Spreading the Word* theory, the interactions of higher- and lower-order attitudes plays this role. To see how this works, let us consider the following argument:

- 1) If kicking babies is wrong, then kicking toddlers is wrong.
- 2) Kicking babies is wrong.
- 3) Therefore, kicking toddlers is wrong.

Let us, for the moment, grant Blackburn that what is being expressed in this argument is actually²²³:

223 H!(x) expresses approval of x (as in “Hooray, x!”); B!(x) expresses disapproval of x (as in “Boo, x!”); T!(x) expresses tolerance of x; vertical lines (“||”) are used to denote reference to rather than expression of an attitude; the semicolon (“;”) denotes combination of two mental states, along the lines of “having the latter upon possessing the former.”

4) H!(B!(kicking babies);|B!(kicking toddlers))

5) B!(kicking babies)

6) B!(kicking toddlers)

We might then make the following table:

Fig. 5.3

2)	3)	1)
T	T	T
T	F	F
F	T	T
F	F	T

What if we were to make a similar table for 4)-6)? In place of T and F for truth and falsity, let us use H and U, for “happily possessable given the other attitudes” and “only unhappily possessable given the other attitudes.” If all goes well, the chart should come out looking like this:

Fig. 5.4

5)	6)	4)
H	H	H
H	U	U
U	H	H
U	U	H

The asymmetry on this chart may seem a bit odd, but this can be taken care of by nuancing how we read “;”; if we take it as something like “having the latter given the former” rather than “having the latter and the former”, then the chart comes out right.

All is well, then; isomorphism has been achieved. Or has it? The tables here do seem to illustrate how the assumption of propositional form by an expressive language might help, both rhetorically and practically, in tracking and maintaining happy combinations of attitudes; this much of the theory seems to work. But in line with our original worry about this theory, we might be concerned about whether H and U possess the kind of stability that we need in order to make this kind of chart. For example, what if the individual in question also possessed an extremely strong disapproval towards having more than three attitudes, this one included? In this case, the chart would become:

Fig. 5.5

5)	6)	4)
H	H	U
H	U	U
U	H	H
U	U	H

The upper right cell of the chart changes from H to U, since three Hs, plus the dislike of having too many attitudes, would make four attitudes, which would be too many.

Yet the truth charts for the expressions will demonstrate no parallel effect; while there

can be facts outside the scope of a given argument that contradict one or more lines of the argument, overtly psychologistic effects like this one have no influence. One might resort to something like, “there can be at most three moral truths”, but even that does not quite work; such a fact would simply obliterate the top row of the table, whereas the psychological quirk that we have been discussing could render the top row unhappy without any “contradiction”.

Thus, although the higher-order attitudes approach can-- and probably should-- make use of propositional reflection, at the same time, the idea helps us to see what is wrong with the approach. The flaws do not mean that higher-order attitudes are completely unsuited to the use of propositional reflection; our tables show that at least in *some* circumstances, the device works well enough. What it does not do is work well enough to play the role of a driving force in our explanation. While we might see propositional reflections of higher-order attitudes arising derivatively, copying the maneuver from somewhere else, it seems implausible that we would start the practice of using propositional form just for a result of this quality.

Let us turn our attention, then, to our next candidate: Blackburn's later theory, from *Ruling Passions* (1998).

Blackburn initiates his discussion by introducing the idea of “emotional ascent”. The concept is as follows: attitudes come in a wide variety of types, which can be placed on a continuum by the degree of involvement they carry for the one who has them, ranging from simple throwaway preferences on one hand, to full-blooded

commitments on the other, attitudes that we not only have, but promote, defend, etc. Increasing complexity in a sequence of reactions drives progress along the continuum. In the simplest attitudes, the reaction chain has a length of 1: one is presented with a stimulus and reacts to it, and that's all there is to it. For many reactions, though, others see fit to react to the reaction; in which case, one might end the chain with a reaction about the appropriateness of the second reaction: *or*, the chain might continue, with further dispute, or concurrence, even, in some cases, leading to one encouraging others to share in one of the earlier reactions, to react negatively to those who do not, even to one becoming hostile to those who do not conform. At the far end, there are attitudes where, unless one is so numbed by exhaustion that one has become completely unreactive, one will continue the chain indefinitely with every available conversant, should the subject come up (or even without). These tend to be the clear, trademark moral cases. The longer the reaction chain, the greater the degree of emotional ascent involved in the attitude.

Blackburn tries emotional ascent as a means of placing, within our psychology, a system of interconnections reflected by the apparent logical relations between our attitude-expressions. The view is that we have commitments as described by the theory of emotional ascent; the makeup of these commitments leaves them not independent, but rather capable of some combinations and incapable of others. Blackburn claims that we may find ourselves “‘tied to a tree’: only able to endorse some combination of attitude and belief”²²⁴; what this means is that an attitude, for as long as we have it, can

224 Blackburn (1998) p. 71

restrict what other attitudes we are able to have. The key, though, is that Blackburn does not naively assert that, as a brute fact, that we cannot have clashing attitudes. Rather, the restrictions are built into the very attribution conditions of the attitudes in question.

In this view, the relevant attitudes are defined in part in terms of their inferential role. The theory of emotional ascent explains how such entities could come about: an attitude without the inferential role expected of its propositional reflection would not *be* the attitude expressed by that sentence, since the two would differ in their degree of emotional ascent. The isomorphism tables are forced to come out the right way, in that *making* them come out the right way is baked into the account of the complex sentences that the expressivist gives in terms of the attitudes expressed by those sentences. For example, suppose that “if murdering humans is wrong, then training your cat to murder humans is wrong” is not *just* $H!(|B!(mh)|;|B!(cmh)|)$, but rather, that it is expressing a similar but emotionally ascended attitude. Because the attitude is emotionally ascended, it commits us, as part of its attribution conditions, to accepting that training your cat to murder humans is wrong upon accepting that murdering humans is wrong. One who fails to accept the latter upon accepting the former cannot have the emotionally ascended attitude; if it was previously possessed, it must be withdrawn from.

And so, rather than 4)-6), we get, using “^^” and square brackets to indicate emotionally ascended attitudes:

7) $\text{H}[(\text{B}!(\text{kicking babies});\text{B}!(\text{kicking toddlers}))]$

8) $\text{B}!(\text{kicking babies})$

9) $\text{B}!(\text{kicking toddlers})$

...which results in the following table:

Fig. 5.6

8)	9)	7)
P	P	P
P	I	I
I	P	P
I	I	P

We switch from using “H” and “U” to using “P” and “I”: “P” indicates that an attitude is possible when other possible attitudes on the row are possessed without the violation of any other commitment, while “I” indicates that it is not. Now in our new chart, the upper right hand cell is stable; there is no chance for the value of the cell to be “overridden” without some contradiction. So, for now, it seems we have done better than we did with the higher-order attitudes view.

Another advantage of Blackburn's later view is that it integrates some of the rhetorical advantages of propositional form into the account of propositional reflection. The seriousness and unwillingness to brook divergence that attend to propositional form at the rhetorical level are the hallmarks of emotional ascent. And they are exactly what we expect from morality. Whether or not I am merely expressing an attitude when I say that “pushing unwary philosophers onto the subway tracks is

wrong,” I certainly do not want to present it as *merely* an expression of an attitude, nor to tolerate the presence of any other attitude on the matter anywhere in my vicinity. And thus I express my attitude-- my *commitment*-- as “pushing unwary philosophers onto the subway tracks is wrong” and not merely as “boo for pushing unwary philosophers onto the subway tracks.”

Unfortunately, this view will not suffice. Sonderholm (2005) raises a serious problem: a commitment is *active*, in a way that the expressions of moral attitudes seems not to reflect. For example, take Blackburn’s explanation for disjunction, as “a state in which *if* one side is closed off to me, I am to switch to the other-- or withdraw the commitment”²²⁵: holding a disjunction is explained as a readiness to *do* one thing or another. Now there *is* a reading of Sonderholm’s problem which is fixable; he argues that the commitment structure cannot capture cases such as $(p \wedge n \wedge q) \vee (p \wedge n \wedge r)$ implying p , which is a red herring, since it focuses on Blackburn’s specific example for disjunction. The structure of relevant commitments can be made more complex to capture what is desired; indeed, Elstein (2007) attempts to do just that. There is, however, a deeper reading of the problem from which the view cannot so easily recover: on the active-commitment model, something always has to happen; we must fulfill our commitment, or lose it. But any action requires some processing time; there is no incoherence in not having followed through *quite yet*, even though there would be incoherence in never following through (which would mean that you did not have the commitment) and yet having the commitment. The (supposed) propositional

225 Blackburn (1998) p. 71

reflections, though, are typically treated as having no such delay; indeed, logic is not something that has a *time* in any non-metaphorical sense. To use the metaphor of a gap, it seems that the *Ruling Passions* view attempts to close it altogether, where what we desire is for it to be bridged.

This disanalogy threatens to weaken the view's attempt to draw plausibility from its isomorphism claim. While it is built into the *Ruling Passions* view that, as we saw on our chart, the overall patterns of our ascended attitudes will match up with the apparent logic of their propositional reflections, we are still in danger of being misled by the surface forms. We know, or think we know, that it is possible to believe A and A->B but not (yet) believe B, or even be in any way disposed to believe B (this is easy enough; all that it requires is that the person not happen to think about A and A->B at the same time). Yet, on this view, someone who believed 1) and 2) would have to believe 3); if he did not, then we would have no grounds on which to attribute to him the emotionally ascended attitudes expressed in 1) and 2).

As with the previous view, perhaps the analogy is close enough; if a good-enough case can be made for expressivism on other grounds, even an imperfect analogy could be deployed to explain away apparent counterexamples from our practice, so long as the strain on one's credulity imposed by the analogy is less than the plausibility of expressivism earned for it by other arguments. That said, if we can come up with a theory with a closer analogy, our quasi-realism will be of more help to the expressivist. For the sake of exploration, we will assume the need for a resemblance

that is as useful as possible.

5.3 Another Propositional Reflection

An alternative theory aims to combine some of Blackburn's insights with a different theoretical structure, one where the relationship between attitudes that we are tracking with the use of propositional reflections is not a relationship between the attitudes themselves but something else²²⁶, in this case a further mental state definitionally dependent on our attitudes.

Suppose that, in addition to beliefs and attitudes, we have another kind of mental state—"commitments." Commitments relate beliefs and attitudes to actions; having a commitment to a belief or an attitude is to stand ready to act on it. A commitment, in this sense, is not itself a plan, but rather, the intermediary between other mental states and plans. The relation between commitments and plans is a close one; it is commitment that determines how our attitudes, hypothetical and otherwise, determine our plans. And like plans, commitments are restricted in how they combine by what is and is not compossible in the world. But committing is not, in itself, to make a plan; we are committed prior to (in both the temporal and explanatory senses) planning, and are all too often in the unfortunate situation of having to express our committed attitudes (in action) without any particular plan to see us through. We might find it useful to use plans as an explanatory tool in describing particular

²²⁶ The view of Gibbard (2003) shares this structural attribute, though it builds out of rather different materials. I don't seek to address Gibbard's position here, though I see no reason why some of its elements could not be incorporated into the pluralistic view that I ultimately endorse.

commitments, insofar as listing some of the plans that might result from the commitment in question might help our interlocutors to grasp the import of the commitment, and the interactions between the imports of different commitments; but of course an expository technique should not be confused with an exhaustive characterization.

A commitment in this sense is not meant to be an alien technical artifact. There is a natural sense of commitment that is, at the very least, in the vicinity of the one used here. We often speak of commitment to principles: an individual who is committed to a principle allows that principle to direct his actions, while the endorsement of a principle without commitment often comes to nothing. The sense in which I am using “commitment” is either identical to or a close relative of the sense in which it is used in these thoughts.

What we express, in this view, is still our attitudes. In order to be consistent about what we are expressing, let us allow that this is the case for both simple and complex sentences. Let us borrow Blackburn's higher-order attitudes for the purpose. Putting an attitude into propositional form expresses that the attitude is held in a committed fashion. Acting on a higher-order attitude is done by resolving one's other attitudes into accord with it.

Could some notion of involvement, where having one means having another, hold between different commitments? It would seem, given that one course of action can relate back to many different attitudes, overlap between commitments is to be

expected. Take, for example, a pro-attitude towards all coffee, and a pro-attitude towards Kona coffee in particular: attitudes that we might express as “All coffee is good” and “Kona coffee is good”, respectively. What is the relationship between these attitudes? Well, one might have the first without having any inkling that coffee is grown in Kona; one who has never heard of Kona coffee would never even consider the second attitude, even while possessing the first. At the same time, one who is selective about his coffee not only *could* possess the more specific example without the more general one, but is, in fact, quite likely to be in such a state. The two attitudes, then, are independent. But what of commitments to these attitudes? One who endorses any coffee that he comes across will endorse Kona coffee, should he encounter it; the actions that endorse and promote the consumption of Kona coffee are actions of endorsing and promoting coffee (note that we are talking about the kind of action that endorsing Kona falls under, not whether or not its promotion is good for the coffee industry as a whole—it is possible to carry out a pro-coffee action that, in fact, fails to promote coffee). A commitment to endorsing all coffee, then, includes a commitment to endorsing Kona coffee; the narrower commitment is a subset of the wider one. Even though the attitudes themselves are independent, their commitments are not. The propositional reflection of the relationship between the commitments is that “All coffee is good” implies that “Kona coffee is good.”

To make the point further, consider a disjunctive example:

- 10) Either this coffee is organically grown, or it's no good.
- 11) This coffee isn't organically grown.

12) It's no good.

