

ARE WE PLAYING A MORAL LOTTERY?

MORAL DISAGREEMENT FROM A METASEMANTIC PERSPECTIVE

SINAN DOGRAMACI

The University of Texas at Austin

If someone disagrees with my moral views, or more generally if I'm in a group of n people who all disagree with each other, but I don't have any special evidence or basis for my epistemic superiority, then it's at best a 1-in- n chance that my views are correct. The skeptical threat from disagreement is thus a kind of moral lottery, to adapt a similar metaphor from Sharon Street. Her own genealogical debunking argument, as I discuss, relies on a premise of such disagreement among evolutionary counterparts.

In this paper, I resist the threat from disagreement by showing that, on some of the most influential and most attractive theories of content determination, the premise of moral disagreement cannot serve any skeptical or revisionary purposes. I examine and criticize attempts, made by Gilbert Harman and Sharon Street, to argue from disagreement to relativism by relying on a theory of content determination that involves a principle that, within certain constraints, maximizes the attribution to us of true beliefs. And I examine and criticize Robert Williams's attempt to show there is moral disagreement by relying on a theory of content determination that involves a principle that instead maximizes the attribution to us of rationality. My overall aim is to defend commonsense moral realism via a careful look at the theory of content and concepts.

1. Introduction

1.1. Minimal Moral Realism

It's wrong to torture or kill a pig just for the pleasure it will bring you. It's not wrong-*for-me* but maybe okay-*for-you*. It's simply true that it's simply wrong.

Contact: Sinan Dogramaci <sinan.dogramaci@gmail.com>

And it remains true no matter what any of us says, thinks, or wants. And we all know it.

These and similar claims characterize the minimal kind of *moral realism* I'll defend here. Moral truths are absolute, mind-independent, and easily known. I call it 'minimal' because, as a minimalist about truth, facts, and properties, I see no further issue over whether there are moral truths, facts, or properties.¹ I'll group all alternative views under the label *moral revisionism*. This group includes relativism and error-theory. (I don't mean to include expressivism, a position I take to be consistent with this minimal kind of realism.²)

My aim is to defend realism, but not to positively argue in its support. No argument is needed to support the realist view. It's obvious at the outset. It would take an extraordinarily powerful case to overturn its default plausibility. I'll defend realism against some alleged contenders. My target is a family of arguments for revisionism that all rely on a premise that certain rival moral communities are, or would be, in disagreement with each other.

1.2. *Disagreement and the Moral Lottery*

My target includes the influential evolutionary debunking argument. Though its advocates don't usually advertise it as relying on a premise about disagreement, it's easy to see, if we look, that a premise about disagreement is critical to the debunkers' case.

The debunkers' argument goes like this: the full causal explanation of why we have our moral beliefs will mention that they are evolutionarily adaptive beliefs, but it will not leave it probable that they are true beliefs. Debunkers highlight the process that produced our actual beliefs, and they contrast this with possibilities where we have alternative beliefs. As Sharon Street says:

[C]onsider the question of all those normative judgements that human beings *could* make but don't. . . . [T]he universe of logically possible evaluative judgements is huge, and we must think of all the possible evaluative judgements that we *don't* see—from the judgement that infanticide is laudable, to the judgement that plants are more valuable than human beings, to the judgement that the fact that something is

1. See Horwich (1998).

2. See Gibbard (2003: x, 18–19) on expressivism and minimalism. But see Schroeter and Schroeter (2005) for a dissenting argument that Gibbard's expressivism cannot vindicate realism.

purple is a reason to scream at it. (2006: 133; see also 2016: 305, 314–316; and ms. §11)

The thing I want to point out here is that in order for this observation about the universe of possibilities to serve Street's purposes at all, it must be that along the different evolutionary paths that she invites us to imagine we could have taken, the batch of judgments that we make on any one path must, on the whole, be *inconsistent* with the batch of judgements we make on any other path. Her reasoning is: we are as likely to have ended up on any one of zillions of paths, and on nearly all of these paths we hold false moral views, and therefore it's unlikely our moral views are accurate. For this reasoning to work, our possible counterparts must be disagreeing with each other. Without this premise of disagreement, Street and the debunkers lose the threatening intuition that we participated in a kind of moral lottery. As she uses that metaphor:

[T]he realist is committed not only to a somewhat mysterious epistemology, but also to the much more radical conclusion that—given the apparent odds against having won the normative lottery and the lack of any (internal) reason to think we did—we are in all likelihood hopeless at discovering the normative truth. (2016: 329; see also 313–16)

Disagreement, then, is a crucial premise for the debunkers. It's not usually made explicit, but it's relied on. And it's not argued for. Debunkers just take it as intuitive.³

You might reason that there is disagreement in the *possible* cases because you think there is disagreement in *actual* cases. And, of course, a premise of actual disagreement features in the more traditional argument from disagreement to revisionism. Mackie puts it this way:

In short, the argument from relativity has some force simply because the actual variations in the moral codes are more readily explained by the hypothesis that they reflect ways of life than by the hypothesis that they express perceptions, most of them seriously inadequate and badly distorted, of objective values. (1977: 37)

At the end of that quote, Mackie is claiming that if different groups were expressing attempts to perceive objective values, then "most of them" would be "seriously inadequate and badly distorted", and he is right that this follows from his

3. Bogardus (2016: 655ff.) has argued that disagreement is central to the evolutionary debunking argument.

premise that different groups are genuinely disagreeing and thus at most one of them is getting things right. So, again, the argument for revisionism here rests on the idea that we're in a moral lottery, and it would require some special reasons, which we lack, to think we're not just another loser.⁴

1.3. *How Do I Tell If We're Disagreeing?*

But *are* we all disagreeing? How am I supposed to tell whether other moral communities, or my evolutionary counterparts, are really in a disagreement with me or in disagreements with each other?

Many philosophers do say *it's just intuitive*. Well, what's intuitive to me is that I disagree with my colleagues and my congressman. But the issue here is disagreement between members of radically different communities, even across times and possible worlds. Do I disagree with my infanticidal evolutionary counterparts? It's not obvious to me literally *what* they would be thinking. Is there even disagreement between, say, ancient Egyptians and Aztecs over whether it's "permissible" to remove the heart from a human sacrifice? We're supposed to have an intuition of disagreement in several classic thought experiments.⁵ But many people like me *don't* have a very firm intuition about this stuff, certainly not one that's powerful enough to overturn those (very intuitive!) moral views I started with about hurting a pig just for the pleasure.

