

INTRODUCTION

Simon Blackburn's quasi-realist project is, at bottom, an attempt to reconcile an apparent division between two philosophical intuitions: naturalism and moral realism. While those intuitions are similar in that each seems to suggest itself in our experiences of the world, the division appears when we analyze the features of those experiences. Suppose you were a witness to a murder, which you took to be immoral, and when asked by the police to describe your experience, you listed its features in the following way: “I saw a woman shoot a gun”, “I heard the gunshot's report”, “I smelled the gun powder”, and so on. But when asked to list the *moral* features of the event, what would you point to, and by what faculty would you say you experienced them? You might also have said that you tried to intercede between shooter and victim before the shots were fired because you had wanted to stop an immoral act before it happened. But if moral properties and natural properties are similar, why should an event's moral features have motivated you in ways that its physical features did not? From these epistemological and metaphysical distinctions, it seems clear enough that natural properties are not so similar to the properties posited by moral realism as our initial inclinations might suggest.ⁱ

As a brand of anti-realism, Simon Blackburn's quasi-realism sides with the naturalists on this question — affirming a physical world of scientifically explicable causes and effects and denying the existence of moral properties that are, in J.L. Mackie's words, “utterly different from anything else in the universe”

(Mackie 38). But as the name suggests, quasi-realism makes a unique move among anti-realists to preserve certain features of moral realism — namely, the authority that moral beliefs are taken to have. Under many realist paradigms, that authority accrues to moral beliefs — and thus to the moral judgments we make with them — because their truth depends upon moral facts about the world and not upon the minds of those who have the beliefs; the truth of our moral beliefs is supposed to be *mind-independent*. Anti-realism threatens moral authority in positing that our moral beliefs are, in fact, expressions of non-cognitive attitudes, thus doing away with the question of moral truth all together.

In attempting to preserve moral authority, the quasi-realist faces two obstacles: the mind-independence of the truth of our moral beliefs and the Embedding Problem. Mind-independence is important to Blackburn because it justifies the authority of moral beliefs, and because it tacitly advocates moral realism, in that it is difficult to imagine that moral beliefs could express non-cognitive attitudes and retain the authority that comes with mind-independence. Thus, if Blackburn could achieve mind-independence (or something like it) for the truth of non-cognitive moral attitudes, he will have simultaneously bolstered the authority of morality under his theory and weakened our inclinations toward moral realism. Blackburn's second obstacle, the Embedding Problem, suggests that there are contexts in which we use moral language where it does not seem like the meaning of that language could consist simply in a projection of non-cognitive attitudes. But since attitudes are not projected in those contexts, they cannot be

projected in any contexts — e.g. between lines in logical proofs. If Blackburn can offer a convincing response to the Embedding Problem, he will have gained access to truth-functional contexts for his non-cognitive attitude-semantics that will allow that semantics to work in many of the same ways that our conventional cognitive belief-semantics are taken to work.

With these interests in mind, Blackburn responds to each concern. For mind-independence, he argues that the conjunction between a requirement for attitudinal consistency and minimalism about truth can achieve mind-independence for the truth of his moral attitudes. His response to the Embedding Problem shows how we *already* express non-cognitive moral attitudes in truth-functional contexts, as well as how we have become so mistaken about the supposedly realist properties of our moral semantics in the first place. Blackburn argues that these efforts do a great deal of work in securing moral authority for his ultimately anti-realist theory — insofar as those efforts succeed, anyway.

It is upon that degree of success that I wish to focus this paper: does Blackburn's theory “earn the right” to continue to use moral language in the traditionally realist ways? I will argue that such success is unlikely, at least in terms of securing moral authority through minimalism about truth and through his response to the Embedding Problem. First, I will attempt to show that Blackburn's response to the Embedding Problem fundamentally misreads the problem's import, and thus misses the mark in terms of what that response *actually* secures for his theory. Second, I will argue that the problem with his use

of minimalism is that it becomes difficult to see how it secures anything even approaching mind-independence for moral attitudes. Toward those ends, I will lay out a more positive version of the quasi-realist project, including Blackburn's minimalism about truth, which will require some discussion of the underlying expressivist motivations. I will then move to a more comprehensive discussion of the Embedding Problem and Blackburn's response to it, setting up the problems with both minimalism and the Embedding Problem that I have already mentioned.