Here we have a disjunction 10); since we are trying to be expressivists as straightforwardly as possible, let us consider it as expressing an attitude, in particular one concerning the relationship between believing that the coffee is organically grown, and disapproving of it-- perhaps something like approval of having at least one of the mental states expressed by the disjuncts. Committing to the attitude in 10) tells us to make plans based either on a belief that the coffee is organic or on a con-attitude towards it. At the same time, though, we have 11) expressing rejection of the belief that the coffee is organically grown. If we commit to 10), we will plan based on the belief that the coffee is not organic; once we do so, though, the only ways remaining that our commitment to 10) can be fulfilled are through plans based on a con-attitude towards the coffee. To plan based on a con-attitude towards coffee, in turn, is to commit to that attitude: in other words, to commit to 12). Thus, a commitment to both 10) and 11) is a commitment to 12). The propositional reflection of that fact is that 10) and 11) jointly imply 12).

This kind of theory holds that inferential import is a matter of "commitments". Commitments can involve one another; a Venn diagram could illustrate how multiple commitments intersect to commit us to further attitudes and rule out commitment to others²²⁷. This theory is unlike either of Blackburn's in that it, on one hand, allows acceptance of the premises of a valid argument without acceptance (yet) of the

227 Gibbard (2003) also uses the metaphor of the Venn diagram to illustrate the operation of a (quasi-)inferential structure.

conclusion, while at the same time allowing our expressions of attitudes to reflect a structure that is constantly present and immune to the kind of capricious attitude-targeting attitudes that plagued the higher-order attitudes view. We can be committed to an attitude in advance of when we actually come to possess that attitude, as in the case of the committed endorser of all coffee being committed to endorsing Kona coffee, even before he knows that such a thing exists; there is nothing in the theory to make it impossible to accept the premises of a valid argument and not yet accept the conclusion. Yet being committed to a conclusion is not an empty attribute; if one is committed, when the time of decision is at hand, what attitude will guide one's actions is already set, unless one withdraws from the commitment.

If we had to pick just one relationship to be what we track with the logic of our propositionalized expressions of attitudes, the relationships between the attendant commitments of the attitudes seems promising; it seems like the right sort of relationship for us to track in this way. Since it is possible to have a commitment without (yet) having the associated attitude, the view does not suffer from the same problem as the ascended attitudes view. At the same time, there are plain facts about the relationships between commitments, and so, unlike the older of Blackburn's views, the commitment-tracking view is not dependent on a relationship that is subject to psychological caprice. Isomorphism is achieved thanks to the appropriate acts of inference being built into the commitments associated with higher-order attitudes. This view, then, seems to be the best of those we have thus far considered.

On the other hand, what the commitment-tracking theory gains in terms of the ability to avoid the pitfalls that claimed its predecessors, it loses in elegance. Adding commitments to the theory effectively doubles its count of moving parts. Moreover, it claims that the fundamental concern behind the logic of moral arguments is keeping track of which attitudes we might simultaneously place in charge of our action-producing mechanisms and which we cannot. This concern *is* important, but is it really so important as to be ever-present in the consideration of moral arguments?²²⁸

Whether our combinations of attitudes are happy or unhappy (as in the higher-order attitudes view) is something that can and does always concern us, inner disharmony being, after all, a miserable state. The ascended attitudes view also passes this test, since by hypothesis we are always working with such attitudes in moral discourse (and when we are not, propositional form, and the attendant logical appearances, are not used), and serious concern is built into emotional ascent.

In other words, the simpler views, despite their failings, have more basic plausibility. They resonate with our everyday concerns in a way that the commitment-tracking view does not, or at least does not do as well. There is a way we can save the parts of the simpler views that we liked despite the mismatch, but it requires a different sort of theory.

5.4 The Pluralistic View

Let us call the “different sort of theory” the “pluralistic view” (the somewhat

²²⁸ Remember that importance was one of our original desiderata.

awkward “pluralistic” is used because the more natural term “pluralism” refers to an unrelated view in first-order ethics). The pluralistic view starts with the realization that the kinds of concerns that the mimicking expressivist appeals to are ones that overgenerate: if it is useful for us to appropriate propositional form to track and make serious one form of relationship between attitudes, why should we suppose ourselves to have left it there? Couldn't we do the same thing again with a different relationship? There seems to be no reason to suppose that there is exactly one relationship at the heart of propositional reflection. Not even theoretical simplicity favors such a view, since a one-relationship view would need to answer the question of what *stops* us from taking advantage of propositional form more than once.

The pluralistic view claims that there is no *one* relationship that is tracked by the logic of moral arguments. Rather, we use propositional form for our attitudes when there is some relationship that we are interested in tracking and making serious through the use of propositional form. On such a theory, we might also use propositional form when we feel like taking advantage of the general practice to try and further a particular conclusion even in the absence of any such relationship (just because we want the assumption of propositional form to be generally non-deceptive does not mean that we cannot or should not allow that we could sometimes, derivatively, use it in such a deceptive way). At the same time, there might be cases where a number of different relationships simultaneously contribute motivation for bringing our attitudes into conformity with a pattern isomorphic to the logic of their

expressions. Here, no one relationship would suffice to motivate propositional reflection. But since there is nothing to preclude using the tool of propositional reflection in cases of joint sufficiency between a number of relationships, and good practical reasons for us to do so, there is no reason to think that we would not do so. The theme of the pluralistic view is that propositional reflection is a tool that we possess. We do not restrict our use of this tool to the satisfaction of any one primary purpose; rather, we use it whenever we want, and can think of many different situations when it might be beneficial for us to do so.

The pluralistic strategy for embracing the doctrine of propositional reflection is to proceed by examples. Its explanations of the suitability of the surface form will proceed on a case-by-case basis, appealing to whichever sub-theories seem applicable. This might seem like cheating-- since the theory is essentially helping itself to the luxury of always having an explanation that works-- but, in fact, an explanation of this style can be argued against in the same ways as any other. If no explanation is forthcoming for an argument that seems sound, or if the only explanation on offer fails to deliver the intensity or strictness that we expect for the argument, our confidence in the theory could be compromised, just as with any of the non-pluralistic alternatives.

There are differences, of course. The pluralistic theory has the option of continuing to search for alternatives, rather than simply taking the blow directly. If one explanation fails, it can be claimed that it was the wrong explanation, rather than admit a problem with the entire picture. In addition, if an argument fails to be

amenable to any of many options, we might begin to wonder if there is not something wrong with the form or a premise of the argument (especially if a conditional premise is present). Nonetheless, the possibility for failure still exists; the theory has not rendered itself tautologous. It simply has an easier task facing it than its non-pluralistic competition.

What the pluralistic theory cannot do is offer a single blanket explanation, to handle not only all of the cases at hand but also any others that might be encountered. It does not present a global account capable of settling, once and for all, the question of propriety for the transmutation of attitude-expressions into propositions. This might seem like a weakness, but, once again, the overgeneration argument is relevant. A single, global account does not make for a better theory when we have reason to believe that no such thing exists. It would be quite odd-- and require explanation-- if, having come up with the useful, powerful device of propositional reflection, we decided to pick one of the possible applications of the technique, use it, and then stop forever. And even if we did stop, and are using propositional reflection in only one case, there does not seem to be anything limiting it to that. In such a case, we might well come up with a story that can be told across all instances of attitude-expressions taking propositional form, only for it to start failing later on. Even if there *were* a true one-relationship view, it would only be true accidentally, while the pluralistic view faithfully tracks what is necessary and possible. The pluralistic view gives us no way of identifying the relevant relationship for every possible moral argument because there

is no way of doing so; thus, this aspect of the view is a strength, not a weakness.

5.5 The pluralistic view and Richards

The pluralistic view can be bolstered by some of Richards' rhetorical considerations. We needn't accept his philosophical views to benefit from his observations as to the various mechanisms which affect the expressive content of terms or phrases; we might be reluctant to call every bit of expressive content a "meaning" as freely as Richards does, but how we label a particular rhetorical effect does not affect whether or not that effect exists²²⁹. Furthermore, if a rhetorical effect exists, it has the potential to influence the outcome of a practice-- such an effect can have philosophical relevance for theories that cite practice as part of an explanation. Mimicking expressivism is just such a theory; it cites practical benefits for the adoption of propositional form in the expression of attitudes. Thus, we need not accept that the kinds of effects examined by Richards are to have any standing in the theory of meaning in order for them to be of relevance to mimicking expressivism.

The pluralistic view holds that attitudes can relate to one another in a number of different ways suitable for the motivation of propositional form. If we introduce observations from Richards, we can bolster this view by noting that multiple uses of

²²⁹ This is not to say, of course, that how we name something cannot have a rhetorical effect of its own, or that this further rhetorical effect cannot bolster or undermine the effect which it names-- because these are certainly possible. Nonetheless, in our soon-to-be-cruelly-terminated state of pre-quasi-realistic innocence, we can permit ourselves some expressive inaccuracy in exchange for descriptive accuracy. We will see, in the next chapter, that such a maneuver is not half as innocuous as it appears; our considerations there will, I think, vindicate Richards' terminology, though not the theory of language by which he himself justifies it.

propositional reflection are not independent, but are mutually supporting. That is to say, our various employments of propositional form for the expression of attitudes interanimate one another. We can apply propositional form to the expression of attitudes to import the serious and steadfast nature of description to that expression; having done so, we can then use propositional form to import some of the serious and steadfast nature of the previous expression of an attitude to another expression of an attitude. Each use of propositional form reminds us, perhaps not consciously, but emotionally, of the others.

Furthermore, given the intensity, and perhaps emotional ascendance, of many of the attitudes involved with morality (think about, for instance, our disapproval of setting small children on fire for entertainment), it seems to stand to reason that the use of moral language should have unusual power as a form of expression. Expressive content is stubborn even in the face of changing (descriptive meaning); it seems that at least some expressive content can attach to a word rather than to the word's descriptive content. Stevenson's persuasive definitions rely on this effect; when a term receives a persuasive definition, its descriptive meaning is changed, but the expressive content, having become attached to the word itself, remains, and thus becomes associated with the new descriptive meaning. We will return to the recalcitrance of expressive content, with examples, in the next chapter; for now, let us note that a pluralistic theory can exploit this effect.

Our attitudes, on a pluralistic theory, can relate to one another in many

different ways; you can, on such a theory, have two moral *modus ponens* arguments where the relationships between the attitudes expressed in the premises and those expressed in the conclusions are not the same. For example, consider:

- 1) If kicking babies is wrong, then kicking toddlers is wrong.
- 2) Kicking babies is wrong.
- 3) Therefore, kicking toddlers is wrong.

- 4) If kicking babies is wrong, then kicking premature babies is wrong.
- 5) Kicking babies is wrong.
- 6) Therefore, kicking premature babies is wrong.

An advantage of a pluralistic theory is that it can admit that the attitudes expressed in 1) and 2) do not relate to that of 3) in the same way as those of 4) and 5) relate to that of 6). 1)-3) seems like a good candidate for analysis by the ascended-attitudes approach of *Ruling Passions*. 4), on the other hand, seems like an inferior candidate for emotional ascent; it needn't be the same kind of serious, conformity-forcing attitude as 1). The pluralistic theory can admit that 4) is unlike 1), and yet claim that 4)-6) works as an argument nonetheless. The similarity in expression is justified by the functional similarity, even though the mechanism underlying the function-- emotional ascent in the first case, and possibly commitment involvement in the second-- is different.

We might wonder how it is that on a pluralistic theory, complete seamlessness in functional role can be maintained, and why we do not see moral arguments, especially those which are reflections of interactions between attitudes that are less perfect in their isomorphism, begin to diverge away from the patterns of reasoning suggested by their expressions. The recalcitrance of expressive content can answer this worry. We should expect, given the persistence of expressive content in the face of changes in descriptive meaning, that valid forms of moral argument will establish patterns of attitude-acceptance that are themselves resistant to change. An expressive force becomes attached to the role of being the premise or conclusion of a moral argument by the various forms of attitude-interaction that we track, and, once thus attached, it is sustained by the common presence of moral language. We need not have any pre-philosophical access to the tracked relationships at all; we become conditioned, when we are enculturated into moral practice, to have attitudes in accordance with the correct patterns, and that is all there is to it.

One strength of the pluralistic theory is that there needn't always be an underlying relationship between the attitudes in question. Once the practice of moral argument has been established, there is nothing to stop individuals from co-opting moral language for their own rhetorical purposes. If no appropriate relationship can be identified between the attitudes expressed by the lines of a moral argument, it may be that there simply is not one; the pluralistic theory is able to accept, and even expect, that possibility. The stubbornness of expressive content is sufficient to sustain such

parasitic cases.

Continuing in the spirit of the fast track, we can extend this solution downwards to cover embeddings. Our various models, drawn from Ayer and Stevenson and various generations of Blackburn, with the addition of the commitment-involvement view, demonstrate the need for and legitimacy of moral sentences that play the functional roles illustrated by their propositional forms. Once propositional form has been adopted, any form of embedding is available, even if no analysis in terms of attitudes is possible. In such cases, we use the sentence for its functional purpose, and there is nothing more to it than that it is an assemblage of words that performs that purpose. Its ability to perform that purpose is, in turn, sustained by the general practice of ordering our attitudes by their propositional form, which continues to sustain itself according to the uses outlined by the models, as well as the more general forms of rhetorical effect that we have already noted. The stubbornness of expressive content binds the effectiveness of the proper and defective cases together.