We could conduct a survey of other people's opinions. Apparently a lot of folks out there think there isn't a real disagreement between rival moral parties.⁶ But why should I trust other people on this? How do *they* know whether there's disagreement?

The issue is foggy also because there are different kinds of disagreement. My interest is in what's called disagreement *in content*. I believe *p*, you believe *q*, and it's not the case that *p*-and-*q*.⁷ The threat of a lottery requires this kind of disagreement, but there are other kinds of disagreement too. There is also disagreement over what to do,⁸ and there is disagreement over how to use our

4. The style of argument is not just Mackie and Street's. Another example is Velleman (2013: 45).

5. Hare (1952: §9.4); Horgan and Timmons (1990–91).

6. See Sarkissian, Park, Tien, Wright, and Knobe (2011), and Khoo and Knobe (2018).

7. This is so-called "material" inconsistency. Familiar forms of inconsistency are strictly stronger, and so they are included in my target. For example, if *p* and *q* are *logically* inconsistent, then they're also materially inconsistent. I.e., if it's not the case that *p*-and-*q* in all logical possibilities, then it's not the case that *p*-and-*q*.

8. See Stevenson (1944).

words.⁹ So, even if I did sense a disagreement, how can I be sure what kind of disagreement it is?

We need to get back to the basics. We have a problem of radical interpretation. It really is radical because we cannot interpret moral terms by consulting definitions given in more basic non-moral terms—there are no such definitions. (In this paper I'll talk about *morality*, but you can easily re-run the whole discussion for the more generally normative.) So, how do you solve a problem of radical interpretation? We need a theory of content determination, a theory that tells us how intentional facts, such as about what people mean and believe, are determined by the non-intentional facts.

A theory of content determination offers hope of saving us from the moral lottery. To get their disagreement premise and thus the lottery intuition, Street and the debunkers require that our counterparts on each evolutionary path are using the same moral concepts as our own concepts. If I say torture is wrong₁, and what you deny is only that torture is wrong₂—a different concept—then, I take it,¹⁰ there's no inconsistency, no ordinary disagreement in content between us. When Street imagines our counterparts judging that infanticide is laudable, that plants are more valuable, that being purple gives a reason, and so on, she is assuming that these judgements involve our own actual concepts of laudability and so on. But in saying that, Street took a stand on non-trivial issues in the theory of content. Evolution, and all the rest of the non-intentional facts, don't just hand us one batch of beliefs out of a big bucket mostly full of false beliefs. The correct theory of content determination will have the non-intentional facts assign us certain beliefs, and in doing so they also assign us our moral concepts. And depending on what the theory tells us, it may then turn out to have been all but guaranteed that our history and our environment hand each of us moral concepts that correctly apply to just the things we typically apply those concepts to. Far from making us play a lottery, our contingent predicament might rather have served as the insurance that guarantees we're reliable.

The rest of this paper examines two approaches to the theory of content determination, both of which I find appealing. (Others authors have examined other theories of content determination to see what they imply about moral disagreement, and also reached the conclusion that there's no clear disagreement between rival moral communities.¹¹)

9. See Plunkett and Sundell (2013) and Khoo and Knobe (2018: §3 & §4).

10. I'll be relying on the mild assumptions that concepts exist, and that agreement or disagreement in content requires shared concepts. I believe my main interlocutors (Harman, Street, Williams) would not object to my arguments by objecting to this assumption, but some philosophers (e.g., ones who think all representation and all content is coarse and unstructured) may independently deny these assumptions.

11. Dowell (2016) considers Millikan's teleological approach. Schroeter and Schroeter (2013: 2019) consider an "Ideal Accessibility" approach drawn from Jackson and Chalmers.

The first approach I'll look at is one I'll call *Truth-Max*, which says that, within certain constraints, a person holds maximally true beliefs. I'm interested in *Truth-Max* not only because it's an influential and plausible view, but because Gilbert Harman and Sharon Street are led by it to a moral view I find very revisionary: moral relativism. My goal will be to argue that *Truth-Max* does *not* lead to relativism but instead has us stick with absolutist, mind-independent realism.

The second approach I'll look at is one I'll call *Rationality-Max*, which says that, within certain constraints, a person holds maximally rational beliefs. I'm interested in this again not only because it's another influential and attractive view, but also because Robbie Williams (2018) recently argued that this approach in particular leads to the view that rival moral parties do disagree. I'll argue, against Williams, that *Rationality-Max* does *not* imply that rival moral parties disagree.

Historically, both of these two approaches were influentially discussed under the name of "The Principle of Charity".¹² I've re-named them here to be clearer about which version of that famous but murky principle I'm talking about in each context.

2. Truth-Max

2.1. What Is Truth-Max?

The intentional facts are determined by (i.e., metaphysically necessitated by) the non-intentional facts. As I'm interested in it, *Truth-Max* is not a full theory of content determination. *Truth-Max* tells us just one part of how the overall determination goes. (That's why I call it an "approach" to the theory of content determination.) It says a creature has maximally true beliefs, but of course that claim can't be the *whole* story: nobody believes every truth and zero falsehoods. So, I'll work with just the following:

Truth-Max: within certain constraints (specified elsewhere in the theory of content determination), the correct assignment of contents to a creature's beliefs is the assignment that maximizes attribution of true contents.

12. The name "Principle of Charity" was originally introduced in Wilson (1959), who defined it this way: "We select as designatum that individual which will make the largest possible number of Charles's statements true" (1959: 532). Charity, on this *Truth-Max* reading, was then popularized by Quine and Davidson. See Quine (1960) and Davidson (1984), both of which have index entries for "principle of charity". Davidson comments on the importance of the principle in Quine's views. The adjustment from truth to rationality was most clearly featured and popularized by Lewis (1974), though he still used the name "Charity".

The idea here is that there is significant “pressure” on creatures to end up with true beliefs. I don’t mean there’s pressure on you or me, as onlookers, to think something about some creature we’re trying to interpret (though that may also be the case). I mean it’s a feature of reality itself, of the way reality puts creatures in representational states, that there is a bias or pressure toward giving them true beliefs.