I. ORIGINS OF QUASI-REALISM

The motivations for quasi-realism emerge from earlier forms of expressivism that found it difficult to formulate responses to the Embedding Problem. In order to adequately place quasi-realism in this expressivist pantheon, I will spend a little time discussing the metaethical work of C.L. Stevenson, one of Blackburn's *emotive* predecessors.

In his article "The Emotive Meaning of Ethical Terms", Stevenson offers an expressive analysis of our moral semantics. Using the term "goodness" he arrives at three criteria necessary for the analysis: (1) Upon arriving at what we believe to be a satisfactory definition of "x is good," we must be able to disagree about whether x, in fact, is good; (2) The definition must have magnetism, or the ability to motivate; and (3) Scientific observation *alone* cannot offer a definition. He employs these criteria to evaluate a number of different ethical theories (different brands of subjectivism, scientific theories, and social relativism),

concluding that more clarification of what is meant by “x is good” is necessary.

For that clarification Stevenson distinguishes between descriptive and dynamic uses of language, wherein the former refers to language that *describes* a state of affairs, and the latter *expresses* emotions, attitudes, feelings, etc. The meanings inherent to both dynamic and descriptive uses of language also differ, in that the meaning of descriptive uses of language has truth value, while dynamic uses of language carry meaning that Stevenson calls “emotive.” Emotive meaning refers to the influence that our attitudinal expressions attempt to exert over others. For example, uttering the phrase “Boo Christopher Columbus!” expresses *both* a negative attitude toward Christopher Columbus, as *well* as an attempt to influence others to think similarly.

Stevenson then considers “x is good” in terms of its emotive — rather than descriptive — meaning, substituting the attitudinal expression “We like x — do so as well” for “x is good.” This phrase has a descriptive use, in that it describes my attitude about x, and it has a dynamic use, in that it attempts to motivate others to accept my beliefs. He concludes that such a substitution also fulfills his three criteria for an analysis of “goodness”: (1) *Disagreement*: Different people can have different attitudes about x, and they can disagree about those attitudes (Imagine an Indians' fan sitting next to a Yankees' fan, each cheering his respective team throughout the game); (2) *Magnetism*: Stevenson takes our non-cognitive attitudes to be intrinsically motivational, and “do so as well” attempts to capture the sense in which the phrase “x is good” motivates others to alter their

beliefs to accord with *x*; (3) *Scientific Observation*: physical facts alone cannot have the final say over whether *x* is good, since our attitudes make the ultimate determination.ⁱⁱ

That the substitution “We like *x* — do so as well” fulfills Stevenson's three criteria suggests that emotive translations of “*x* is good” get closer to a definition of goodness than do many other ethical theories. But even so, Stevenson also admits that “We like *x* — do so as well” does not quite capture the “immediate aura of feeling” that surrounds “*x* is good” (Stevenson 77). This apparent incongruity suggests that Stevenson's analysis may not have adequately accounted for some features of what we mean by “goodness,” and objectors seem to have agreed. In particular, they have questioned whether emotivism can adequately account for ethical disagreement, since in our ethical disputes, we seldom take ourselves to be disagreeing about attitudes.

Stevenson does allow for some degree of ethical disagreement — though significantly altered from its non-emotive form. He argues that parties agreeing on an emotive attitude can still disagree about facts and category distinctions corresponding to that attitude. For example, two separate parties who share negative *attitudes* toward murder can disagree on *the facts* as to whether abortion falls under the category of murder. But if those two parties had different attitudes about whether murder was wrong in the first place, they could agree on the facts of whether abortion was murder and have nothing left to argue about, since their disagreement about an attitude toward murder is not open for debate.

And this issue will constitute a major sticking point for Blackburn, for when he reduces what we have conventionally thought of as moral truth to a combination of scientific inquiry and attitude expression, he essentially does away with the notion of moral truth all together.