It would be convenient if no such sentences existed-- that is, if a satisfactory secondary explanation could be found for all the constructions of moral language-- but the universe is under no obligation to refrain from containing occurrences which offend our sense of theoretical beauty. The pluralistic approach is compatible with such raw, functional sentences, and can make explanatory resort to them if necessary. In this manner, it achieves the fast-track's goal of avoiding the Ptolemaic drudgery that Blackburn (1988 [1993]) feared; the expressive theory can be established without

having to have an account of every possible embedding.

The stubbornness of expressive content is important. So far, we have not examined it in much depth, but have, rather, focused on its effects. The next chapter, in dealing with quasi-realism, will provide us with an opportunity to return to the question of expressive stubbornness, allowing us to examine the ways in which expressive content resists various forms of verbal manipulation in more detail.

Chapter 6. Quasi-Realism

“Quasi-realism” is the practice of adopting, after some kind of manipulation to ensure compatibility, realist meta-ethical claims into an expressivist theory. Varieties of quasi-realism can be distinguished by what sort of manipulation they use.

Let us call “simple quasi-realism” the use of minimalism in pursuit of quasi-realism. The simple quasi-realist will thus claim, for instance, that truth is minimal, and so expressions of attitudes can be true or false²³⁰; that being a proposition or a property or an attribution of a property is also minimal, and so a moral sentence can very well assert a proposition attributing a moral property to some object; and so forth. We can see the attraction of quasi-realism in that it lets us say such things; we can be an expressivist, the quasi-realist contends, and still vindicate such ordinary practices as calling a moral statement true, or a fact. “Nuanced quasi-realism”, by contrast, is quasi-realism that proceeds by means of *expressivism*-- that is, it seeks to render at least some of the distinctive claims of the realist position in meta-ethics compatible with expressivism by offering expressivist analyses of those claims.

Before proceeding, we must disambiguate. Minimalism, as I mentioned, permits many variations, but there is one particular distinction that we need to keep track of here. Some writers on minimalism, both for and against, conceive of it as a way of

²³⁰ Truth minimalism, or deflationism about truth, comes in many forms; for our purposes, we needn't focus on any one in particular. Let us, instead, take as an example the disquotational account-- “s” is true just in case s, and that's all there is to it-- and couple it with the proviso that similar ideas, so long as they remain metaphysically lightweight, should be permitted to use the same label.

characterizing the concepts in question; a concept can be either minimal or robust²³¹, but not both. Others see minimalism as offering a sense of the relevant terms, but also that they may also take robust senses²³². Let us call the former category *hard* minimalism, and the latter *soft* minimalism.

The central idea of this paper is that simple quasi-realism faces a dilemma. Either it must endorse soft minimalism, which seems untenable, or endorse hard minimalism and pay a substantial cost. I will argue that nuanced quasi-realism, on the other hand, is a stronger position. True, nuanced quasi-realism must make a similar choice-- between, on one hand, an approach that parallels hard minimalism by chooses a certain number of realist-sounding claims to analyze as always expressive and accept without qualification, and denies the rest, and on the other hand, an approach that parallels soft minimalism by positing that many or even all of the realist's claims can take distinct expressive and descriptive senses, where they are to be accepted in the former but denied in the latter. My goal in the latter part of this paper will be to show that the expressive equivalent of soft minimalism can defend itself against the kinds of objections that are brought against soft minimalism itself, and can therefore provide us with a useful and workable form of quasi-realism.

231 This sort of thinking is found in *e.g.* Dreier (1996, 2004), Boghossian (1990), Wright (1987, 1992), and Sinclair (2007).

232 *e.g.* Stoljar (1993), Lenman (2003), and considered (though not accepted) in Holton (2000), Dreier (2004), Sinclair (2007), Jackson, Oppy, and Smith (1994), and Asay (2013). Interestingly, this also seems to be the position adopted by Ayer in his (1954, p. 231), though he had earlier, in his (1936 [1952]), considered *and rejected* applying an early form of minimalism to moderate his position. Such a move is in keeping with the presence of other advanced features in Ayer's later works, such as the application of expressivism at the meta-ethical level in his (1954, p. 248) and his (1984, p. 34), a move that mirrors Blackburn's position on mind-independence in his *e.g.* (1988a [1993]), (1984), (1998), and which will be involved in nuanced quasi-realism.

6.1. Some difficulties for quasi-realism

Let us start our examination with a study of the problems that a successful quasi-realism must overcome. First, all quasi-realist versions of expressivism face the question of how they are to distinguish themselves from realism. We can see the problem most clearly for a simple quasi-realist who employs hard minimalism. Such a theory works by dividing the realist's claims into two groups. One group of claims are to be analyzed in accordance with hard minimalism and thereby agreed with. The other group are not to be so analyzed, and are instead to be denied. Putting more claims into the first category will help the theory to more closely copy realism, and thereby answer the charge of implausibility. Putting claims into the second category allows the theory to differentiate itself from realism.²³³ A theory that puts all of the realist's claims into the first category is itself realist, though also minimalist. An expressivist theory that puts all of the realist's claims into the second category is still expressivist, but not quasi-realist. One way for the realist to collapse the expressivist's position, then, is to argue that the expressivist is committed to minimalism on the claims in the second category, and that his theory is consistent with those claims when they are interpreted minimally. The second half of this procedure is generally easy, so

²³³ For an example of a hard theory in action, consider the "ecumenical expressivism" of Ridge e.g. (2006a, b). Ridge suggests that the quasi-realist should accept that the sentences of moral discourse express representational states, but should claim that these states are such that they do not necessarily bequeath the truth value of their representation upon their expressions. The initial acceptance makes the position palatable, or at any rate less *unpalatable* than it would have been had it denied what does appear to be a feature of moral discourse. The subsequent qualification locates the difference between realism and quasi-realism.

we will focus on the first.

We can distinguish two ways of collapsing the simple quasi-realist theory that uses hard minimalism. One is to argue that minimalism is the correct theory of the concepts that the quasi-realist placed in the second category. Call this the *aggressive collapse strategy*²³⁴. The other is to argue that, whether or not minimalism is the true theory of the concepts in the second category, minimalism on those concepts follows from the quasi-realist's other commitments, typically his commitment to the minimalism of the concepts in the first category. Call this the *creeping collapse strategy*²³⁵.

My focus here is on the creeping collapse strategy. While I think that the sort of considerations that I will present make the aggressive collapse strategy less appealing, it is not my aim to rule out every possible aggressive collapse argument. That, after all, would involve creating a blanket ban on all unwanted minimalisms, which is too ambitious of a task for me to attempt here. Insofar as the theory I aim to construct proves attractive, it might count against the sort of considerations that the aggressive collapse strategist presents, simply by virtue of offering an alternative theory with its own advantages. But I cannot, within the scope of this project, address every possible argument that an aggressive minimalist might deploy against the position. I can only hope to defeat the more modest creeping minimalism, which does not rely on pressure from outside arguments.

234 e.g. Wright (1996), Divers and Miller (1994, 1995)

235 e.g. Dreier (2004)

If we are aiming to defuse the worry from creeping minimalism, we should look more closely at how the objection is supposed to work. Developing the problem, as it turns out, takes a bit of interpretive finesse. Dreier (2004) presents it as, first of all, not being a problem for expressivism, but for philosophers to solve in pursuit of reducing their own confusions²³⁶. He also presents it not as a logical consequence of theories, but as a result of historical trends²³⁷ that have encouraged expressivists, in the name of plausibility, to keep applying more and more quasi-realism, which we have often done in the form of simple quasi-realism. As expressivists go minimal on more and more claims, the story goes, it has become harder and harder to see them as having a distinct view. Formulated thusly, though, the argument is only a dilemma for the expressivist insofar as he is compelled to apply too much quasi-realism on pain of implausibility resulting from whatever claim he denies. If he does not mind denying some of our ordinary claims, or can make the point that the claims are not, in fact, ordinary (and how often do members of the folk who are not real estate agents talk about properties, let alone moral ones?), then he can escape.

That said, the problem can be sharpened into a genuine objection. Doing so requires adding a component from Divers and Miller (1994), Sinclair (2007), and Chrisman (2008). This component is the idea that minimalisms are not independent; the expressivist is not free to pick and choose as he pleases. Rather, the minimalisms are interconnected, and to other notions besides. For instance, as the aforementioned

236 Ibid p. 31

237 Ibid p. 25

point out, if the expressivist uses minimalism as a way of rendering moral statements truth-apt, and to assert is to make a truth-apt statement, then the expressivist is committed to the view that moral statements are assertions. If to assert is to express a belief, then moral statements express beliefs. And thus, the above authors argue, the expressivist must either allow his view to collapse, or else extend his minimalism to these other concepts as well. Thus can Dreier's statement that "Once Minimalism gets creeping, it's hard to see how to stop it"²³⁸ be converted into an argument: the simple quasi-realist who uses hard minimalism needs to divide claims into ones where he applies minimalism and ones where he does not, so that his view will be both plausible and distinct, but it seems like there is enough interdefinition between the relevant concepts that the quasi-realist cannot keep things that he would like to place on the "denial" side of the divide from migrating to the other list.

It should be noted that those who consider this argument, such as Dreier (2004) and Asay (2013), have not seen the difficulty as only targeting simple quasi-realism that uses hard minimalism. They also wish to close off soft minimalism, of which they both choose Timmons (1999) to be the representative, though in doing so they continue a debate which previously involved Boghossian (1990) and Stoljar (1993). Soft minimalism is not subject to the creeping minimalism worry without further argument, because it is not an either/or proposition; the simple quasi-realist using soft minimalism may put every claim of the realist's in his list of claims to accept, and at the same time put as many of them as he pleases in his list of claims to deny, so long

238 Dreier (2004) p. 29

as he accepts them in the minimal sense and denies them in the substantial. The lists can be stable so long as the minimal claims are only interdefined with other minimal claims, and the robust claims are only interdefined with other robust claims.

Plausibility can be maintained so long as it can be argued that practice needs no more than the ordinary sense to avoid error. Distinctness can be maintained so long as it can be argued that the realist is committed to the relevant claims in the robust sense.

Of the arguments that have been deployed against soft minimalism, some are better than others. Boghossian (1990) and Asay (2013), for instance, both claim that terms like “truth” present a clear appearance of univocality, rendering soft minimalism implausible. But I see no reason whatsoever to agree with them about how truth and its compatriots *seem*. On the contrary; it *seems* to me like it is a concept that is used in many ways, and if there is a univocal core *meaning* underlying some or all of those uses, we shall have to do work before we can see it. For instance, one might respond to the claim that obelisks look best in obsidian with “that’s true”, while simultaneously believing, apparently without irrationality, that the coloring of obelisks is a matter of mere preference and thus not something where there is a “truth of the matter.” Now “true,” “belief”, etc. might not look *particularly* ambiguous, but that does not matter. Some of the early figures of relevance to the expressivist tradition, such as Stevenson (1944), and his influence Richards (e.g. 1936) saw ambiguity everywhere, and wondered whether single, determinate meanings, outside of parts of the hard sciences, were even worth looking for. For present purposes, we needn't worry about whether or not we

agree with them; the mere fact that their views exist, without being entirely ignorant or insane, is enough to show that a concept appearing ordinary is not the same as it appearing univocal (though it may well be). So, whether we are ultimately soft minimalists or not, the univocality of “true” and the other concepts that are on the line will have to be settled by theory, not by appearances.

The answer given by Dreier (2004) and Holton (2000) is better. For the soft minimalist to have a position, he must be able to differentiate minimal and robust senses of the relevant terms. There are various different forms of minimalism on offer, but specifying the robust sense can be more difficult, especially if the expressivist is an expressivist because he does not find realism to make sense in the first place. Not only that, but the robust sense must be something which the expressivist can deny but to which the realist is committed, and the realist himself might not put any effort into proclaiming his robustness²³⁹. Dreier (2004) and Asay (2013) both discuss this issue using Timmons' (1999) notation of terms in all-caps to denote robust senses; hence assertion and ASSERTION, belief and BELIEF, properties and PROPERTIES, etc. The question then becomes how the soft minimalist can distinguish ASSERTION from assertion, BELIEF from belief, and so forth, without violating the provision that the all-caps term must be something that the realist is committed to and the expressivist can deny.

Soft minimalists have, of course, had things to say by way of elaboration.

²³⁹ Indeed, he might even go so far as to follow Wright in endorsing minimalism, and thereby attempt an aggressive strategy against the expressivist.

Lenman (2003), for instance, distinguishes minimal and robust senses of truth-aptness by recourse to whether or not the sentence could be transformed into a non-truth-apt form. What we have here is, after all, a challenge, not a ban. In this regard, it is similar to the argument from creeping minimalism against the hard minimalist: the hard minimalist faces the challenge of identifying a place for minimalism to stop. There, too, we see attempted solutions; Asay (2013), for instance, thinks that the idea of “truthmakers” allows us to identify a reasonable endpoint. In neither case is the expressivist required to give up without a fight.

Nor, though, is the challenge trivial²⁴⁰. The proponent of the collapse objection might, for instance, reply that Lenman does not try hard enough to find non-truth-apt transformations of straightforward descriptive sentences (compare “treedog!” upon seeing a dog in a tree). Alternatively, he might point out that achieving reducibility in Lenman's sense commits the expressivist to a very simple view of the attitudes involved in morality, one from which even Ayer (1936 [1952]) and Stevenson (1944) both took pains to distance themselves. If the expressivist wishes to hold, with Ayer, that moral attitudes are distinctive²⁴¹, or, with Stevenson, that there is no precise emotive reduction of moral sentences²⁴², or, with the later Blackburn, that the attitudes expressed in moral practice differ from simple exclamations by virtue of emotional

²⁴⁰ The replies given here, I should note, are not meant as anything more than tentative sketches.