We’re also going to be interested in the assignment of *concepts*, since agreement and disagreement in content require shared concepts.¹³ But an assignment of contents will thereby assign a thinker concepts. You have whatever concepts are needed in order to believe (or otherwise represent) the contents you believe. For example, if you believe that water is identical to H₂O, then you need, and you have, the concepts of water, of H₂O, and of the relation of identity.

Many big questions aren’t answered by *Truth-Max* as I’ve stated it. What determines that a creature has mental states, in particular belief states, in the first place? What determines the contents of other intentional mental states like desires or perceptual experiences, or of assertions? What about the grounds of content attributions, the facts in virtue of which they hold, as opposed to mere determinants? How exactly does *Truth-Max* determine which concepts we possess, if different permutations of different concepts could equally well satisfy it? I’ve left all those questions unsettled because they won’t matter for my argument against Harman and Street below. Although the indented statement of *Truth-Max* above may look like it doesn’t tell us much at all, it’s enough for our purpose here, which is to examine whether *Truth-Max* can help make a case for moral revisionism.

Advocates of the *Truth-Max* approach hold very different views about what these “certain constraints” are that *Truth-Max* says our beliefs are maximally true within. The classic Quinean idea is that a person’s beliefs *about certain topics* will be maximally true, in particular beliefs about their own immediately observable environment. Richard Boyd, in his seminal application of metasemantics to metaethics, says we’ll tend to get true beliefs about the things we’re causally connected with. He famously put it this way:

Roughly, and for nondegenerate cases, a term *t* refers to a kind (property, relation, etc.) *k* just in case there exist causal mechanisms whose tendency is to bring it about, over time, that what is predicated of the term *t* will be approximately true of *k* . . . (1988: 195)

Boyd’s statement shows, among other things, how *Truth-Max* is not vulnerable to the objection that it rules out intuitive possibilities of massive error. Can’t

13. See Footnote 10 above

most or even all of a person's beliefs, including their beliefs about their local environment, be wrong? Yes, but, if *Truth-Max* is right, not all the time. You can be unwittingly kidnapped overnight and turned into a brain in a vat, and then have mostly false beliefs about your new environment, but you'll eventually once again have true beliefs about this new environment, and in the meanwhile you'll have true memories about your old one. These are not implausible restrictions on how much error is possible.

More recently Timothy Williamson (2007: ch. 8) gives a new and appealing suggestion for a way of charitably maximizing truth attribution. He proposes a person has those attitudes that would (within certain constraints, as always) maximize the attribution to them of *knowledge*. (His chapter is titled "Knowledge Maximization".) Since knowledge implies truth, this leads to a version of *Truth-Max* (and also to a version of the approach I discuss later in the paper, *Rationality-Max*, since knowledge implies rational belief). Since our empirical knowledge seems to require causal connections to the environment, Williamson's version could remain close to the Quinean idea about what topics we'll tend to have maximally true beliefs on, though Williamson argues his proposal is more defensible than traditional belief-centered versions.¹⁴

I'll remain neutral between a Quinean, Williamsonian or other version of *Truth-Max*. And I'm not specifically interested in our beliefs about the local environment that we're causally in touch with, at least not as such. My interest is in our ordinary moral beliefs, which I don't assume are about causally efficacious natural properties, as Boyd thinks—I'm officially neutral on that. I'm specifically interested in what *Truth-Max* might imply or suggest about our moral beliefs, and in particular whether *Truth-Max* could support moral revisionism. Harman and Street make an argument for moral relativism that, we'll see, turns on an appeal to *Truth-Max* as applied to our moral beliefs in particular. How exactly the constraints around *Truth-Max* allow it to apply to our moral beliefs (to make them in particular come out true) will not matter to my assessment of Harman and Street's case for relativism.

While I'm officially neutral on whether moral properties are, as Boyd thinks, natural properties, I'll briefly mention now that the best motivation in support of *Truth-Max* is one that may favor Boyd's naturalistic approach. The best motivation for *Truth-Max*, at least in my view, is that it features in an excellent explanation of why humans use the intentional concepts they do, in particular why we use our concept of belief. One function of our practice of attributing beliefs to each other is that it allows us to make use of testimony. If our practice of belief attribution is constrained by *Truth-Max*, then the practice will be one that we tend to apply to true beliefs, particularly in whatever

14. See McGlynn (2012) for a critical assessment of Williamson's proposal.

range *Truth-Max* applies to, and that would help to make sure that testimony is useful. Williamson's knowledge-maximizing version of *Truth-Max* would be especially well motivated by this fact about the function of intentional concepts. It could be nicely packaged with Edward Craig's view that it's the special function of knowledge attributions to help us identify and publicly flag trustworthy testifiers.¹⁵ However, for many people, this testimony-centered motivation for *Truth-Max* will sit awkwardly with the application of *Truth-Max* to our moral beliefs, because, as many people think,¹⁶ moral testimony is not trustworthy in quite the same way testimony about ordinary empirical matters is trustworthy. But this motivation will seem more reasonable to people like Boyd, who assimilate our moral beliefs to our empirical beliefs about the natural world, and who thus also presumably won't see moral testimony as especially puzzling.

In any case, my interest isn't in one or another particular motivation for *Truth-Max* or how to fill out the many details it leaves unspecified. My interest is in Harman and Street's inference from it to relativism. It's interesting to point out here, though, that Harman once said, "I continue to think of myself as a moral relativist. What I think I have learned from Sturgeon and Boyd is that Moral Relativism is best treated as a version of Cornell Moral Realism!" (2019: 2)

2.2. Harman and Street: *Truth-Max* Supports Relativism

It's widely thought that an approach like *Truth-Max*, and in particular Boyd's view, implies that rival moral communities do *not* disagree. This is the main point of Horgan and Timmons (1990–91), who consider this a cost (because they claim that's counter-intuitive), but I'm viewing it as a benefit: it helps us escape the moral lottery.

The question I'm interested in now, though, is this. *If* we endorse *Truth-Max*, and we apply it to our moral beliefs, and we gladly accept that rival moral communities are not disagreeing (in content), *then* should we be relativists or can, and should, we remain absolutists, like on my favored form of realism? Does pressure to assign true moral beliefs to rival communities lead to pressure to go relativist?