From this background, it becomes easier to see where Blackburn utilizes and alters the expressivism underlying his quasi-realism. Consider his earliest and most consistent definition, which comes from *Spreading the Word*: “I call the enterprise of showing that...even on anti-realist grounds there is nothing improper, nothing 'diseased' in projected predicates the enterprise of *quasi-realism*. The point is that it tries to earn, on the slender basis, the features of moral language...which tempt people toward realism” (Blackburn 171, *STW*). In a note on this definition, Blackburn further clarifies quasi-realism as a defense of projectivism, which he takes to be “the philosophy of evaluation which says that evaluative properties are projections of our own sentiments (emotions, reactions, attitudes, commendations).” As a defense, quasi-realism becomes:

the enterprise of explaining why our discourse has the shape it does, in particular by way of treating evaluative predicates like others, if projectivism is true. It thus seeks to explain and justify the realistic-seeming nature of our talk of evaluations — the way we think we can be wrong about them, that there is a truth to be found, and so on. (*STW* 180)

Like Stevenson's emotivism, quasi-realism maintains that, despite their realist trappings, our moral semantics lack reference to moral properties, because such properties do not exist for our semantics to refer to; instead, our moral utterances are projections of our own sentiments and complexes of those sentiments onto

the world.

Keep in mind that a naturalistic metaphysics underwrites Blackburn's entire enterprise; he is trying to construct a moral theory that comports with the idea that, to modify Wittgenstein, the physical world is all that is the case. For Blackburn, such a physicalist picture casts human beings as “frail complexes of perishable tissue, and so part of the natural order,” and it refuses “unexplained appeals to mind or spirit, and unexplained appeals to knowledge of a Platonic Order of Forms or Norms,” which is to say that it “refuse[s] any appeal to a supernatural order” (*RP* 48-49).

Toward that end, Blackburn recasts 'the moral problem' as a problem of conflicting intuitions — one toward moral realism and another toward naturalism. He writes that “we nearly all want to be naturalists and we all want a theory of ethics. So the problem is one of finding room for ethics within the disenchanting, non-ethical order which we inhabit, and of which we are a part” (49). From this, it is not difficult to see that Blackburn would take an appeal to a mind-independent notion of moral truth as an appeal to just the sort of “supernatural order” that naturalism scorns. Surprisingly, the route he elects is one that attempts to provide the truth of our moral judgements with the authority presented by mind-independence by combining two claims: (1) that moral judgments express non-cognitive moral attitudes, and (2) that we can still think of those expressions as true and false in a mind-independent way.

But why is Blackburn so intent on securing a mind-independent notion of

moral truth for his theory? Or, put differently, why is Blackburn so intent on avoiding a mind-dependent notion of moral truth for his theory? The problem with mind-dependence is that it might seem to rob his theory of the authority that morality is supposed to have over thoughts and actions, for mind-dependent notions of moral truth very quickly raise the specter of moral subjectivism. As an example, suppose my friend Alex trips an old woman. I am inclined to say that his action was wrong because tripping old women is wrong. However, if the truth of the claim “tripping old women is wrong” depends upon Alex's and my attitudes toward the assertion, we each will appeal for the truth of our moral beliefs to the mind-dependent truth of our moral attitudes. And while he may insist that I should follow his moral attitude and I may insist that he follow mine, neither of us will have a reason beyond the truth of our own attitudes to make the alteration. In such a scenario, the fervor and cunning with which we project our attitudes — not the 'facts of the matter' — will decide the dispute. As we see in this example, the worry is that a mind-dependent notion of moral truth seems to support a kind of moral subjectivism, which would leave our moral disputes in very precarious situations.

As seen with Stevenson, other expressivists avoid this problem by simply claiming that expressions of moral attitudes do not have truth-value at all, which eliminates the problem of mind-independence. However, the problem of lost moral authority persists, and as far as our intuitions about moral realism go, such a loss may undermine inclinations toward the projectivist program. But since

quasi-realism works toward preserving that authority and allowing us to use moral language in truth-functional contexts, this is just the sort of worry that Blackburn's notion of moral truth attempts to defend against.

This is not to say that developing such a notion of moral truth is an easy task, and the difficulty becomes clearer as we consider the tight-rope that Blackburn is walking. As mentioned, Blackburn's theory of truth is attempting to secure mind-independence (contra Subjectivists), but it is also trying to secure bivalence (contra Stevenson) *and* avoid correspondence theories.

Correspondence theories pose a particular problem for the quasi-realist, because, in the first place, they fit very well into many realist paradigms. Furthermore, if an expressivist were also a correspondence theorist, it would seem difficult for her to secure moral authority in any robust sense, because, on the expressivist account, the truth of moral judgments would likely *correspond* with non-cognitive attitudes, the truth of which is supposed to be mind-dependent.