Ultimately, I think that nuanced quasi-realism's advantages will be enough to render it a superior reply to the creeping minimalism challenge even if the minimalist is still able to offer an alternative. Thus, I do not mind if the simple quasi-realist is able to win here; in fact, it might be convenient for the expressivist to have a live minimalist option as a potential fallback.

²⁴¹ Ayer (1936 [1952]) p. 113

²⁴² Stevenson (1944) p. 82

ascent, and therefore permit no strict reduction²⁴³, then the only way to accomplish Lenman's reduction is to take Ayer's strategy of appeal to magical punctuation marks, such as a special !!²⁴⁴; but if we allow !! the power to transform whatever we feed it into clearly non-truth-apt form, without the need for equivalence between what comes out aside from the !! and what went in, then we will be able to feed it any straightforward descriptive sentence as well. So I am not sure that Lenman is able to make the kind of distinction that he wants to make.

To summarize this section: versions of quasi-realist expressivism using either hard or soft minimalism face the task of preserving the distinctness of their view from realism. The hard minimalist straightforwardly denies some of the realist's claims, while giving a minimalist analysis to others. The realist can respond by trying to show that the quasi-realist's chosen applications of minimalism commit him to minimalism on those claims that he wishes to deny, and that once those claims have been analyzed minimalistically it is no longer plausible for the quasi-realist to deny them. The soft minimalist, on the other hand, distinguishes minimal and robust senses of the relevant claims, accepting them in the minimal sense while denying them in the robust sense. The realist can respond by trying to collapse the distinction between the two senses.

6.2. More difficulties for quasi-realism

If we were to limit our considerations to the concern raised in the previous

243 Blackburn (1998) pp. 12-14

244 Ayer (1936 [1952]) p. 107

section, it would most likely seem like hard minimalism is the most promising option. What the hard minimalist needs to defend against the objection from creeping minimalism is a way of drawing a line in the sand, so to speak: he must be able to come up with a rule or principle to separate realist-seeming metaethical claims that should be understood minimally from those that should be understood robustly and while I am uncertain as to how he might do so, at the same time, the idea that there is so much interdefinition between the relevant concepts that dividing the minimal from the minimal is impossible needs more support before we can consider it proven. Unfortunately for the hard minimalist, there is another kind of objection that gives us strong reason to search for an alternative.

Let us introduce this second concern by means of an illustration. Suppose we follow Asay (2013) and allow the quasi-realist to differ from the realist in his account of the truthmaking relation. Asay's view, in our taxonomy, is classified as a form of hard minimalism; the quasi-realist, on this view, achieves distinctness when he denies that which the realist (presumably) would like to say about what makes (true) moral judgments true. Faced with such a view, the realist might challenge whether the quasi-realist's alternative is *adequate*-- that is, if the quasi-realist can deny that there is a mind-independent truthmaking relation between the moral truths and reality without making us wonder if he isn't taking away an important contributor to the seriousness of morality. Contrast Asay's position on truthmaking with Blackburn's position on mind-independence: surely, speaking in the spirit of Blackburn, we cannot accept, for

instance (using Asay's example), that “kicking dogs is wrong” is only made true by the fact that it causes them pain because of our attitudes and the projection of attitudes onto the world. That would be to allow our attitudes to play about with our reason-relations on a whim, which we cannot, plausibly, countenance. The recourse to truthmakers pushes the undesirable mind-dependence back by one level of explanation, but that is all; it does not eliminate it, and thus also does not allow the expressivist to take onboard his necessary burden of realism while not sinking into realism himself.

The basic problem that faces Asay's view-- and all other quasi-realisms-- is that achieving the quasi-realist's goal seems to be an all or nothing proposition. In order that we might understand why that is so, we must first remind ourselves of the reason why an expressivist might choose to become a quasi-realist. The reason, quite simply, is that realism seems right, at least before we think too much about it. For example, we react negatively to claims that moral arguments are not really capable of validity, or that there are no moral properties that really exist in the world; we take our negative reaction to such claims to be an intuition telling us that those claims are mistaken, and that there really are logically valid moral arguments and moral properties that exist in the world. We also want consistency in our picture of language; moral sentences look like ordinary descriptive sentences, and moral arguments like ordinary arguments, and so it offends our sense of neatness to have to say that moral language works in a fundamentally different way. To an extent, the latter concern can be assuaged by a

story, told at the practical level, of why we might want to use moral language to express our attitudes as if we were describing the world. But even then we might wonder as to whether such a practice would be able to persist while relying on our continuing to misunderstand the nature of our own language. The purpose of quasi-realism is to allow the expressivist to bring his view into conformity with the realist appearance of moral language, rather than having to accept the considerable burden inherent in arguing against the veracity of that appearance.

Right away, we can see one area where we might worry about hard minimalism. Hard minimalist quasi-realism encourages the expressivist to preserve his identity by denying some of the realist's claims. Yet, insofar as he does so, he will be going against the appearance of realism. Whatever distinctness his view possesses takes away from the extent to which it is able to be quasi-realist, and whatever extent to which it is quasi-realist takes away from its ability to remain distinct. Quasi-realism might not *have* to be perfectly faithful to the appearance of realism, but it would be nice if it could, since each denial comes at a cost of plausibility. Our unease with Asay's view is an example: in disagreeing with the realist about truthmakers for moral claims, that view seems to give us an incorrect answer to the question of mind-dependence. For an ambitious quasi-realist who seeks to render his view as plausible as possible, then, it might make sense to look into alternatives to hard minimalism.

Another, potentially more serious problem for hard minimalism presents itself when we consider a variation on the creeping minimalism argument, which we will call

“creeping boldness” (since the expressivist must be bold to deny openly that which the realist asserts). The danger of creeping minimalism is that quasi-realist expressivism will collapse into realism. But there is also the worry that hard minimalist quasi-realist expressivism will collapse into plain expressivism. This worry exists because of the same connections between the various components of realism as motivated the creeping minimalism worry; where there exists an inconsistent triad between a claim the quasi-realist accepts (though analysis), one that the quasi-realist denies, and a definitionally motivated biconditional between the two, resolution may be achieved by the expressivist coming to accept that which he denied (creeping minimalism) or by his coming to deny that which he accepted (creeping boldness). Avoiding both creeping minimalism and creeping boldness requires taking a third path: denying the biconditional to form a break point between those claims that the quasi-realist analyzes and accepts and those that he denies. But it is difficult to see how this could be done.

Examples can help us understand this argument. Dreier (2004) divides the continuum of positions between the boldest expressivism and full-blooded realism into four kinds of positions²⁴⁵: “old emotivism”, which denies the existence of moral truths, facts, and propositions, also denies that moral statements are assertions or expressions of beliefs, and further denies that moral statements describe or express representational states²⁴⁶; “new expressivism”, which accepts moral

²⁴⁵ The chart on page 28 of Dreier's article is helpful in keeping track of the various positions, though, as that article is readily available, I will not reproduce the chart here.

²⁴⁶ Dreier bases this position off of the 1930s version of Ayer's emotivism.

truths/facts/propositions but maintains the other denials of old emotivism²⁴⁷, “cognitivist expressivism”, which is similar to new expressivism except that it now accepts, rather than denies, that moral statements are assertions and express beliefs²⁴⁸, and “raving realism”, which accepts all of the claims that old emotivism denies. Dreier's worry is that it seems like cognitivist expressivism will have a hard time distinguishing itself from raving realism. But there is a similar worry in the other direction. As Dreier (2004) tells the story, cognitivist expressivism emerges out of new expressivism via the denial of the idea that assertion and belief are conceptually linked to description-- thus does the cognitivist expressivist attempt to earn the right to claim that the sentences of moral discourse are assertions and express beliefs, even though they are neither

247 Dreier sees this position as modeled after those of Gibbard and Blackburn, though I think those positions work much better if read as employing soft nuanced quasi-realism, for reasons that will occupy much of the remainder of this chapter. A word of warning, however: for reasons that we will not get around to discussing until later in this chapter, a soft nuanced quasi-realist will often sound like a hard nuanced quasi-realist, to the extent that all but the most careful reading is likely to confuse them.

For example, in his 1988b paper “How To Be an Ethical Anti-Realist,” Blackburn addresses an objection from Cassam (1986), who alleges that Blackburn’s quasi-realist response to the question of mind-dependence improperly disregards the “external” reading of the question. Blackburn, in response, seems to deny that there is an external reading (he writes: “there *would* be an external reading if realism were true”; Blackburn 1988b [1993] p. 173) at all. But we must be more careful in how we read him. His response to the charge of willful deafness is that “there is only one proper way to take the question” (ibid); in the next paragraph, he writes:

This is not, of course, to deny that ‘external’ questions make sense-- the projectivist plus quasi-realist package is an external philosophical theory about the nature of morality. But external questions must be conducted in a different key once this package is brought in. (ibid)

This quote suggests that we should not dismiss the “proper” in Blackburn’s answer as a disposable word for taking care of trivial responses, but that we should treat it as something more substantial. The suggestion that we must proceed in a “different key” favors an interpretation on which Blackburn sees only a single reading as being available not because he endorses hard quasi-realism but because there is some kind of other consideration that can, at least sometimes, rule out the external reading from even being part of a debate. Unfortunately, Blackburn does not tell us what kind of other consideration can do that. A significant portion of the remainder of this chapter will go towards rectifying that omission.

248 Dreier models this position off of Timmons (1999). The later Horgan and Timmons (2008) is also a good read.

descriptions nor expressions of representational states. This strategy is the right way of going about drawing a line between what the quasi-realist should accept and what he should deny-- that is, it resolves the inconsistent triad by denying the definitionally motivated biconditional. But it's also easier said than done. Majors (2008) argues that any supposed non-descriptive belief is not, in fact, a belief at all, since there are no other grounds-- and in particular no phenomenological grounds-- by which it might be justifiably categorized as such. For my part, I find it hard to understand the notion of a non-descriptive belief because both minimalism and boldness can creep so easily. If it is given that such a thing is not descriptive or representational, it seems tempting to want to call it something other than a belief. At the same time, if it is given that a state is a belief, it is tempting to allow whatever is believed by someone in that state to count, at least minimally, as the descriptive or representational content of that belief. As much as the cognitive expressivist might want to argue against the connection between belief and representation, doing so seems to be, at best, an uphill battle.

So far we have looked at the problems that boldness and creeping boldness create for the hard minimalist. These sorts of worries are not exclusive to hard minimalist expressivism, though; a similar problem faces soft minimalism. Like hard minimalism, soft minimalism seems to risk, first of all, implausibility from when it does robust meta-ethics and disagrees with the realist, and secondly, a collapse into either realism or plain expressivism. We considered the latter earlier, as a branch of the creeping minimalism problem; unless the quasi-realist can make the difference

between minimal and robust senses clear, he is in danger of becoming a realist (if the only coherent sense is the minimal) or a straightforward expressivist (if the only coherent sense is the robust). Let us now consider the way in which it stands to fail at achieving its goal of satisfying our realist intuitions.

Soft minimalism tend to be vulnerable to an attack by stipulation. There is nothing to stop the realist from simply stipulating that he is speaking in the robust rather than minimal sense, and when he does so, he does not seem to diminish thereby the realist appearance of moral discourse. Adding the proviso “in the robust sense” does nothing to take the sting out of losing moral properties or the validity of moral arguments; our intuitions seem to remain steadfastly in favor of realism even when we make explicit that we are talking about *robust* realism (whatever that winds up being). Quasi-realism that uses soft minimalism, then, is in danger of being useless; the quasi-realist expressivist gains nothing by endorsing realist claims in the minimal sense if his denials of those claims in the robust sense constitute just as much of an affront to our intuitions as did the original unqualified denials of the plain expressivist.

All in all, it seems that both forms of minimalist quasi-realism fall short of achieving what the quasi-realist wants. Neither hard nor soft minimalism seems successful in relieving the expressivist of the burden of plausibility that he carries from denying initially attractive realist claims. Furthermore, we can create a variation on the problem of creeping minimalism, the problem of creeping boldness, which threatens the quasi-realist's ability to be anything more than a plain expressivist. Given these

problems, it seems reasonable to investigate the alternative to simple quasi-realism.

6.3. Nuanced quasi-realism

Nuanced quasi-realism-- for which Blackburn (1984, 1993, 1998) is the primary source-- differs from simple quasi-realism in using expressivist, rather than minimalist, analyses to adopt realist claims. Like simple quasi-realism, it can come in hard or soft forms; hard nuanced quasi-realism chooses some realist claims to analyze expressivistically and accept, and others to simply deny, while soft nuanced quasi-realism posits that many or all of the realist's claims admit of both expressive and descriptive readings, and are to be accepted on the former but denied on the latter. My eventual goal in this paper is to defend soft nuanced quasi-realism.

At first, nuanced quasi-realism might seem less promising than simple quasi-realism. This is so because it requires more work-- the nuanced quasi-realist must come up with plausible expressivist understandings of realist meta-ethical claims-- while seeming to suffer from parallel problems. In this section, we will, first, outline how nuanced quasi-realism is meant to work, and then examine the problems that it faces.

Coming up with expressivist analyses of realist meta-ethical claims is more of an art than a science. Simple quasi-realism can draw upon the minimalist formula, but nuanced quasi-realism has no such rigid procedure. Instead, the nuanced quasi-realist proceeds by thinking about what it is that motivates us to defend realist metaethics, and in particular what we want for or fear happening to the practice of moral discourse

and the exercise of morality in general. By means of these considerations, he hopes to generate expressivist analyses of the realist's claims which are plausible by virtue of their connection to concerns that we already recognize as present in meta-ethical theorizing.