Gilbert Harman and Sharon Street endorse relativism, and they do so on the basis of an application of *Truth-Max* (as I'll explain). (Street also calls herself a "constructivist", but in (2019) she clarifies that her view is a form of relativism.)

15. See Craig (1990).

16. See Hills (2009) and McGrath (2011).

I'll now present their view and afterward I'll argue that their appeal to *Truth-Max* would better support absolutist realism.¹⁷

What is relativism?

For simplicity, I'll just talk in terms of (moral) reasons and attributions of reasons. (It'll be easy to generalize the discussion to good, right, permitted, etc.) We can then state Harman and Street's favored understanding of relativism as this conjunction:

Relativism: (i) people once generally believed that the truth-conditions of our attributions of reasons required an n -place relation to hold, but (ii) in fact, the truth-conditions of our attributions of reasons require an $(n+1)$ -place relation to hold, and (iii) only the latter relation actually holds.

For example, perhaps we thought our attributions required a 3-place relation to hold, x is a reason for y to z , but now, it's claimed, we see that we require, and there does indeed hold, a 4-place relation, x is a reason for y to z given m . That last relatum, m , will be what Harman calls a moral "framework" or sometimes "coordinates" or "code", the agent's (y 's) moral framework/coordinates/code on the agent-relative form of relativism that Harman and Street both prefer. Street takes m to be the agent's (global set of) "evaluative [or normative] attitudes".¹⁸

Conjunct (i) only concerns what we believed, not what we meant, because Harman and Street both say relativism isn't built into the semantics of our word, or built into our concept of a reason.¹⁹ But they both explicitly say, as (ii) indicates, that we must now revise our understanding of the *truth-conditions* of our attributions (even though our words and concepts have not changed). For Harman, the revision occurs after the discovery of moral diversity in the world; for Street, it's reflection on the contingency of our causal history, exemplified by Darwinian evolution, that prompts the revision.²⁰ The last conjunct, (iii), reflects their view that these relativist truth-conditions are met, as opposed to error-theory.²¹

How then do Harman and Street argue for relativism?

17. Harman wrote about relativism in many places over a long period of time, and his views evolved. (The quotation I gave earlier from Harman (2019) is an example of him even announcing a change in view.) Below (as my citations will indicate) I will draw primarily from his lengthiest work, which was his 1996 book (a debate with Judith Thomson) in order to interpret what Harman takes relativism to be and what his main argument for it is.

18. Harman (1996: §§5.3–5.4); Street (2006: §10). Harman also considers, tentatively, a "quasi-absolutist" moral language. This is a use of language relativists might construct in which moral evaluations express the speaker's own attitudes of approval and disapproval (see 1996: 34ff., 49). See Judith Thomson's reply (1996: esp. 206–8) for a skeptical assessment of the quasi-absolutist proposal.

19. Harman (1996: 4); Street (2019: secs. 6–7).

20. Harman (1996: §§1.2–1.3); Street (2006: esp. §10; 2016: esp. §12.11).

21. Harman (1996: §1.1.1); Street (2016: esp. §12.1).

Neither explicitly invokes the name ‘principle of charity’ or any similar name when they argue for relativism. And neither explicitly says they’re resting their arguments on a premise about content determination. But both of them make an argument that calls for a revisionary interpretation of the truth-conditional content of our attributions, and both are advocating for that re-interpretation on the grounds that we need it in order to make our judgments come out true. The form of argument is: we should, or must, interpret people’s moral judgements as generally true, and this requires that the truth-conditions involve the kind of $(n+1)$ -place relation characteristic of relativism. (Boghossian also says an appeal to the “principle of charity”, is “Harman’s main argument for his view”.²²) Let’s look a bit at what each of them say.

Harman argues by first considering relativism about mass and simultaneity on Einstein’s theory of relativity, and then extending that case to morality. An appeal to *Truth-Max* occurs in his discussion of Einstein when he writes “[I]t would be mean-spirited to invoke an error theory and conclude that these pre-Einsteinian judgments [about mass or simultaneity] were all false!” (1996: 4). The idea is that Einstein showed us that only a, let’s say, $(m+1)$ -place relation holds, but by *Truth-Max* we want to say pre-Einsteinian people made true judgements, so Harman infers that their truth-conditions, and ours too, involve this $(m+1)$ -place relation.²³

But, most reasonable people would agree, there is no Einstein for morality. So why does Harman think only an $(n+1)$ -place relation holds in the world when it comes to moral reasons?

Harman’s case for moral relativism is an inference to the best explanation based on the fact of moral diversity: “moral relativism is (part of) the most plausible explanation of the range of moral diversity that actually exists” (1991: 13; also see 1996: §§1.2–1.3). What is the *data* Harman aims to explain here? The data is the important fact that “moral disagreements that rest on basic differences in moral outlook cannot be rationally resolved” (1996: 12). (Harman does not mean that these are cases of disagreement in content; he will propose to explain this datum by interpreting both rival parties as making true judgments, so the judged contents must be compatible in the end.) What’s the *explanation*? He first tells us what the absolutist could say: their “explanation might be that some people are not well placed to discover the right answers to moral questions” (1996: 12). This amounts to the suggestion that the absolutist is forced to say that there are many losers in the moral lottery. What’s the relativist’s competing explanation? What

22. Boghossian (2006: 17); see also (2011: 54–55).

23. Boghossian (2006; 2011) objects: Harman is being charitable about people’s judgements about mass and simultaneity, but uncharitable about their judgements about the truth-conditions of their thoughts and assertions. Perhaps Harman could reply to this objection somehow. I’ll set it aside.

Harman says at this critical point in his presentation of the argument is just a statement of moral relativism. He states relativism about motion and then he says

A relativistic answer is also plausible in the moral case. Moral right and wrong are relative matters. A given act can be right with respect to one system of moral coordinates and wrong with respect to another system of moral coordinates. And nothing is absolutely right or wrong, apart from any system of moral coordinates. (1996: 13; and see §1.3)

What's the reasoning here? Harman views the best explanation of the situation as one on which we're all making true moral judgments. Why exactly does he think that explanation is the best? Here, it seems to me, he is just relying on *Truth-Max*, and not on any further argument. He simply thinks it's less plausible to impute widespread errors to people, and more plausible to be charitable and follow *Truth-Max* by imputing true beliefs—at least in the absence of other over-riding considerations. And I think we should, defeasibly, agree this use of *Truth-Max* does look fairly reasonable.