The problem for such a theory arises from the fact that most of us want the truth of our moral judgments to influence the moral judgements of others. Therefore, as the example of Alex tripping the old woman attempts to illustrate, if the truth of those judgements depends upon my mind, there is no reason for my judgements to influence anyone else's judgements. This, roughly, represents the problem of bivalence for the non-cognitivists, in that, on the non-cognitive view, the truth or falsity of moral judgements on will depend on the attitudes being

expressed. The problem of bivalence arises when we have two contradictory attitudes (e.g. Alex's and mine); we would have to say are both simultaneously true, which seems to violate the condition of bivalence.

So, in order to secure bivalence for moral attitudes, Blackburn's theory of truth will either have to allow for moral attitudes to be *treated* as though they were more broadly bivalent—that is, bivalent in a way that does not depend upon the mind of the person with the attitude—when they are, in fact, more narrowly bivalent, or it will have to show how a broader bivalence really does accrue to those attitudes. As David Lewis has argued, the former option implies a fictionalism about moral truth that Lewis takes to underwrite quasi-realism's moral program. Lewis argues that quasi-realists attempt to earn the right to say the things that moral realists say, all the while stipulating those things are false. On Lewis' view, the quasi-realist will advocate for the same claims that the realist will, but will preface the entire set of realist statements with a phrase like, “Let's make-believe moral realism is true, even though we know we know it isn't.” Thus, our morality becomes *quasi*-realist when we combine the actions (and statements) of moral realism with anti-realism's beliefs about the truth of realism (Lewis 318-319).

Of course, since Blackburn does not take himself to be a fictionalist,ⁱⁱⁱ he argues that moral attitudes are in fact bivalent, showing that some moral propositions will have affirmative truth value (i.e. they are *true*), thus skirting Lewis' charge of fictionalism. In order to secure that bivalence while avoiding

correspondence theories, then, Blackburn argues for a minimalist theory of truth.

Blackburn brands his flavor of minimalism “quietism about truth,” but he is ambiguous in trying to pin down a definition. In *Spreading the Word*, he discusses quietism as an attitude toward truth “which urges that at some particular point the debate [over truth] is not a real one, and that we are only offered, for instance, metaphors and images from which we can profit as we please” (*STW* 146). Elsewhere, he treats quietism as “the conservative, pessimistic, and perhaps Wittgensteinian view that we cannot do much with our language except speak it, or at best put down rules for building up meanings by establishing rules that govern the components of sentences with those meanings.” And if such a theory of moral truth should yield “the placid truths that 'good' is satisfied by good things, 'chance' refers to chance, and that *A*, *B* satisfies 'x causes y' when *A* causes *B*, we should rest and be thankful” (Blackburn, “Opinions and Chances” 76).

While perhaps not an actual theory of truth, such quietist notions might be taken to imply a particular attitude that emerges from a minimalist theory of truth, the definition of which Blackburn is more explicit about: “*p* is true if and only if *p*.” Under such a theory, the phrase “is true” adds nothing to our utterances, denying 'truth' the status of “a robust or substantive property” (*RP* 75). So, statements like “It's true that elephants are grey” are reducible to “Elephants are grey,” the truth of which is guaranteed by the fact that elephants are grey.

What, then, is minimalism about truth supposed to gain for quasi-realism?

Since it seems to reduce the truth of a sentence to the meaning of the sentence, or to its surface grammar, minimalism allows the quasi-realist to talk about moral statements that express attitudes as if those statements were bivalent. It shifts the burden of truth in our moral discourse from any kind of correspondence with ethical properties to the surface grammar of the statement itself, thus stripping moral statements of some of their realist implications. Put a different way, while we may balk at the suggestion that our moral statements are bivalent in virtue of their truth or falsity depending upon our attitudes, minimalism overcomes this balking, because it allows the quasi-realist to continue to treat our moral statements as true or false, regardless of the fact that attitudes and sentiments ultimately underwrite them.