We can understand the approach better if we consider some examples. Take, as our first case, that of a nuanced quasi-realist who would like to claim that there exist moral properties. He might find it fruitful to look at the sort of emotional upset that is generated by the denial of moral properties-- that we worry that a morality without properties would not "count" or be good enough, that there would be nothing to make us take it seriously or treat it as mind-independent. The idea is that the meta-ethical claim that such properties actually exist is a reflection of exactly those concerns; it expresses our attitude of disapproval towards un-serious or offensively mind-dependent treatment of moral issues.

Blackburn's (1984) and (1988a [1993]) argument against moral relativism is also a useful example. Blackburn sees the denial of the relativist thesis as an expression of the limits of our tolerance: we do not condone an unlimited variability in the attitudes of others, though their circumstances and individual propensities might vary quite extensively. Rather, there are many cases where we not only disapprove of a type of action, but also disapprove of anyone's approving of such actions, regardless of the other person's other attitudes or whatever social or cultural factors might encourage the approval. Our discomfort with the relativist thesis is a result of these disapprovals,

and our rejection of the thesis is an expression of them. Given that we do feel such discomfort, and that it seems well-placed to motivate us in our denial of relativism, the expressivist analysis of the relativism issue seems to be a good match for the phenomenology of the debate.

Nuanced quasi-realism is subject to problems that parallel those threatening simple quasi-realism. Like simple quasi-realist expressivism, nuanced quasi-realist expressivism must be prepared to defend its distinctness against potential collapses into either straightforward expressivism or realism. Hard nuanced quasi-realism depends on there being enough independence between the various claims of realism that it is coherent to analyze and accept some of them, but not all; if this proves impossible, it will have to either accept or deny the entire set. At the same time, hard nuanced quasi-realism suffers from the same inadequacy as hard simple quasi-realism; it only partially resolves the awkwardness that comes to threaten the expressivist when he denies seemingly plausible realist claims, since it cannot accommodate all such claims without becoming indistinguishable from realism.

Soft nuanced quasi-realism seems to face even worse problems; since it uses expressive analyses, but at the same time acknowledges a non-expressive reading, it needs to justify the relevance of the expressive analysis in answering philosophical questions. For example, if the quasi-realist wishes, as Blackburn does, to claim that the question of mind-dependence should be understood as a moral question, he will have to defend himself against the realist's insistence that there is a perfectly viable

understanding of mind-dependence available which is primarily descriptive, that expressivism is committed to mind-dependence in the descriptive sense, and that being committed to mind-dependence in the descriptive sense is a cost. Given that the realist objector was not intending to make a moral claim, but rather a metaphysical one, presumption is not on the quasi-realist's side. And yet the soft nuanced quasi-realist seems to give up his greatest tool-- the ability to say that there is only one interpretation of the claims about mind-dependence, and if the realist objector insists that he does not mean the claims in that way, then he is simply speaking gibberish-- for dealing with the situation.

For that matter, there seems to be nothing to stop the realist objector from simply *stipulating* that he intends the external reading. Such an objector might state: "On a purely metaphysical level, expressivism is committed to the mind-dependence of moral facts. Yet moral facts are mind-independent. Therefore, expressivism is false." This argument is dangerous for soft nuanced quasi-realists because our faith in the mind-independence of moral facts is not shaken in the slightest by the stipulation that we are talking purely metaphysically. The realist objector does not even need an idea for what the puerly mind-independence thesis would be; he need only specify that he is using whatever external reading the soft nuanced quasi-realist is willing to countenance.

Nuanced quasi-realism, then, does not initially seem to represent much of an improvement over simple quasi-realism. I will argue, however, that we have a way of

saving soft nuanced quasi-realism. If the arguments I present in the following sections are compelling, they will show that the move made by the realist objector in attempting to insist on an external reading is unsuccessful; try as he might, the objector does not succeed in escaping from the moral domain. Once we understand the difficulties in attempting to work with external readings, we will, I think, be able to defend the legitimacy of the quasi-realist's focus on internal, that is, expressive ways of understanding typically realist meta-ethical claims.

6.4. Richards and Stevenson again

As enlightened moderns, we may not feel particularly inclined to accept many of Richards' positions. Nonetheless, there are lessons to be learned from him, some of which are relevant to the topic of quasi-realism. These are:

First, returning to the motivation for the project of multiple definition in Richards (1933), stipulation has its limits²⁴⁹, and if it is weak anywhere, it is weak on

249 Richards writes:

Every word in every philosopher's abstract vocabulary has, I take it, a number of different senses to mislead his readers (and, alas, too often, himself). A remedy for this would seem to be to have more words; and to put "one word, one sense; one sense, one word" before us as a slogan. But, as we all know, the new words tend to take over the ambiguities of the old, and the result is only more words which need still more careful watching. (Richards (1933) p. 31)

Richards worries that we might attempt to disambiguate philosophical terminology through stipulation-- to create new words for each sense, and thereby have one sense per word. He thinks that this approach is likely to be unsuccessful, because the new words, having inherited their meanings (here I use the term in his sense) from particular senses of older words, will become associated with those older words, and that that, in turn, will cause shades of meaning to leach over into the new terms.

We needn't be as extreme as Richards in this regard. To go so far as to ban new terminology seems excessive. Likewise, we needn't insist that whatever theory of meaning we might otherwise favor bequeath the title of "meaning" to any effect on the use or causal effects of a word brought about

emotive meaning²⁵⁰. Attempting to stipulate a term that is like another term but without inheriting its emotive meaning, or with a new emotive meaning, typically fails. As a demonstration, let us now invent a new form of poetry. In this form of poetry, which we will call stipu-poetry, rather than exploiting traditional poetic devices, we make use of new terms, which are stipulated to carry the emotive meaning that would have been created by those devices. In that vein, I present a stipu-poem entitled “Bob.” It goes as follows:

-Begin-

Let “xvf” carry the same descriptive meaning as “sad”, but let it be extremely evocative; let it express extreme sorrow.

Bob was xvf.

-End-

Somehow, I doubt this form of poetry will catch on.

We should note that, per Richards' comments on the matter, the weakness can be either positive or negative. Our example above demonstrates that emotive meaning is difficult to generate.

by this leaching. Yet there is certainly *some* truth to his comment, as we will see from various examples.

250 As per my earlier notes on Richards' terminology, keep in mind that the label of “meaning” is optional, and *not* intended as a claim that such functions satisfy the criteria for meaning endorsed by any theory of meaning save Richards' own. Feel free to substitute “expressive content,” if that phrase is more comfortable. As a reminder, I will do so occasionally.

Richards' argument against the introduction of new terms claims that it is also hard to prevent. Likewise, I see no reason to suppose that stipulating emotive meaning out, of, for example a poem, will go any better than stipulating it in. We can motivate this idea by means of another demonstration.²⁵¹ For reasons of propriety, this demonstration will be stated as a procedure rather than an example. To start with, choose an ethnic slur. Next, imagine a new term, which we construct by taking the ethnic slur and adding an asterisk. Stipulate that this new term is defined as having the same descriptive content as the previous term, but none of the expressive content. Now consider the new term. Such consideration will reveal that the stipulation-*out* of the expressive content has (most likely) failed.

Furthermore, even if we do succeed in removing expressive content through the creation of a new term, it may return in short order. Take, for example, the replacement of "mentally retarded" with "mentally challenged"; it seemingly took no time whatsoever for the new term to become usable as a pejorative, despite the new term having been introduced specifically for the purposes of avoiding pejorative connotations. In fact, it seems that we sometimes struggle to create non-pejorative terms for certain groups, such as those referred to by the above phrases; names for them acquire negative expressive content faster than we can come up with new ones. That we can find ourselves in such a struggle against our own language, and potentially *lose*, says interesting things about our ability to control expressive content, or lack thereof.

²⁵¹ With thanks to Josh Dever for the case.

We needn't take all that much from Richards to learn this lesson from him. His theory could explain why stipulation is weak-- the new term not having participated in the right sorts of contexts, nor possessing the interanimative links of the original, while at the same time possessing its own links in virtue of being its own word-- but we do not need to keep his reasons. What we need is the weakness itself, which, our attention having been called to it by Richards, we can observe on our own.

An interesting consequence of these observations is that if an area of discourse is shot through with emotive meaning, de-emotivising it is no simple task. Introducing new terminology will not do the job. Accordingly, when working in such an area, there may be times when it is difficult to avoid making a statement with emotive meaning. If we *also* take a moral from Richards-- though once again we may disagree with him as to the source if necessary-- that the setting exercises a great deal of control over the meaning (broadly construed) of language, even ordinary language without explicit openings for context, and we combine this idea with the intractability of emotive meaning, we wind up with a picture on which, potentially, there can exist areas of discourse where the control of the speaker over his emotive meaning, at least if he proceeds by purely prosaic techniques, will be quite low. And at the same time, if we follow Richards, we will regard the emotive meaning as worthy of taking seriously, despite its non-descriptiveness, and will do so whether it is inherent or derived, alone or accompanied by other kinds of meaning-- to do otherwise is just to be misled by the various Myths to which Richards calls our attention.

The next moral that seems worth taking from Richards is the pervasiveness of complexity. Of course, once again, we may have to adjust some terminology, and we may disagree about the lower-level mechanisms that explain all of this, but we cannot easily disregard his observations about the subtle influences on the overall effect of bits of language. Very simple language-- what we might call “purely descriptive and univocal”-- is, for Richards, confined to parts of the hard sciences; everywhere else, there exists-- and we *embrace*-- complexity.

The thing to note on this point is that philosophy is part of the *everywhere else*. Of course, if we want to be expressivists, we must have some sympathy for that idea already; after all, to be an expressivist is in part to take a branch of philosophy-- namely ethics-- and conceive of it primarily in terms of emotive meaning. But what of other areas? Stevenson, following Richards, argues that “meaning” is itself rich in (positive) emotive meaning²⁵². Likewise, he argues that “validity” is quite strongly emotive. As we discussed in Chapter 4, he has the insight that, when compared to that which is “valid,” other methods of argument pale; they become “merely”-methods (“merely persuasive,” “merely rhetorical,” and so on)-- and so seem inappropriate for ethics.

We cannot simply assume that there is a non-derivative expressive content to these terms. It may well be that denying the validity of moral arguments, for instance, upsets us not because it expresses a cavalier attitude towards moral reasoning, but

252 Stevenson (1944) p. 42. My reluctance to discard the outdated phrase “emotive meaning” is due to this effect. We shall see shortly why it is that the effect must be respected.

because it is descriptively false. Distinguishing independent expressive content from descriptive content that provokes a reaction is a difficult question, and one that we will not fully answer even when we return to it in section 6 of this chapter. For the time being, however, let us grant the claim-- drawing not only on the insights of Richards and Stevenson but also on the nature of the expressivist project itself-- that expressions of attitudes can not only be the object of philosophical study, but can, in fact, *participate* in philosophical theories.

Put so bluntly, the claim seems surprising, but it is nothing that the expressivist has not been saying all along. After all, the expressivist does not suddenly turn realist when the scope of a moral claim increases. If bits of small-scale moralizing-- such as “using the blender to liquefy your pet hamster is wrong”-- are to be understood as expressions of attitudes (“boo for liquefying your hamster in the blender!”, perhaps), then so too are bits of large-scale moralizing, like “it is wrong to treat others as mere means to one’s own end”. Such large-scale moralizations are the building blocks of first-order ethical theories. Thus, expressions of attitudes are part of the composition of such theories.

Nuanced quasi-realism claims that at least part of meta-ethics is also composed of expressions of attitudes. At the same time, it also adds a complication by granting, as all forms of quasi-realism do, that expressions of attitudes may be true or false. Truth-aptness tends to be a high priority for quasi-realists. It is rare to find an expressivist with so much **boldness** as to deny that moral statements may be legitimately

regarded as true or false; as we saw in Chapter 3, even Ayer, by the mid 1950s, was willing to countenance some employment of truth and falsity within the moral realm. This maneuver will extend to those expressions of attitudes that constitute parts of philosophical theories, including those that participate in meta-ethics. Understanding the soft nuanced quasi-realist answer to the challenge facing it requires that we understand how to treat theories with truth-apt expressions of attitudes as components.

6.5. Dialectical constraints

We are now in a position to establish, making use of the expressivist's inheritance, certain constraints for the dialectic process when expressions of attitudes participate in theory.

Emotive Accuracy: The emotive accuracy constraint is derived from the premise that emotive meaning is common, even in philosophy (which is itself derived from the combination of Stevenson's observations about the likes of “meaning” and “validity”, Richards' observations on the ubiquity of emotive meaning in general, and the expressivist's willingness to see emotive meaning in ethics) and the premise that emotive meaning is important (derived from the elimination of Richards' myths and the quasi-realist's *desire* to, somehow or other, enable emotive statements to be true or false). It claims that, just as a theoretician should not make descriptive claims with

which he disagrees, so too he should not make emotive claims with which he disagrees-- and, thanks to the ubiquity of emotive meaning, this is a real danger in philosophy. If our quasi-realist plans to apply his quasi-realism to the subject of truth, he must be prepared to place himself in a situation where he can say something *false*-- and no less so than in the ordinary case-- if it expresses the wrong attitude.

Accordingly, when theorizing, he must be just as careful about his emotive meaning as his descriptive meaning.

The point here bears repeating, because it is important. The expressivist must accept that expressions of emotion can themselves be part of a philosophical theory, unless he wishes to abolish theorizing in ethics altogether-- which goes against the quasi-realist's goal of modesty. If he attains truth-aptness for such expressions in his pursuit of modesty, then those parts of his theory, like all the others, will be capable of truth or falsity. If he gets those parts of his theory wrong, he is in the same sort of situation-- that of having false claims in his theory-- as if he had gotten any other parts of his theory wrong. The application of quasi-realism to the subject of truth, combined with the fact that an expression of an attitude may be a genuine part of a theory, rather than a mere consequence, creates a parity between descriptive and expressive error.