Street's argument for relativism also includes an appeal to *Truth-Max*. If absolutism is true then we're in the moral lottery and our overall perspective undermines itself: our ending up with true moral beliefs would be a coincidence that we can't rationally believe in. But we must think we do have true moral views, so we're *not* in a lottery, so absolutism is false.²⁴ But what's her argument that our moral views are true? Street, like Harman, doesn't explicitly cite *Truth-Max* or any metasemantic views, but she's clear that, in her view, we *need* to give an interpretation of the truth-conditions of our judgements that has them come out true—a revisionary interpretation if necessary. Street thinks we must charitably interpret our normative beliefs in particular. This is because normative assessment is “inescapable” in our lives:

The practical standpoint is inescapable because we cannot avoid asking and answering questions about our normative reasons. Even if our thinking about such matters is often only implicit and superficial, we cannot avoid considering what to do later today, or tomorrow morning when we wake up. (2016: 294)

In effect, when it comes to the interpretation of ourselves, *Truth-Max* rises to the level of undeniability: we must have normative views, and we have to think our

24. See Street (2006; 2016). There's a short statement in (2016: 305).

own views are fairly reliable, and so we cannot attribute (truth-conditional) contents to our own views that make them come out too unreliable.

So, in short, Street's argument goes like this: we need a view of the contents of our normative beliefs that, when combined with the contingency of evolution, isn't self-defeating—and *Truth-Max* appears in that "isn't self-defeating" part—and that means we must endorse relativism.

2.3. My Proposal: Conceptual Diversity

Actual and possible moral diversity do pose a problem. If we're all disagreeing, and nothing makes us special, then we're in a moral lottery. Most groups, likely including ourselves, will hold overall false views. Harman and Street aim to avoid this outcome. We can avoid it by obeying *Truth-Max*: interpret the truth-conditional contents of what rival moral groups believe in a way that makes everyone reliable. In particular, rivals aren't disagreeing. Relativism *can* get that for us.

But the inference to relativism is too quick. Harman and Street have made an argument by elimination. They argued for relativism by disputing the whole conjunction of realism and the thesis of disagreement.²⁵ There's an option they neglected.

Undercutting the threat of the moral lottery by interpreting everyone in the lottery as a winner is a great idea. The idea that we're likely a loser, getting it all wrong morally, is too implausible to believe—we *know* it's wrong to hurt the pig just for our pleasure. But relativism is radically revisionary, so if it could also be avoided, it should be. Can we do both, undercut the lottery and avoid relativism?

We can. Consider again the relativist's move of turning to an $(n+1)$ -place relation as the truth-conditional content of our moral language for attributing reasons.²⁶ So, the relativist says that we use a moral predicate ("is a reason" or something like that) to express a relation that consists of $(n+1)$ -length tuples, $\langle x_1, x_2, \dots, x_n, m \rangle$, with the newly discovered relatum in the last position being an agent's "moral framework" or something like that. What I propose is to move the role of the moral framework out of the content of our moral concept and

25. Others have similarly said only relativists, not absolutists, can remove disagreement in content from otherwise disagreeing (in plan, or only apparently) rival moral parties, e.g., Wedgwood (2019: 97, 106).

26. I am using "relation" for the truth-conditional content of a predicate. I mean the same thing as what's also standardly called an intension, which is just an extension for each possible world, which is just a set (of ordered tuples) for each possible world.

relocate it in the conditions that determine the content of an n -place moral concept shared by a moral community. My proposed view is this:

Conceptual Diversity (C-D): **(i)** any moral community m that coherently uses a term of moral praise and blame expresses a moral concept, C_m , whose truth-conditional content is a relation, R_m , but **(ii)** this concept and this content do *not* contain the moral code or framework of m as a relatum, rather **(iii)** the content is an n -place relation, $\langle x_1, x_2, \dots, x_n \rangle_m$ and **(iv)** for each community, m , the n -place relation that community expresses (marked here by the sub-script index m) does hold of the things in the world that the community typically applies their predicate to.

First a few comments to clarify and explain this view, (C-D). Then I'll argue that, if *Truth-Max* is our guide for assigning content to moral beliefs, then (C-D) is preferable to relativism.

Why does (C-D) assign concepts to whole communities, as opposed, say, to individual agents or to contexts of utterance? The answer is that semantic deference seems to occur between individuals and their own communities. In the case of moral concepts, this is manifested in (what I find to be) the intuitive fact that individuals in a shared community easily disagree, and specifically in such a way that one of them must be mistaken. A stubborn anti-vegan or a racist congressman doesn't so easily get to use a concept that makes him right. I do, however, take the actual world and its history to contain a wide range of moral communities that use diverse moral concepts, though the communities' boundaries are hard to specify. (C-D) is analogous to a popular kind of moral contextualism, but moral contextualists tend to treat *utterances* as shifting truth-conditional contents, where (C-D) treats shifts in content as sensitive only to whole communities over a long stretch of time. Two notable recent defenses of contextualism are Björnsson and Finlay (2010) and Brogaard (2008). However, their characterizations of contextualism do not feature the critical logical shift that I'm emphasizing in (C-D). They seem to leave the shifting relatum m inside the content of the proposition, rather than moving it out in the way I do.²⁷ Moving m out of the content will be important to my argument in favor of (C-D) in the next section.

What makes a term, or concept, one of *moral* evaluation? I leave this question to others.²⁸ (My inclination is to think there's some coordinative function in

27. See Björnsson & Finlay (2010: 7, 25), and Brogaard (2008: 400).

28. See Wong (2006) for one rich exploration of how human societies need and use distinctively moral concepts.

social animals that moral language serves, a function that different words and concepts could equally well serve.)

What guarantees, as conjunct (iv) says, that the relation obtains in the world? Aren't moral properties and relations queer entities that we should be skeptical about? No, not for minimal moral realism. On the minimalist view, properties in general are cheap. For any set of possible objects, including any set of possible actions, there's some property that they, and they alone, have. (Similarly, for any set of possible n -tuples, there is a relation.) And there is also a potential predicate expressing that property (or relation), and true of just those possible actions. So, the set of possible actions that utilitarianism deems right, or the set a Kantian theory deems right, or any other set of possible actions, is a suitable candidate for the property picked out by a coherently deployed predicate of moral praise or blame.²⁹ Perhaps the requirement of coherent usage will rule out extremely gerrymandered sets that cannot be promoted through any coherent practice of moral praise, but I don't know what such a case would be. Like Street suggests, I believe we could have used *moral* predicates and concepts—though I deny they would be our own actual concepts—to laud infanticide, to value plants over humans, or to praise screaming at purple.