However, it is not entirely clear that adopting minimalism solves the problem of bivalence for Blackburn, since it is not clear how a minimalist theory of truth allows him to deal with the problem of contradictory sets of moral attitudes. For instance, suppose that I make the claim, "Working on Sunday is immoral" and you make the claim, "Working on Sunday is not immoral," how is Blackburn's project served by his claim that the truth of the contradictory statements depends upon the surface grammar of each? Certainly, it does not seem that Blackburn's minimalist theory of moral truth provides adequate grounds for adjudicating moral disputes.

Blackburn does offer a rejoinder to this problem, and it comes in the form of what he calls "maximally consistent sets of attitudes," a notion that he

developed in *Spreading the Word*. This notion seems to take as a given — as what we might consider a psychological constraint upon our set of moral attitudes — that we each have an interest in each of the moral attitudes being consistent with the rest of our moral attitudes. Maximally consistent sets of attitudes, then, are sets of attitudes in which no single member contradicts another member, and on Blackburn's view, these are the sets of attitudes that we all strive toward. Blackburn's move to combat the problem of bivalence emerges here: he argues that, on pain of violating the consistency constraint, we could never admit of any attitudes that contradict our own. So, in our dispute about the Sabbath, it may be the case that the truth-values of our statements are determined by their surface grammar (as per minimalism), but unless one of us alters our attitudes, neither of us will be able to admit that the other person's attitude about the Sabbath could, in fact, be true (*StW* 197-202).

II. THE EMBEDDING PROBLEM AND BLACKBURN'S RESPONSE

Minimalism has implications beyond the security of bivalence, and discussion of those implications will involve Embedding Problem, to which much of Blackburn's theory of moral language responds. In short, the problem points out that our moral semantics are not compatible with moral attitudes in the ways that some anti-realists might argue that they are.

That incompatibility poses a problem for earlier expressivist/projectivist theories in general and for Stevenson's expressivism in particular, which has little

to say on the similarity between our ethical language and our descriptive language. Detractors of anti-realism are essentially calling for an error theory hereiv: Why, they ask, does our ethical language *appear* as though it employs properties and facts if those properties and facts do not exist? The objection attempts to attack the projectivist/quasi-realist thesis (that our moral language expresses attitudes) by offering a counter-example in which our moral language seems to express something other than an attitude. The counter-example is known as the Embedding Problem, and it was raised by P.T. Geach, who attributed it to Frege (Geach 223). Blackburn employs this form of the counter-example:

- P1** It is wrong to lie.
- P2** If it is wrong to lie, then getting your little brother to lie is wrong.
- C3** It is wrong to get your little brother to lie. (Blackburn, *STW* 190)

This is problematic for the quasi-realist, because on pain of equivocation, this valid argument must apply definitional identity between instances of the word “wrong” in each line. The problem arises in that the way that “wrong” is used in P2 does not express an attitude; rather, a person making the conditional claim simply states the conditional — she does not express her attitude toward it one way or another. Since the meaning of the phrase must be identical in each line, and since “It is wrong to lie” in P2 does not express an attitude, the instance in

P1 cannot express an attitude. Thus, if this counter-example holds, our moral language does not seem to express attitudes.

Blackburn indicates that this particular challenge poses two questions of the quasi-realist: (1) "What are we up to" when we arrange our attitudes in truth-functional contexts (e.g. conditionals, conjunctions, disjunctions, etc.)? (2) If attitudes are not truth-functional, *why* would we arrange them in accordance with a truth-functional semantics? Regarding the first question, Blackburn argues that we have an interest in conjoining our attitudes in certain ways so as to express an attitude about that particular arrangement of attitudes. For instance, consider the following conjunction: "Baseball is fun, and we play it in the summer." While the second conjunct expresses a cognitive belief about when people play baseball, the first expresses my unique attitude toward the sport generally. What this example attempts to illustrate is that it is possible to use a conventionally truth-functional semantic form (like the conjunction) between non-truth-functional conjuncts. Thus, it is not necessarily the case that couching our moral language in an apparently truth-functional semantics (e.g. "*If* it is wrong to lie, *then* getting your little brother to lie is wrong.") *requires* the semantics to be truth-functional (191). Furthermore, Blackburn's minimalism about truth has secured for him the ability to treat attitude statements as bivalent, so it becomes difficult to see why attitudes would not work in the logical contexts related to the Embedding Problem.