Emotive Subtlety: The emotive subtlety constraint is the result of the fact that emotive meaning is difficult to control (which we derived above from observation and the

concerns of Richards) and the emotive accuracy constraint. The expressivist must recognize the emotive meaning that his claims will carry in a particular setting and, if he is also following the emotive accuracy constraint, adapt accordingly. This constraint is of particular relevance when emotive meaning and descriptive meaning are both present in a claim. Not only may the expressivist not miss either one, but there may not be an easy and straightforward way of separating the two, or at any rate not one that can be done through definition, stipulation, or perhaps even, more generally, the kinds of techniques that would allow you to convey exactly what you want to convey whenever you want to convey it.

Although we have already talked about this difficulty, it is worth lingering on the matter a while longer to explore the workings of the emotive subtlety constraint. The expressivist must maintain emotive accuracy over expressive content that he cannot easily control. Accordingly, the proper statement of an expressivist view might be indirect, and require some unpacking to understand. I must confess that I have not always given this constraint the respect that it deserves. In particular, when discussing mimicking expressivism, I have made claims that are, in their expressive import, false; my aim was to be more straightforward in my description of the fundamentals of expressivism, but, due to the intractability of expressive content, the fact that I was only interested in the descriptive level of my presentation does not eliminate the expressive. Blackburn, on the other hand, is far more careful, though less descriptively direct as a result; my presentation often diverges from his due to this difference.

Emotive Primacy: An expressivist subject to the prior constraints will often find himself in a difficult situation. He will face questions that seem to have no right answer: where there is a descriptive content with which he disagrees, but an emotive content (or suggestion or implication) with which he agrees, or visa versa. Realists needn't be concerned about such situations; they do not elevate expressions of attitudes to the level of participating in theory, after all, so they are not subject to the emotive accuracy constraint. For the expressivist, on the other hand, it seems that there is a problem.

Resolving such cases is complicated by the other constraints. According to the emotive accuracy constraint, there is no significant difference in the situations of a theory that includes a false descriptive claim and one that includes a false (or unacceptable) emotive claim; the same kind and magnitude of error is committed in both cases. At the same time, the emotive subtlety constraint bars any quick and easy way of disentangling the descriptive and emotive. We cannot make the emotive effect of a term appear or disappear through stipulation. Even creating new terms will not help, because, as Richards writes, “[A]s we all know, the new words tend to take over the ambiguities of the old, and the result is only more words which need still more careful watching.”²⁵³ That there is no effortless way to get exactly the mixture of descriptive meaning and emotive effect that one desires would, no doubt, seem a trivially obvious observation to any student of rhetoric, literature, or poetry; philosophers, too, could benefit from keeping that fact in mind.

253 Richards (1933) p. 31

Owing to the emotive subtlety constraint, a philosopher who seeks to challenge a theory comprised of both descriptive and emotive claims, some of which use the same words as one another, does not possess full ownership of his challenge. What I mean when I say that he does not possess full ownership of the challenge is that he has no way, short of employing artful rhetoric, control over the context, and masterful self-discipline, to control whether he uses, reacts to, or even thinks with the expressive rather than descriptive sense of a claim that permits both expressive and descriptive readings. Enquiry is required to discover whether an intuition that a certain metaethical claim is required, or on the other hand implausible, stems from the descriptive or emotive reading of the claim-- and such enquiry will require deep humanistic thinking owing to the unavailability of the fast and easy methods typically employed by philosophers.

It stands to reason, then, that anything like coming up with a definitive *rule* for identifying when an objector has slipped into thinking in the internal (expressive) sense would be a far more massive undertaking than we have the liberty to attempt here. Instead, in the next section I will aim for the more humble goal of providing some reasons to think that many of the objections that might be posed against expressivism hinge on the expressive rather than descriptive content of meta-ethical claims. Before then, though, we need to work out how the expressivist is to treat mixed cases.

The emotive primacy constraint states that the asymmetry in intractability between descriptive and expressive content requires that we give first priority to the

expressive content in answering meta-ethical questions. That is, when asked to accept or deny a particular meta-ethical claim, where that claim takes both expressive and descriptive readings, the expressivist *must* answer based on the expressive reading. The reason why expressive content must be given priority is that doing so is the only way to achieve clarity. We cannot *easily* consider the descriptive content on its own, thanks to the intractability of expressive content; descriptive content, on the other hand, is amenable to the usual philosophical manipulations. That is not to say that we must endorse quietism in large areas of descriptive meta-ethics. Emotive meaning is stubborn, but not entirely uncontrollable. But unless we can be sure we have undertaken the proper rhetorical and contextual manipulations to engage solely with the descriptive sense, discussion must be constrained to the expressive.

The emotive primacy constraint is useful to the quasi-realist. The soft nuanced quasi-realist can use the emotive primacy constraint to address worries such as those of Cassam (1986), who alleges that the quasi-realist is willfully blind to the purely descriptive sense of the mind-dependence objection. Blackburn's response to Cassam (see footnote #247) is that there is only one "proper" way to understand the question of mind-dependence. The emotive primacy constraint gives us a way of making sense of Blackburn's response: there is, indeed, an external sense of mind-dependence, but that sense is barred from being part of the debate by emotive primacy. By exploiting the constraint of emotive primacy, the soft nuanced quasi-realism can be whole-hearted in his adoption of realist claims in spite of his softness; there is, for him, only *one* answer

to the relevant meta-ethical questions-- only, this answer is subject to supplementation by an external explanation, and revision in certain, peculiarly cold contexts.

The quasi-realist can exploit the emotive primacy constraint in another way as well. When a question is posed with both expressive and descriptive content, the expressivist must answer in accordance with the expressive content. Only if the expressive content is stripped away can the descriptive content be debated. According to *emotive subtlety*, it is difficult to know when this has occurred; the objector cannot bring it about by stipulation. The expressivist may thus make the following argument: by default, the meta-ethical debate proceeds according to the expressive sense (emotive primacy). We do not know *for sure* that we have moved on to the purely descriptive sense so long as any significant expressive effect is present. Yet, any cases where there is no significant expressive effects will be ones in which any intuitions to suggest that the expressivist's (descriptive) position is costly are no longer present. So long as the expressivist's non-expressive description of his position feels wrong, it is possible that we have not really escaped from the expressive discourse forced upon us by emotive primacy, in which case the expressivist may answer by presenting his internal position.

This argument answers the primary concern that we had with soft nuanced quasi-realism: that the objector may stipulate a purely descriptive reading of his objection to the expressivist's descriptive position, and the stipulation does not take the force out of the objection. According to the reply I have just offered, so long as the expression has its force, we cannot know that we are actually arguing in the descriptive

sense, and so it would be premature to conclude that the problem can really survive the shift to purely descriptive language.

6.6 Some signs that meta-ethics is itself expressive

In the previous section, I introduced the constraints of emotive accuracy, emotive subtlety, and emotive primacy. The emotive accuracy constraint applies whenever a theory contains truth-apt expressions of attitudes, the emotive subtlety constraint applies to all expressive content, and the emotive primacy constraint applies to those cases where a claim that is part of a theory possesses both descriptive and truth-apt expressive content *and* the theory wishes to agree with only one of the two meanings. Soft nuanced quasi-realism will want to claim that all three constraints apply to meta-ethics, because many or all realist meta-ethical claims can be understood both as robust, metaphysical claims, and as expressive claims. The soft nuanced strategy is to apply the emotive primacy constraint in order to argue that the meta-ethical debate is actually about the claims in their expressive senses. Were it not, quasi-realism would be irrelevant, since it would not have accepted the realist claims in the senses in which we are interested in preserving them; but if it *is* about the expressive senses, then the expressivist can gain plausibility while continuing to remain distinct from realism.

In order to make the soft nuanced position compelling, the quasi-realist must give us some reasons to think that his expressive analyses are correct. In giving such

reasons, the goal is not to prove that meta-ethical discourse is *necessarily* expressive-- that would lead us into hard quasi-realism, which cannot offer a complete capture of realism while maintaining distinctness. Nor is it to show that *in practice* all meta-ethical discourse up to this point has been expressive, for there are other possibilities, such as the rare properly descriptive case and the case of the theorist who talks nonsense, which are not ruled out. At any rate, either of these two goals would require much stronger reasons than those that are actually available to the quasi-realist. Rather, the quasi-realist's strategy will be to try and give us reasons that suggest key meta-ethical claims have expressive readings, and then lean on the principle of emotive primacy to show that insofar as our discourse in meta-ethics is not hopelessly garbled, it must use those expressive readings.

Finding expressive content within meta-ethics is not a new project. As we have already seen, Richards, Stevenson, and even Ayer had insights that will help the quasi-realist here. For example, consider Stevenson's observation about validity. As a mimicking expressivist, when I claim that moral arguments are amenable to validity, I am applying quasi-realism; what I mean is that the arguments can have forms that I endorse, and think we ought to regard as compelling, in the same way as we do valid descriptive arguments. Now suppose that a realist claims that this is not the kind of validity that we are interested in when we debate metaethics. He might claim, for instance, that our intuitions tell us that moral arguments are potentially valid, and that the "validity" that he is talking about when he makes such a claim is simply a matter of

having one of the logical forms that we recognize as valid for descriptive arguments.

Whom should we trust about the subject of the debate?

Before continuing, I should point out that the quasi-realist needn't, and perhaps shouldn't, fight this particular battle, thanks to the recursive nature of the quasi-realist tactic. Perhaps he might say that moral arguments really *do* have their apparent forms, where this claim is to be understood as a moral claim-- that is, as a claim about how moral arguments are to be treated. According to emotive primacy, the quasi-realist may be compelled to answer in this way; indeed, there may well be dialectical situations where an external claim never becomes available. But let us cheat a little bit for the sake of the example; we will suppose, for the time being, that the dialectical situation allows for sufficient control of expressive content that the quasi-realist may make his external claim without thereby violating the emotive accuracy constraint. Allow me to suggest, though it is understandable if the reasons for this suggestion do not yet seem clear, that if the quasi-realist's stance in the following discussion seems utterly bizarre, it is because our attempt at dialectical cheating has been caught, rather than because of an actual problem with the quasi-realist's position.

The quasi-realist answer to the question with which we ended the paragraph before last comes in two parts. First, he must offer reasons to suppose that there is an expressive sense of "valid", and not simply an emotional reaction to the descriptive content. Second, he will deploy the emotive primacy constraint to block any attempt by the realist to make the debate about anything but the expressive sense. Let us consider

these parts in turn.

With regards to the expressive notion of validity, Stevenson's insight is important, but it is not the only evidence available. We can also observe the use of "valid" amongst the folk, most of whom have no training in logic and thus no inkling of any technical sense of "valid". The folk, if my observations are representative, use "valid" to praise pieces of reasoning, statements used in arguments, and conclusions. We should also make some observations about the intuition to which the realist wishes to appeal. First, the strength of the intuition seems to vary in proportion to the extent to which the mimicking position is phrased in such a way as to seem to take something away from morality. "Moral arguments are not really capable of validity" strikes us as less offensive than "moral arguments are all invalid", which is less offensive than "moral arguments are all just rhetoric." Second, calling a moral argument invalid seems to trouble us without having a particular notion of validity in mind; it even (and perhaps especially) seems to do so on the senses of "validity" used by members of the folk who are wholly unschooled in the terminology of logic. Thirdly, we should be aware of the feeling of the intuition itself. When thinking about meta-ethical issues, our experience of contemplation is somewhat different from in other areas of philosophy. We often feel worried, or troubled-- like a meta-ethical theory is *bad*, not just mistaken. Suggestions of taking something away from ethics disturb us, rather than seeming absurd or silly. Questions about validity are anything but an exception to this trend. All of these observations suggest that our intuition is responding to an expressive rather than a

descriptive claim.

The second half of the quasi-realist argument is to deploy emotive primacy. The realist will want to argue that the quasi-realist's talk of an expressive sense of validity is irrelevant, because we can simply stipulate that we are using the technical, purely descriptive sense. But the constraints suggest that the expressive sense is inescapable; the realist attempts to talk in a purely descriptive sense, but he does not succeed. So long as the expressive sense, which resists the realist's attempt to stipulate it away, is present, it is, according to the emotive primacy constraint, the sense that we are talking about. Were this not the case, we would be unable to control what we are talking about; we would thus be unable to know what to make of the results of the debate, and would be forced to discard them. Only the results of the expressive debate-- in which the quasi-realist's claim accords with our intuition-- can be counted.

This is not to say that we cannot *ever* use the descriptive sense. If we are certain that the expressive content, for all of its stubbornness, has been removed, then we can speak externally. Yet it is very difficult to be sure that we have created an environment where descriptive philosophical debate can proceed without confusion from expressive content so long as we observe emotive effects, which we do so long as the intuition to which the realist appeals remains. Perhaps, when our knowledge of emotive meaning is more advanced than it is now, we will be able to identify other conditions under which non-derivative expressive content has been eradicated, in which case we would have to revisit the case of validity. But, given the evidence to suggest that the intuition tracks

the expressive sense of validity, and the lack of evidence to suggest that this intuition is ever replaced by a different one, one dependent on the descriptive meaning alone, at any stage of the debate, I see no reason to believe that such future research will change our verdict.