The basic idea behind (C-D) is, I hope, simple. Different moral communities use a diversity of absolutist moral terms and concepts that they correctly apply to different things.

Two analogies may further help convey the view and, hopefully, its initial appeal. Will the complete causal explanation, including evolution, of why a community ended up with the beliefs they have about their particular local foods make it likely those beliefs are true? Yes, of course. And it's obvious that there will be a tremendous amount of diversity in the concepts used by different communities: each community will develop a distinct set of concepts for their own local plants and meats, and they will not use concepts for ones that are unavailable locally. It's true of actual communities as well. Aztecs said true things about avocados and they didn't talk about olives at all, and vice versa for ancient Egyptians. No relativism and no surprising ($f+1$)-place relations are called for here though.

The other analogy is color concepts. This analogy is suggested by van Roojen (2017) in a paper presenting (though not fully endorsing or positively arguing for) an idea similar to (C-D). He points out that on certain theories of color, in particular the theory of so-called "anthropic realism" developed by Hilbert (1987), colors are mind-independent properties, but it is highly contingent which of the many possible color concepts we evolve to use. Different evolutionary

29. Schroeter and Schroeter (2019) give a lengthier defense of such a minimalist line on properties in their response to the Benacerraf problem for metaethics.

paths lead creatures to think, talk, and form reliable beliefs about different mind-independently pre-existing colors.³⁰

I think conceptual diversity, (C-D), is a plausible model for moral concepts. When an Aztec applies a concept of moral permission to an act of human sacrifice by cutting out the heart, this is a different concept than the one used by the ancient Egyptian who applies his concept of moral permission only to sacrifices that leave the heart in the mummified body. I think of these as both in the family of moral concepts, but I insist they are distinct concepts: it is possible for both parties to be speaking or judging the truth when the Aztec is applying the one concept to heart extraction and the Egyptian is embedding the other concept in a denial about heart extraction. (C-D) says these two people have two distinct concepts of permission, and the contents of their words and concepts are different n -place relations.

2.4. Why Prefer (C-D) to Relativism?

(C-D), like relativism, would allow us to happily accept all the contingencies of our evolutionary history, and even the contingency of the diversity of actual moral communities. Everyone can win the lottery now. So, if we accept *Truth-Max*, how could we then decide between relativism or (C-D)? I favor (C-D), and my main argument, now, is that (C-D) avoids the mind-dependence of relativism, and thereby avoids the costly revisionism of relativism.

Moral relativists endorse contents that wildly defy common sense. Street says that an ideally coherent Caligula has no reason to stop torturing people, and Harman says similar things about (possibly fictionalized) Hitler and his reasons.³¹ Street and Harman have to say these things because of that extra relatum's appearance in the content. The reasons Caligula possesses, or lacks, depend on his m , and they thus depend on his mind. The dependence is built right into the logical form of the contents of these reasons-attributions: whether a given $(n+1)$ -length tuple appears in the $(n+1)$ -place relation, $\langle x_1, x_2, \dots, x_n, m \rangle$, partly depends on the value appearing in the last position, which could be the moral code fixed by Caligula's mind.

The dependence that the relativist builds in here is subtly by very importantly different from any dependence found on (C-D). For (C-D), whether I've said something true when I attribute reasons to Caligula does, in a way, depend

30. Horgan and Timmons (1996: 15, 22–23) themselves sketched, and dismissed, the basic idea of (C-D).

31. See Street (2016: 297; 2019: esp. §6); Harman (1996: 61–63; and 2000: 8–10). (Harman (1996) says Hitler would have no "objective reason". Again, Harman there also tentatively considers a "quasi-absolutist" language.)

on what concept I'm using; it depends on the sub-scripted m in $\langle x_1, x_2, \dots, x_n \rangle_m$. But *that* is the innocuous dependence that all claims have on our language and concepts. It's the same as the way that whether "Snow is white" is true depends on whether I'm speaking English. But that sentence is only about snow and whiteness, and not *about* English.

Relativists, on the other hand, are committed to something beyond innocuous conceptual dependence. It's a worldly form of dependence. They say we're talking *about*, among other things, m , when we attribute reasons. The inclusion of m inside the content gives our reasons attributions a kind of dependence we originally didn't think they have. And relativists must *say* there is dependence, because for any standard predicate attributing a relation, you can explicitly shift what relata you're talking about. When I say "If you are a student, you should listen to me.", I use "student" to implicitly attribute a two-place x -is-a-student-of- y relation, and I can make the implicit y relatum explicit and shift it in a way that generates something new and false: "If you are a student of anarchism, you should listen to me". The relativist likewise says, "If your mind is one way, you have reasons not to torture, but if your mind is like Caligula's, you lack such reasons."

If we endorse (C-D), we won't endorse anything so absurd as what relativists say about Caligula. We won't because we can't. By removing the relativist's $n+1$ th relatum from the content of our moral assertions and judgments, we cannot use our moral concepts to talk about it. We'd need to use some different concepts altogether. So, when I say Caligula has no reason to hurt you or me or that pig, I say something true, something that concerns n relata, and something that remains true however those n relata are further related to whatever " $n+1$ th relatum" you want. And that gets us the absolutism, and the sensible kind of mind-*independence*, that we wanted for minimal moral realism.

Objection: (C-D) lands us in a new and worse lottery! It's a lottery of concepts, and I could easily have ended up with a concept that applies to screaming at purple! That's horrifying.

Response: The worry here animates a recent project of Eklund and Dasgupta: how do we know we're using the *right* moral concepts?³² And my response is what Eklund (dismissively!) labels *the complacent reaction* (2017: 6). The complacent reaction has it that there are only metaphorical ways of gesturing at this allegedly metaphysically important notion of the "right" concepts, and therefore, as I believe, the whole notion is a mirage.³³ If the question is what the right concepts to use are, the answer is settled by the theory of content determination's assignment of a content to 'right', in whatever sense we just used that

32. See Eklund (2017) and Dasgupta (2018).

33. Eklund (2017: 6ff.) agrees he faces a serious inexpressibility problem in stating his project.

word earlier in this sentence. It's just the same as how the facts about the right ways to treat a pig are settled by the content of the word 'right' in this sentence.