What, then, are we doing when we employ a truth-functional semantics for

what are our purportedly attitudinal moral claims? Blackburn argues that what we are doing is expressing an attitude (or a sensibility) about a particular arrangement of other attitudes. Recall the lying example. The conditional “If telling lies is wrong, then getting your little brother to tell lies is wrong” arranges two attitudes in a particular way. The first attitude is that lying is wrong, and the second is that getting your little brother to lie is wrong. Thus, the conditional expresses a negative attitude toward simultaneously believing that lying is wrong while encouraging your little brother to lie. And while this amounts to a negative attitude toward attitudinal inconsistency, it does not carry the sort of logical requirements extant in truth-functional claims.

Adding other bells and whistles to that kind of attitudinal semantics (combination symbols, division symbols, grouping symbols, etc.), Blackburn arrives at a substitute expressive language that allows him to reformulate the former modus ponens about lying as the following:

B!(lying)

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

B!(getting little brother to lie)

In this formulation, the speaker expresses a positive attitude toward the consistency of endorsing a sensibility that expresses attitudes both against lying and against getting one's little brother to lie. Such an arrangement of attitudes conveys a pragmatic (though not logical) constraint of consistency upon our attitudes, such that we would not approve of an inconsistent set or pairing of

attitudes (193-95).

Furthermore, the expressive semantics bears a striking resemblance to our previous truth-functional property-talk. Blackburn takes this similarity as a reason why we might have made the mistake of confusing actual property-talk (“The grass is green”) with our moral expressions (“Murder is wrong”) in the first place: given the fact that even in an expressive version of the English language it would be difficult to see how use of moral language differed significantly in form from our truth-functional language, it becomes less difficult to see how we could have been mistaken in the semantics we actually use. Thus, with more difficulty in distinguishing between our actual property language and our moral attitude language — and with the help of philosophers whom we trusted and who made mistakes — we became far more susceptible to collapsing the distinction between the two.

III. PRAGMATISM AND MORAL PSYCHOLOGY

Blackburn argues that his response to the Embedding Problem allows him to use an attitudinal moral semantics in the same ways that we have always used a belief-oriented moral semantics. But what does “in the same ways” mean? Certainly, one of those ways will be in the logical, truth-functional contexts that Blackburn's response to the Embedding Problem has secured for the theory. But in ethics, this is hardly the end of the story. So while we may use moral language in those truth-functional contexts, the reason we do so is because we

also want to use moral claims in order to form beliefs based upon the outcome of disputes in those logical contexts.

Consider the following scenario in terms of a belief-based moral semantics: In offering his opinion on the role of morality in decisions regarding the U.S. military, Barack Obama says, “If military leaders propose to me an immoral course of action, I will not support it.” Watching the press conference on television, I form a belief that, if the military should make an immoral proposal to Barack Obama, he will not support it. The next day, I read in the newspaper that Obama called continuing operations at the military prison Guantanamo Bay “immoral’, despite counter-arguments made by military leaders.” Thus, in accordance with my belief from the previous day, about Obama and the morality of military leaders’ proposals, it would be reasonable for me to believe that Obama will not support continued operations at the prison.

Now if quasi-realism succeeds, we ought to be able to use moral attitudes in all of the ways that we have normally used moral beliefs. So, in the above example, the truth of the judgment that the military proposal is immoral will depend, initially, upon a cognitive belief that Obama has about the proposal—about whether the actions contained in it promote unwarranted suffering or are, in fact, cruel. But the question of whether cruelty and undue suffering are *immoral* is going to be based upon attitudes that Obama has about cruelty and undue suffering. Thus, at the most important level of the debate, the truth of the judgment will depend upon a *non-cognitive attitude* that he has about the morality

of actions advocated for in the proposal. And here difficulties begin to arise. For if the judgement of the morality of the proposal depends upon the truth of Obama's attitude about its actions, then the question of whether he supports the proposal — which is *not* an attitude — will be determined by one. So, too, with my own belief about the way that Obama would act: it hinges upon the orientation of his attitudes.

For this particular example, the quasi-realist seems to have a quick response. For one reason that quasi-realism maintains the influence of attitudes on our moral psychology is that, as per Hume, attitudes motivate our action in ways that beliefs do not; thus, if we want our moral psychology to be somehow motivating, our moral theory will have to have elevate attitudes to a special position, much in the way that quasi-realism and other non-cognitive theories do. In the Obama example, then, the quasi-realist might argue that, if we want to know whether Obama's beliefs or attitudes are the more rational grounding for our beliefs about the world, it would be more rational for us to believe that Obama would close the prison if his judgement were based on an attitude, since an attitude-judgement would ultimately motivate him to close it in ways that a judgement based upon a belief may not.