The quasi-realist will want to make similar arguments for other realist claims. Similar evidence for independent expressive content can be found for other meta-ethical issues. For example, the quasi-realist may wish to adopt the idea that there are moral properties. This is done by understanding the claim about properties expressivistically, as a claim about how we are to treat moral matters as serious and mind-independent. Note that many of our arguments for the intuition in the case of validity tracking expressive content also apply in the case of properties. For instance, our intuitions about denying the existence moral properties vary in strength based on the extent to which the denial is phrased so as to suggest taking away from the serious and steadfast nature of morality. “There are no moral properties” is offensive, but less offensive than “moral properties aren't real,” which is less offensive than “moral properties are nothing more than illusions created by the form of our language.” We also treat questions of moral properties without first having an account of what a property is-- and it doesn't seem to matter. And the phenomenological argument can apply as well here as in the case of validity.

These cases give the quasi-realist license for optimism. This license can be reinforced with a general thought. It seems that we want realism to be true. Insofar as

we can understand why we desire for realism to be the case, we can bolster quasi-realism; each aspect of realism that we choose to adopt, we can understand expressively as voicing the part of our desire for moral practice that stands to be satisfied by that part of realism being true. I do not know whether any intuitive support for realism will remain once these expressive aspects have been accommodated, but, given the conformity of our meta-ethical intuitions to the patterns examined in the two cases above, I see no reason to believe that there will.

6.7 Future

Quasi-realism is of great help to the mimicking theorist, since he no longer has to worry that he is flying too much in the face of our intuitions. The kind of quasi-realism that I have specified will allow him to accommodate those intuitions, while keeping his view distinct. I take the combination of positions that I have specified--mimicking, pluralistic, and softly nuanced quasi-realist-- to constitute the strongest option available to the expressivist.

We might proceed from this point in several ways. First of all, we need to give more attention to emotive meaning, how to detect it, and how to control it. Left uncontrolled, and combined with expressivism, it presents a large opportunity for confusion; if expressions of attitudes are to be part of our considerations, we will need to be more careful in keeping charge of when we make them.

Secondly, if we are sympathetic at all to Richards, we might wish to re-examine

some of the notions that we use in evaluating theories, in order that we might do so in a way that is not led astray by emotional effects. Let me present a suggestion-- not even a claim, but a posit to consider. The suggestion is that we need to be wary of brutalism masquerading as rigor. By "brutalism", I refer to the architectural style renowned for its simplicity and ugliness. Its counterpart in philosophy is the feeling that the plainer a theory is-- and the more it lends itself to a presentation that is "technical" in the sense of being filled with lots of symbols and other arcana, such that the work obtains a visual appearance and "feel" akin to mathematics and the hard sciences-- the more we understand it, the more specific it is, and the more careful its reasoning. But I see no reason to suppose that these desiderata are fundamentally aesthetic in nature; quite the opposite: surely rigor and clarity are not a matter of *appearing* serious. Likewise, a work that appropriates the aesthetics of mathematics and the sciences does not thereby inherit their rigor or usefulness.

If we are to avoid such problems, we will need to avoid making assumptions as to what is clear, and what form a piece of information has to be transformed into for us to count as understanding it. We will need not only to approach these issues critically, but to first study all of our potential biases in this regard-- that is, to continue the work that Ogden and Richards began-- so that we can keep watch for them, and compensate appropriately, in our reexamination. This would be a significant endeavor, and one with a significant empirical component, but if it opens the door to changes or alternatives in methodology, and thereby helps us to solve ongoing

problems, it could prove worthwhile.

Thirdly and finally, it is worth dwelling for a moment on the idea of expressions of attitudes participating in philosophical theories. As soon as we permit expressivism, we permit expressions of attitudes to participate in our theories; even if we reject expressivism in meta-ethics, there are still other areas, such as aesthetics (or meta-aesthetics) and meta-epistemology, which we have not considered. If we are expressivists in our analysis of any subject within which there exist philosophical theories, then we will allow attitude-expression participation. If we allow those expressions to be truth-apt, and thereby activate the emotive accuracy constraint, then those expressions will have to be taken seriously as part of our theories.

The situation described in the previous paragraph is likely to seem bizarre to us. But before I discuss the potential advantages to allowing it, I would like to question whether it really is as bizarre as it seems. If we have any sympathy for Richards' observations, we might wonder to what extent our aversion to the participation of expressive content in our theories is justified, and to what extent, on the other hand, it is a symptom of scientism. Science, after all, is descriptive, and insofar as philosophy is not, it is not like science. It may be upsetting to learn, after thinking oneself engaged in the lofty enterprise of describing the world, that one is doing nothing of the sort. We must beware the potential for bias created by this effect.

Whether or not we consider theories made up of expressions of emotions to be bizarre, the view does have certain advantages, which we will only sketch rather than

examine fully (at least for the time being). For one, if a theory is emotive in nature, it becomes un-mysterious how intuitions can play the role of evidence for or against that theory.²⁵⁴ We might wonder how intuitions come to be reliable indicators of whatever properties a purely descriptive theory concerns-- which we can understand for intuitions that are formed as a result of empirical experience, but harder to explain in other cases. On the other hand, for a theory comprised of expressive claims, there is no gap between the intuition and the subject matter. For example, we learn about right and wrong, in part, by using thought experiments to reveal our attitudes. Now this is not to say that *what* we learn are facts about our attitudes-- a view where that was the case would be subjectivist rather than expressivist. The fact that we learn about rightness and wrongness through learning about our attitudes does not make those attitudes what we are learning about, because subjectivism is false. As quasi-realist expressivists, we can comfortably say that what we learn about are the properties of rightness and wrongness; quasi-realism allows us to *reconcile* this claim with the claim that there is nothing more, either metaphysically or epistemically, to a moral intuition than a feeling, and nothing more to a moral thought experiment than a story that we hear and react to.

As this work is fast drawing to a close, permit me the luxury of some speculative remarks. In philosophy, we use intuitions in most, if not all, of our sub-fields, and though we may wonder as to their operation, it is difficult to imagine how

²⁵⁴ For a primer on the issue, one might consult Joel Pust's *Stamford Encyclopedia of Philosophy* article on intuitions.

we could do without them. It would, at any rate, be interesting to see what results could be achieved if, drawing upon the considerations of the previous paragraph, and consider that intuitions are uniquely ubiquitous in our field, we defined philosophy as the discipline of crafting theories using expressive language. If nothing else, such a definition would vindicate philosophy's position within the humanities. Where literature allows us to express our emotions in prose, art in image, or poetry in verse, philosophy allows us to do so in theory. Propositional reflection enables theory to work as a useful mode of expression.

This consideration brings us, at long last, back to Ayer (1936 [1952]), who considers the position, which he attributes to C. A. Mace, that metaphysics is to be understood as art:

Among those who recognize that if philosophy is to be accounted a genuine branch of knowledge it must be defined in such a way as to distinguish it from metaphysics, it is fashionable to speak of the metaphysician as a kind of misplaced poet. As his statements have no literal meaning, they are not subject to any criteria of truth or falsehood: but they may still serve to express, or arouse, emotion, and thus be subject to ethical or aesthetic standards. And it is suggested that they may have considerable value, as means of moral inspiration, or even as works of art.²⁵⁵

Ayer objects to this account. He sees the metaphysician's "art" as lacking worth, having been produced by accident by someone intending to describe, rather than an artist who knows what he is doing. The metaphysician, for Ayer, stumbles around in the dark, and while some small beauty may be produced as a result, it will not compare to that which is produced by a master craftsman. If I am right, the combination of

255 Ayer (1936 [1952]) p. 44

quasi-realism and propositional reflection can rescue something like the Mace view. For on my view, there is nothing accidental about what the ethicist, or the philosopher of any other type whose stock in trade is emotive in nature, does. In an emotive field, the descriptive is expressive-- and this is not a matter of consequence, but of identity. The sentences of ethics describe and express, and when they do these things, they do *one* thing, not two. For the question of whether or not a sentence of ethical language describes is itself to be treated expressivistically, as specified by quasi-realism; and, expressivistically speaking, the answer to the question is yes. If all of philosophy is emotive, it needn't be any less descriptive for it, at least so long as we permit the expressive understanding of being descriptive. If we do *not* permit it-- if we reach the rare situation, as specified by our three constraints, where purely external talk is possible; and though I doubt we are in such a situation here, once again, I will cheat a bit for the sake of explanation-- then it will not be so, and on this external reflection, we will note the absence of mysterious properties and mysterious intuitive knowledge. But this external case is, as far as most of our philosophical deliberations are concerned, neither here nor there; it arises only once our concern as to the answer has vanished. As such, it de-mystifies philosophy, but it does not derail it.

If such a state could be achieved for our discipline, it would be a significant accomplishment. Much work, in the direction that I have indicated, will have to be done before this result can be secured. Furthermore, as the present remarks are purely speculative, I cannot guarantee that such a result will be achieved at all. But, at the

very least, the matter seems worth further thought.

Bibliography

- Acton, H.B., "The Expletive Theory of Morals". *Analysis*, Vol. 4, No. 2/3, Dec., 1936, pp. 42-45
- Alston, William P., "Moral Attitudes and Moral Judgments." *Noûs*, Vol. 2, No. 1 (Feb., 1968), pp. 1-23
- Asay, Jamin, "Truthmaking, metaethics, and creeping minimalism." *Philosophical Studies* 163, 2013, pp. 213-232.
- Ayer, A. J., *Language, Truth, and Logic*. New York: Dover, 1936 (1952)
--*Metaphysics and Common Sense*. San Francisco: Freeman, Cooper, & Company, 1967.
--*Philosophical Essays*. London: MacMillan, 1954.
--*Freedom, Morality, and Other Essays*. Oxford: Clarendon Press, 1984.
- Björnsson, G., "Why Emotivists Love Inconsistency." *Philosophical Studies* vol. 104, no. 1, May, 2001, pp. 81-108
- Black, Max "Some Questions About Emotive Meaning," *Philosophical Review*, Vol. 57, no. 2, (March, 1948)
- Blackburn, Simon, *Spreading the Word*. Oxford: Clarendon, 1984 (2004)
--*Essays in Quasi-Realism*. New York: Oxford University Press, 1993
--"Just Causes." *Philosophical Studies* Vol. 61, No. 1/2, February, 1991, pp. 3-17.
Reprinted in *Essays in Quasi-Realism*. New York: Oxford University Press, 1993, pp. 198-209.
--"Practical Tortoise Raising." *Mind* (1995): 695-711; reprinted in Blackburn, S., *Practical Tortoise Raising*. Oxford: Oxford University Press, 2010, pp. 7-25
--"Attitudes and Contents." *Ethics* 98 (3) pp. 501-517, 1988a. Reprinted in *Essays in Quasi-Realism*. New York: Oxford University Press, 1993, pp. 182-197.
--"Moral Realism." In Casey, John (ed) *Morality and Moral Reasoning*. London: Methuen, 1973. Reprinted in *Essays in Quasi-Realism*, New York: OUP, 1993, pp. 111-129.
--Blackburn, Simon, *Ruling Passions*. Oxford: Clarendon Press, 1998
--"How to Be an Ethical Anti-Realist." *Midwest Studies*, Vol. XIII, 1988b. Reprinted in his *Essays in Quasi-Realism*. New York: Oxford University Press, 1993, pp. 166-181.
--"Realism, Quasi or Queasy?" In J. Haldane and C. Wright (eds) *Reality, Representation, and Projection*. New York: Oxford University Press., 1993c
- Boghossian, Paul, "The Status of Content." *The Philosophical Review*, Vol. 99, No. 2, April, 1990.

- Brady, Michael (ed), *New Waves in Metaethics*. Basingstoke: Palgrave Macmillan, 2011
- Brink, David, "Moral Realism and the Sceptical Arguments from Disagreement and Queerness," *Australasian Journal of Philosophy*, Vol. 62., No. 2., June 1984
- Brighouse, M. H., "Blackburn's Projectivism: An Objection." *Philosophical Studies* Vol. 59, No. 2, June 1990.
- Brower, Reuben et al. (eds), *I. A. Richards: Essays in His Honor*. New York: Oxford University Press, 1973
- Budolfson, Mark, "Non-cognitivism and rational inference," *Philosophical Studies* vol. 153, No. 2, March 2011, pp. 243-259
- Burgess-Jackson, Keith, "Rape and Persuasive Definition." *Canadian Journal of Philosophy*, vol. 25, No. 3, September, 1995, pp. 415-454
- Cassam, Quassim, "Necessity and Externality." *Mind* new series Vol. 95, No. 380, October, 1986
- Charlow, Nate, "The problem with the Frege-Geach problem", *Philosophical Studies* Vol. 167, No. 3, 2014
- Chrisman, Matthew, "Expressivism, Inferentialism, and the Theory of Meaning," in Brady (2011) pp. 103-125
 --"Expressivism, Inferentialism, and Saving the Debate." *Philosophy and Phenomenological Research* Vol. 77., No. 2, September 2008, pp. 334-358
- Dall'Agnol, Darlei "Quasi-Realism in Moral Philosophy: An interview with Simon Blackburn." *ethic@* v. 1, n. 2, pp. 101-114, 2002
- Divers, John and Alexander Miller, "Why Expressivists about Value Should Not Love Minimalism about Truth." *Analysis* vol. 54, No. 1, January 1994, pp. 12-19
 --"Platitudes and Attitudes: A Minimalist Conception of Belief." *Analysis* vol. 55, no. 1, January 1995, pp. 37-44.
- Dorr, Cian, "Non-Cognitivism and Wishful Thinking," *Nous* vol. 36, no. 1, 2000, pp. 97-103
- Dreier, James, "Lockean and Logical Truth Conditions," *Analysis* vol. 64, no. 1, Jan. 2004a, pp. 84-91.