Maybe we should feel *lucky* to have ended up with these moral concepts. Maybe or maybe not. It's fine if you feel that way. The important point is: we've now disarmed any argument from disagreement across a moral lottery. Nothing leaves it *unlikely* that you really did end up with those concepts you feel so lucky to have. We can enjoy them all we want.

3. Rationality-Max

Now I turn to the second of the two metasemantic approaches that I consider in this paper, and I'll argue that it, likewise, does not force us to accept moral revisionism.

3.1. *What is Rationality-Max?*

Earlier, when I suggested a motivation for *Truth-Max*, I mentioned that one function we serve with our intentional concepts is that of enabling useful testimony. Another famous function of intentional concepts is to allow us to predict and explain behavior. And this could motivate our second approach to content determination, which I will call *Rationality-Max*. This approach posits a close link between belief and behavior. It also posits a link between belief and something else, evidence, which could also help serve the predictive-explanatory function by helping us get some behavior-independent handle on a person's beliefs.³⁴

The approach, pioneered by David Lewis,³⁵ says a person has those attitudes that would maximize attributions not of truth but of rationality. Here's my statement of it:

Rationality-Max: the correct assignment of contents to a creature's beliefs is the assignment that maximizes the rationality of their beliefs and actions.

34. The posited link between belief and evidence might also help serve the function of supporting testimony. Or, we might hope to blend *Truth-Max* and *Rationality-Max*—our various motivations support such a possibility. I'm sympathetic to that; and Lewis (1974) might hold that overall view himself. Williamson's knowledge maximization approach, mentioned above, is another option for how to blend the two. My arguments, above and to follow, about escaping the moral lottery for each individual theory would apply to any blend I can imagine.

35. Lewis (1974). See Pautz (2013: §4), Schwarz (2014), and Williams (2020) for excellent exposition, elaboration, and more citations.

Lewis breaks *Rationality-Max* up into two sub-principles.³⁶ One has us maximize attribution of beliefs that are *epistemically* rational—given the person’s evidence up to that time. The second principle maximizes the attribution of beliefs (and desires) that make *practically* rational all the person’s actions—given not only the set of actions the person actually does perform but also those they counterfactually would perform.

What happens when attributing more epistemic rationality competes with attributing more practical rationality? Where there are such trade-offs, the correct attribution is indeterminately any of the “ties for best” maximization of both.

Those two “givens” in Lewis’s two sub-principles highlight how *Rationality-Max* rests on some significant presuppositions. It presupposes some independent account of what determines what a person’s *evidence* is and of what *actions* they do and would perform. Since evidence and action are both usually taken to be themselves representational or somehow content-dependent states or activities, *Rationality-Max* thus requires a serious amount of supplementation from other parts of the theory of content determination (as *Truth-Max* did too).

Robbie Williams (2020) develops an excellent book-length defense of the approach, and shows how one might plausibly fill out the theory of content determination here by borrowing from other independent reductive theories of evidence and action.

3.2. *Can Rationality-Max Save Us From the Moral Lottery? — Reason for Hope*

My question now is whether *Rationality-Max* could avoid landing us in the moral lottery. Or does *Rationality-Max* imply that rival moral communities disagree (in content—thus creating lottery conditions)? And when it comes to that question, I now want to draw attention to certain important complications that the answer will depend on. The complications I have in mind arise because of the independent accounts that we just saw *Rationality-Max* must be supplemented with, particularly the account of evidence.

Lewis and Williams take perceptual experiences to be our basic evidence. And Williams says he wants to borrow Karen Neander’s (2017) teleological theory (of what grounds perception’s having its contents) in order to have his reductive account of evidence.

36. Actually Lewis (1974) breaks it up into several principles, but I follow the cleaned-up exposition of Pautz, Schwarz, and Williams. And Lewis (1974) seems to also blend *Rationality-Max* together with *Truth-Max* as defined above, including all these ideas under the banner of the “Principle of Charity”.

But is perception our evidence for our *moral* beliefs? That's certainly controversial. Even if you call the evidence for moral beliefs "perception", it's nothing like what Neander's theory is concerned with. There is thus an important issue here—the issue of what makes moral beliefs epistemically justified. What is *moral evidence*?³⁷

This is especially important for our question about disagreement across rival communities, because evidence is (*very* plausibly) a content-involving mental state, and so if we attribute significantly different evidence to the rival communities, then *Rationality-Max* would seem to recommend attributing very different moral beliefs to them too. For just one example, there is the possibility that moral intuitions are our moral evidence, and that the intuitions' conceptual contents vary across communities. If the particular concepts featuring in the communities' evidence vary, then the concepts in their moral beliefs should vary too, namely in those beliefs *Rationality-Max* would tend to attribute because they are the beliefs it would be epistemically rational for them to have. This will look like the (C-D) view I proposed earlier for those who want to take the *Truth-Max* approach. This is a way that *Rationality-Max* could help us avoid the moral lottery.

I am raising this only as a possibility. My position is this: *Rationality-Max*, without a theory of moral evidence, does not by itself imply, one way or the other, whether rival moral communities disagree. And I do not have any particular theory of moral evidence I want to defend or endorse here.

What I want to do in the remainder of this paper is examine, and argue against, Williams's very different assessment of what *Rationality-Max* implies about disagreement.

3.3. *Or Does Rationality-Max Land Us In the Moral Lottery? — Williams's Argument*

Interestingly, in his clear and valuable article, Williams (2018; also 2020: ch. 4) argues that *Rationality-Max* explains and vindicates the intuition—which he calls "referential stability"—that we easily share the same moral concepts with a rival moral community in such a way that we *do* disagree.

However, although Williams initially advertises his paper as explaining referential stability, he ends up arguing that the intuition that's worth vindicating is one that must be refined. He argues the refinement is minor, but I'll argue the refinement is major, and, if I'm right, it means there's no threat of a moral lottery.