But to push the quasi-realist on the question of the belief/attitude grounding of our judgements, consider this pragmatic question: when I make a moral judgment based upon the advice of another person, *would it be more rational to accept a advice based upon my advisor's beliefs or advice based upon*

his attitudes? Answering this question will require us to distinguish between how we use beliefs and attitudes in our non-moral judgements compared to how we use them in our moral judgements.

Under the Humean picture of moral psychology (of which Blackburn avails himself in a number of other instances) there is a sharp divide between beliefs and attitudes: the truth of beliefs depends upon the world, and the truth of attitudes (if they are bivalent at all) depends upon the mind. And our pragmatic treatment of beliefs and attitudes seems to affirm something similar. Suppose, for instance, that I gamble at the racetrack. Before I make my bet, I ask my friends Ben and Cori how they feel about Secretariat in the seventh race. Ben says, "I would really prefer that Secretariat not win in the seventh," and Cori says, "I think Secretariat is going to win in the seventh, because she runs well late in the day." Since I want my beliefs about the race to align with how the race will turn out, it would be less rational for me to trust Ben's advice, which told me his attitude, than it would be to trust Cori's, which told me her belief.

In the case of moral judgements, if we aimed only to determine the way that the world would be given a certain set of moral considerations, Blackburn's project—as the quasi-realist might argue in the Obama example—and a moral psychology based in attitudes have more appeal than the realist picture of a moral psychology based in beliefs. But, in moral reasoning, predicting how the world will work does not seem to capture fully our intent. Instead, in making moral judgements, we are almost always more interested in how we *should* act,

given a certain set of moral considerations, and not with how we *would* act given those same considerations.

To illustrate, consider two moral advisors, A and B. Now, advisor A's moral psychology is based upon a certain set of moral attitudes and advisor B's is based in beliefs about the world. For example, then, when both advisors say, "Child molestation is wrong" advisor A is expressing a negative attitude toward child molestation, and advisor B means that she has a negative belief about it. Now, suppose that I am seeking advice as to whether I should murder my wife's lover — a question that Blackburn (and presumably Stevenson, as well) would agree is a question of morality and not merely of factual categorization — and I turn to advisors A and B. Now, while A's moral attitudes may serve as a better guide to how A will act than B's beliefs will for how she will act, A's moral attitudes say nothing of how I *should* act. As with Cori and Secretariat, beliefs generally serve as a better guide to the way that the world is than attitudes do, so it ought to be the case that B's beliefs will serve as a better guide to which course of action — to murder or not to murder— is moral and which is immoral. Thus, even if advisor A and advisor B offer me the same advice — that is, if their moral psychologies are exact replicas for one another — it is nonetheless more rational for me to adhere to B's advice than to A's.

In this characterization of attitudes and beliefs, I am still trying to be charitable to Blackburn's more robust notion of attitudes. That is, I do not mean to affirm the rigid distinction between attitudes and beliefs attributed to Hume, for

it seems obvious that there are occasions when our attitudes impact our beliefs and when our beliefs impact our attitudes. And since there are such occasions, I do not mean to say that beliefs map exactly to the world and that attitudes do not map to it at all. The point is that, under normal conditions, we *understand* and *expect* beliefs to map *more* closely to the world than we understand and expect attitudes to do. Thus, when we make judgments *about the world*, we are more apt to rely upon beliefs than we are upon attitudes.

So, recall Obama's statement ("If military leaders propose to me an immoral course of action, I will not support it") and his avowal to close Guantanamo Bay prison because it would be immoral to keep it open as military leaders have proposed. Suppose, as Blackburn argues, that the truth of the judgment that the proposal is immoral is based upon Obama's attitude towards it, rather than his belief about it. Since, in normal circumstances we take attitudes to be less reliable than beliefs, and since in the situation described, Obama's judgment was based upon an *attitude* he had about the military, I would be *less rational* to believe that Obama will close Guantanamo Bay than I would be if the truth of his judgment were based upon a belief.