- “Negation for Expressivists: A Collectoin of Problems with a Suggestion for their Solution,” *Oxford Studies in Metaethics*, vol. 1 Oxford: Oxford University Press, 2006, pp. 217-233
- “Transforming Expressivism.” *Nous* 33:4, 1999, pp. 558-572.
- “Meta-Ethics and the Problem of Creeping Minimalism.” *Noûs* Vol. 38, Supplement: Philosophical Perspectives, 18, *Ethics* (2004b) pp. 23-44
- “Disagreeing (about) What to Do: Negation and Completeness in Gibbard's Norm-Expressivism.” *Philosophy and Phenomenological Research*, Vol. 72, No. 3, May 2006b, pp. 714-721
- “Expressivist Embeddings and Minimal Truth.” *Philosophical Studies* Vol. 83, No. 1, July 1996, pp. 29-51

Dunaway, Billy, “Minimalist semantics in meta-ethical expressivism.” *Philosophical Studies* 151, 2010, pp. 351-371.

Elstein, Daniel, “Against Sonderholm: Still Committed to Expressivism.” *Proceedings of the Aristotelian Society* (Hardback), Volume 107, Issue 1pt1, pages 111–116, April 2007.

Enoch, David, “How is Moral Disagreement a Problem for Realism?” *The Journal of Ethics*, vol. 13, No. 1, 2009, pp. 15-50

Garner, Richard, “On the Genuine Queerness of Moral Properties and Facts,” *Australasian Journal of Philosophy*, Vol. 68, No. 2, 1990, pp. 137-146

Gibbard, Allan, *Thinking How to Live*. Cambridge, MA: Harvard University Press, 2003

- Geach, P. T., “Imperative and Deontic Logic.” *Analysis*, Vol. 18, No. 3 (Jan., 1958), pp. 49-56
- “Ascriptivism”. *The Philosophical Review*, Vol. 69, No. 2 (Apr., 1960), pp. 221-225
- “Assertion”. *The Philosophical Review*, Vol. 74, No. 4 (Oct., 1965), pp. 449-465.

Gert, Joshua, “Naturalistic Metaethics at Half Price,” in *New Waves in Metaethics*, ed. Brady; Basingstoke: Palgrave Macmillan, 2011, pp. 36-61

- “Problems for Moral Twin Earth Arguments,” *Synthese* 150, 2006, pp. 171-183.

Gowans, Christopher, *Moral Disagreements: Classic and Contemporary Readings*. London: Routledge, 2000

Griffiths, A. Phillips, ed. *A. J. Ayer Memorial Essays*. Cambridge: Cambridge University Press, 1991.

Hale, Bob, "The Compleat Projectivist." *The Philosophical Quarterly*, vol. 36, no. 142, Jan. 1986, pp. 65-84

--"Can there be a Logic of Attitudes?" In J. Haldane and C. Wright (eds) *Reality, Representation, and Projection*. New York: Oxford University Press., 1993.

Hampshire, Stuart, W. G. Maclagan and R. M. Hare "Symposium: Freedom of the Will" *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 25, Freedom, Language, and Reality (1951), pp. 161-216

Harcourt, Edward, "Quasi-Realism and Ethical Appearances." *Mind* Vol. 114, No. 454, April 2005, pp. 249-275

Hare, R.M., *The Language of Morals*. New York: OUP, 1952 (1964)

Harman, Gilbert and J. J. Thomson, *Moral Relativism and Moral Objectivity*. Oxford: Blackwell, 1996

Harrison, Jonathan, "Can Ethics Do Without Propositions?" *Mind* vol. 59, No. 235, July, 1950, pp. 358-371

Holton, Richard, "Minimalism and Truth-Value Gaps." *Philosophical Studies* Vol. 97, 2000, pp. 137-168.

Horgan, Terrence and Mark Timmons, "Troubles on Moral Twin Earth: Moral Queerness Revived." *Synthese* Vol. 92, No. 2, Aug. 1992

--"Expressivism and Contrary-Forming Negation." *Philosophical Issues*, Vol. 19, No. 1, October, 2009 pp. 92-112

--"Cognitivist Expressivism." In T. Horgan and M. Timmons, eds., *Metaethics after Moore*. Oxford: Oxford University Press, 2008, pp. 255-298.

Horwich, Paul, "The Essence of Expressivism." *Analysis* Vol. 54, No. 1, January 1994, pp. 19-20

Jackson, Frank and Phillip Pettit: "A Problem for Expressivism" *Analysis* 58(4), 1998 pp. 239-251

--"Locke, Expressivism, Conditionals." *Analysis* 63(1), 2003, pp. 86-92

Jackson, Frank, Graham Oppy and Michael Smith, "Minimalism and Truth Aptness." *Mind*, new series, Vol. 103, No. 411, July, 1994, pp. 287-302

Kahane, Guy, "Evolutionary Debunking Arguments," *Nous*, Vol. 45, No. 1, 2011, pp. 103-125

- Lenman, James, "Disciplined Syntacticism and Moral Expressivism," *Philosophy and Phenomenological Research*, Vol. 66, No. 1, January, 2003, pp. 32-57
- Mabrito, Robert, "Are expressivists guilty of wishful thinking?" *Philosophical Studies* vol. 165, No. 3, September 2013, pp. 1069-1081
- Macagno, Fabrizio and Douglas Walton, "The Argumentative Structure of Persuasive Definition." *Ethical Theory and Moral Practice*, Vol. 11, No. 5, November, 2008, pp, 525-549
- Mackie, J. L., *Ethics: Inventing Right and Wrong*. New York: Penguin, 1977
- Magee, Bryan, *Men of Ideas*. New York: Viking, 1978.
- Majors, Brad, "Cognitivist Expressivism and the Nature of Belief," *Ethical Theory and Moral Practice*, Vol. 11, No. 3, June, 2008, pp. 279-293
- McGill, V.J., *Emotions and Reason*. Springfield; Charles C. Thomas, 1954
- Merli, David, "Expressivism and the Limits of Moral Disagreement." *The Journal of Ethics* 12, 2008: pp. 25-55.
 --"Return to Moral Twin Earth." *Canadian Journal of Philosophy*, Vol. 32, No. 2, June, 2002: pp. 207-240.
- Miller, Alexander, *An Introduction to Contemporary Metaethics*. Cambridge: Polity, 2003.
- Moore, G. E., "The Nature of Moral Philosophy." In his *Philosophical Studies*: Routledge and Kegan Paul (1922) (hypertext version available at <http://www.ditext.com/moore/nmp.html>)
- Ogden, C.K. and I. A. Richards, *The Meaning of Meaning*. New York: Harcourt, Brace and Co, 1923 (1948).
- Pust, Joel, "Intuition", *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2012/entries/intuition/>
- Richards, I. A. *Principles of Literary Criticism*. New York: Harcourt, Brace 1924
 -- *The Philosophy of Rhetoric*. Oxford: OUP, 1936
 -- *Interpretation in Teaching*. New York: Harcourt, Brace 1938

- *Mencius on the Mind*. London: Kegan Paul, 1932
- *Coleridge on Imagination*. London: Routledge & Paul, 1934 [1962]
- *Practical Criticism: A Study of Literary Judgment*. London, Kegan Paul, 1929[1930]
- "Emotive Meaning Again." *The Philosophical Review*, Vol. 57, No. 2 (Mar., 1948) pp. 145-157
- "Multiple Definition". *Proceedings of the Aristotelian Society*, New Series, Vol. 34, (1933), pp. 31-50

- Ridge, Michael, "Saving the Ethical Appearances." *Mind* new series, Vol. 115, No. 459, July 2006, pp. 633-649.
- "Ecumenical Expressivism: Finessing Frege." *Ethics* Vol 116, No. 2, January, 2006

- Rosen, G. "Critical study of Blackburn's *Essays in Quasi-Realism*." *Nous* 32, 1998, pp. 386-405.

- Ross, W. D., *Foundations of Ethics*. London: OUP, 1939

- Sayre-McCord, Geoffrey, "'Good' on Twin Earth." *Philosophical Issues* Vol. 8, 1997, pp. 267-292

- Schroeder, Mark, *Noncognitivism in Ethics*. New York: Routledge, 2010.
- *Being For*. Oxford: Clarendon, 2008a.
- "Expression for Expressivists," *Philosophy and Phenomenological Research*, Vol. 76, No. 1, January 2008b.
- "How Expressivists Can and Should Solve Their Problem with Negation" *Nous* vol. 42, No. 4, Dec. 2008c, pp. 573-599
- "What is the Frege-Geach Problem?" *Philosophy Compass* 3/4, 2008d, pp. 703-720
- "How Not to Avoid Wishful Thinking," in *New Waves in Metaethics*, ed. Brady; Basingstoke: Palgrave Macmillan, 2011, pp. 126-140

- Schueler, G. F., "Modus Ponens and Moral Realism" *Ethics* vol. 98, no. 3, April, 1988, pp. 492-500

- Schuster, C.A., "Rapprochement in Value Theory." *The Journal of Philosophy*, Vol. 50, No. 2, October, 1953, pp. 653-662

- Schwartz, Jeremy and Christopher Hom, "Why the Negation Problem Is Not a Problem for Expressivism." *Noûs* 00:0, 2014, pp. 1-22

- Sonderholm, Jorn, "Why an Expressivist Should Not Commit to Commitment-Semantics." *Proceedings of the Aristotelian Society* New Series, Vol. 105 (2005), pp. 387-393

- Searle, John, "Meaning and Speech Acts." *Philosophical Review*, Vol. 71, No. 4, Oct. 1962, pp. 423-432
- Shafer-Laudau, Russ, "Evolutionary Debunking, Moral Realism, and Moral Knowledge," *Journal of Ethics & Social Philosophy*, Vol. 7, No. 1, December 2012, pp. 1-37
- Silk, Alex, "How to Be an Ethical Expressivist." *Philosophy and Phenomenological Research*, doi:10.1111/phpr.12138, September, 2014. pp. 1-35
- Sinclair, Neil, "Moral expressivism and sentential negation," *Philosophical Studies* 152 (3), 2011, pp. 385-411
- "Propositional Clothing and Belief." *The Philosophical Quarterly*, Vol. 57, No. 228, July, 2007, pp. 342-362.
- Sinnott-Armstrong, Walter "Expressivism and Embedding". *Philosophy and Phenomenological Research*, vol. 61, no. 3, Nov. 2000.
- Skorupski, John, "The Frege-Geach objection to expressivism: still unanswered." *Analysis* Vol. 72, no. 1, January, 2012, pp. 9-18.
- Smith, Michael and Daniel Stoljar, "Is There a Lockean Argument against Expressivism?" *Analysis* vol. 63, No. 1, Jan., 2003, pp. 76-86
- Stevenson, C. L., *Ethics and Language*. New Haven: Yale University Press, 1944
 – *Facts and Values*. New Haven: Yale University Press, 1963 (1967)
 --"Richards on the Theory of Value," in *I. A. Richards: Essays in His Honor*, Reuben Brower et. al (eds), New York: Oxford University Press, 1973, pp. 119–134
 --"Persuasive Definitions." *Mind* 47, 1938; reprinted in Stevenson (1963[1967]), pp. 32-54
 --"Retrospective Comments" (1963) in Stevenson (1963[1967]), pp. 186-232
- Stoljar, Daniel, "Emotivism and Truth Conditions" *Philosophical Studies* Vol. 70, No. 1, April, 1993, pp. 81-101
- Street, Sharon, "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies*, Vol. 127, No. 1, 2006, pp. 109-166
- Sturgeon, Nicholas, "Moral Disagreement and Moral Relativism." *Social Philosophy and Policy*, Vol. 11, No. 1, 1994 pp. 80-115
- Timmons, Mark, *Morality without Foundations*. New York: Oxford University Press, 1999

- Unwin, Nicholas, "Quasi-Realism, Negation and the Frege-Geach Problem." *Philosophical Quarterly* 49(196), 1999, pp.337-352
 --"Norms and Negation: A Problem for Gibbard's Logic". *The Philosophical Quarterly*, Vol. 51, No. 202, January, 2001, pp. 60-75
- Van Roojen, Mark, "Expressivism and Irrationality." *The Philosophical Review* vol. 105, No. 3, Jul. 1996, pp. 311-335
- Walton, Douglas, "Persuasive Definitions and Public Policy Arguments." *Argumentation & Advocacy*, Vol. 37, No. 3, 2001, 117-132
- Wedgwood, R. *The Nature of Normativity*. New York: OUP, 2007
- Wielenberg, E.J., "On the Evolutionary Debunking of Morality," *Ethics* Vol. 120, No. 3, April 2010, pp. 441-464
- Wilks, Colin, *Emotion, Truth and Meaning: In Defense of Ayer and Stevenson*. Dordrecht: Kluwer, 2002.
- Wright, Crispin, "Moral Values, Projection and Secondary Qualities". *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 62 (1988), pp. 1-26
 --"Précis of *Truth and Objectivity*". *Philosophy and Phenomenological Research*, Vol. 56., No. 4, December 1996, pp. 863-868.
 --*Truth and Objectivity*.
 --"Realism, Anti-Realism, Irrealism, Quasi-Realism." *Midwest Studies in Philosophy*, Vol. 12, No. 1, 1988b, pp. 25-49
 --"Comrades Against Quietism: Reply to Simon Blackburn on Truth and Objectivity." *Mind*, new series, Vol. 107, No. 425, January, 1998, pp. 183-203.
- Zangwill, Nick, "Moral Modus Ponens." *Ratio* Vol. 5, No. 2, December, 1992, pp. 177-193