How, then, is this pragmatic problem about rational belief damaging to quasi-realism? Recall that the distinguishing feature between quasi-realism and anti-realism is that quasi-realism is supposed to "earn the right" to two features of realism: (1) the authority of moral language, because it has earned this (2) the ability to use an attitude-semantics in all of ways we used the belief-semantics.

Blackburn takes his response to the Embedding Problem to secure (2) for him, because it earns the right to use the attitude-semantics in truth-functional contexts, and this gains him access to (1) the authority of moral language.

However, such a conclusion seems to misread the consequences of the Embedding Problem, at least on one level. Certainly, the Embedding Problem does involve using an attitude-semantics in logical contexts, but more fundamentally, it suggests that we do not reason with attitudes and beliefs in the same ways when we reason about morality. And while Blackburn may be able to show that we can make inferences with an attitude-semantics in much the same way that we make them with a belief-semantics, securing the logical contexts he is aiming for is only one part of the issue; it does not yet secure the kind of reasoning that we use in making moral judgments.

The pragmatic problem I argue for attempts to draw this point more explicitly. It works against quasi-realism because it suggests that Blackburn cannot use an attitude-semantics in quite the same ways that he can use a belief-semantics, because, in judgement contexts, attitudes cannot be used in the same ways as beliefs can be used. And since those judgment contexts constitute an important segment of our use of moral language, it will be difficult for quasi-realism to earn the right to moral authority by the means that Blackburn has offered in his response to the Embedding Problem.

Works Consulted

Joyce, Richard. The Myth of Morality. New York: Cambridge UP, 2001.3

Smith, Michael A. The Moral Problem. New York: John Wiley & Sons, Inc., 1994.

Works Cited

Blackburn, Simon. "How to Be an Ethical Anti-Realist." Essays in Quasi-Realism. New York: Oxford UP, 1993.

---. "Opinions and Chances." Essays in Quasi-Realism. New York: Oxford UP, 1993.

---. "Quasi-Realism no Fictionalism." Fictionalism in Metaphysics. Ed. Mark E. Kalderon. New York: Oxford UP, 2005. 314-21.

---. Ruling Passions. New York: Oxford UP, 1998.

---. Spreading the Word. New York: Oxford UP, 1984.

Geach, P.T. "Ascriptivism." The Philosophical Review 69.2 (1960): 223. JSTOR.

Lewis, David K. "Quasi-Realism is Fictionalism." Fictionalism in Metaphysics. Ed. Mark E. Kalderon. New York: Oxford UP, 2005. 314-21.

Mackie, J. L. Ethics: Inventing Right and Wrong. New York: Penguin Books, Ltd, 2004.

McDowell, John. "Values and Secondary Qualities." Mind, Value, and Reality. New York: Harvard UP, 2001. 131-151.

Shafer-Landau, Russ. "Error Theory and the Possibility of Normative Ethics."

Philosophical Issues. 15.1 (2005): 107-120. Academic Search Premier.

Stevenson, Charles L. "The Emotive Meaning of Ethical Terms." Moral Discourse and Practice. Eds. Stephen Darwall, Allan Gibbard, and Peter Railton. New York: Oxford UP, 1997.

-
- ⁱ For a key dissent, see John McDowell's "Values and Secondary Properties" (McDowell 131-151).
- ⁱⁱ Return to baseball. The Yankees fan and the Indians fan have attitudes that their respective team is the "better" of the two. While observation can bear out *some* features of that superiority (winning percentage, team ERA, etc.), there is an emotive aspect of the attitude — say, loyalty — that is not scientifically determinable.
- ⁱⁱⁱ Blackburn's quick response can be found in Question 16 in the "Common Questions" appendix to *Ruling Passions* (RP 319). To the contrary, see David Lewis' "Quasi-Realism is Fictionalism", to which Blackburn responds in "Quasi-Realism no Fictionalism" (both in Kalderon 2005).
- ^{iv} I use "error theory" here more broadly than J.L. Mackie's argument that moral beliefs are categorically false. For a cogent discussion of that project, see Russ Shafer-Landau's "Error Theory and the Possibility of Normative Ethics." (Shafer-Landau 107-120). Regarding quasi-realism, I use the phrase "error theory" in reference to Blackburn's attempt to explain the error in our intuition toward realism in our moral semantics.