37. Similar questions arise about moral belief and the other side of *Rationality-Max*: how exactly do moral belief guide and justify action? Is it different from how non-moral belief guides and justifies action? But I won't further explore these important issues here, since my arguments below will only concern the issue of how evidence justifies moral belief.

Let me explain all this, and thereby explain how, contrary to Williams's view, *Rationality-Max* does *not* entail that there easily exist disagreements between rival moral communities.

Here's the outline of Williams's argument from *Rationality-Max* to referential stability, and then the refinement he introduces for that target intuition of stability.

For Williams, having a certain concept is having a mental representation that plays a certain role. There is some role—call it role *R*—such that when a mental representation has that role *R*, then that mental representation—call it *W*—is the concept of wrongness. Williams works with a particular view of what role *R* is, which he says he uses purely for illustration. He pretends that *W* has role *R* when we treat judgements applying *W* as the unique direct basis for blame.³⁸ When Williams says this is the unique *direct* basis for blame, he means there is no intermediate step to the inferential transition; it's like inferring that Tom is an unmarried man from the judgement that Tom is a bachelor.³⁹ Williams also pretends, again for purely illustration, that Kant's categorical imperative is the true moral theory. Williams then tells us that if a person—he calls her Sally—treats judgements applying *W* as the unique direct basis for blame, then, *even if she is an avowed utilitarian*, Sally would still be expressing the concept of wrongness, which, we are supposing, has Kantian truth conditions. Why? Why does Williams attribute to an avowed utilitarian a concept that has Kantian truth-conditions and that she systematically applies incorrectly (in a certain range of cases)? Because, though the attribution of the Kantian concept would make many of Sally's judgements about wrongness *inaccurate*, the attribution would maximize her *rationality*, and that's what *Rationality-Max* demands. That's precisely how *Rationality-Max* differs from *Truth-Max*. (We'll return to Sally shortly.)

Now here's the refinement. Williams considers someone—he calls her Suzy—who is just like Sally with regard to the "elimination" or "exit" rule for her use of *W*, that is, Suzy directly moves to blame just from judgements applying *W*. But, Suzy has an odd "intro" rule for *W*. Suzy takes judgements about failure to maximize utility as the unique *direct basis* for judgements applying *W*. When it comes to Suzy, Williams says it is *not* the case that her judgements determinately express wrongness. So, Suzy and we do not, at least not determinately, share the

38. I talk as if the *bases* of our judgments are mental states, like our judgements applying *W*, but we could, without affecting the present arguments, translate all this so as to talk as if bases are *states of the world* rather than mental states.

39. Williams also characterizes the thinker's inferential transitions as, not only direct but, *indefeasible*. I'm omitting that because it doesn't seem right or anyway helpful to Williams's case. (I think it's not right because no beliefs or inferences exhibit indefeasible justification. All justification is defeasible. All you need is a justifiably trustworthy Guru who is now mistaken.)

same moral concept. Why? Because *Rationality-Max* says we hold those beliefs that maximize the epistemic and practical rationality of our attitudes. While attributing the concept of wrongness to Suzy would make her exit inference practically rational, it would make her intro inference epistemically irrational. When there are such trade-offs, *Rationality-Max* makes no determinate attribution to the person.

But, going back to utilitarian Sally again, why does Williams determinately attribute to *her* our own concept of wrongness? Because, Sally *indirectly* inferred her *W* judgements from the premise of failure to maximize utility. She invoked a background theory of utilitarianism. (See 2018: 53, 61.) And Williams supposes Sally *falsely* but *rationally* believes that background theory.

Here's the question we need to face now: does the argument from disagreement, the debunkers' argument, Harman's argument, and Mackie's argument, capitalize on the existence of apparent disagreements with people like Sally, or does it concern people like Suzy? Did Street, Harman, Mackie, and others argue for moral revisionism by pointing to characters like Sally, or were the characters rather like Suzy? If the answer is Suzy, then, given *Rationality-Max*, the cases were *not* actually cases of determinate disagreements. And then the argument for revisionism fails. The conceptual diversity (C-D) view, or something like it, could hold.

But Williams thinks Suzy's case is rare at best. He says:

Suzy has an extremely strange (almost tonk-like) moral concept, in which she treats a person's failure to maximize hedons as *analytically* grounds for blaming them. More ordinary moral error [is not like Suzy's case].⁴⁰
(2018: 62)

If Williams is right here, and apparent ordinary disagreements are like the Sally case rather than the Suzy case, then the threat of disagreement and the moral lottery would indeed follow from *Rationality-Max*. But *is* Williams right here?

I disagree with Williams, as I'll argue next. I think ordinary moral disagreements involve cases like Suzy, and not Sally. And, worse, on some major philosophical views the Sally case is not even possible, and Williams didn't argue against or even discuss these views.

40. The bracketed expression replaces: "does not afford exceptions to Conceptual Role Determinism". Williams holds that "Suzy provides a case where Conceptual Role Determinism fails" (2018: 62).

Conceptual Role Determinism is here is the claim that the "local" conceptual role of *W*, a role that Suzy *shares* with Sally, determines what concept she has. That shared "local" role is just the direct reasoning between judgements applying *W* and blame. (See 2018: 50, 61.)

3.4. Why Moral Realism Is Still Safe from Rationality-Max

Williams says the Suzy case is strange and, by calling Suzy's inferences tonk-like, he's suggesting it may even be impossible. But I don't see the similarity he alleges between Suzy's inferences and tonk. Reasoning systematically using the tonk rules leads to chaos. Reasoning systematically using Suzy's rules leads to something like Peter Singer or Jeremy Bentham. Bentham himself explicitly held, like Suzy, that utilitarianism is *analytically* true.⁴¹ So, I don't agree with Williams that the Suzy case is strange.

More importantly, what's really controversial is whether the Sally case is even possible. This is because it's controversial whether it's as easy as Williams suggests for someone to be like he says Sally is, falsely but rationally holding her moral theory. Williams describes Sally in tendentiously(!) kind terms:

Sally mistakenly thinks to be wrong is to fail to maximize hedons. However, Sally was responsible in reaching this moral belief: she considered a range of actual and possible cases, and consulted and argued the issue through with trusted kith and kin. Ultimately, weighing her evidence, she endorsed utilitarianism. The contrast with the way that Suzy formed her *W*-beliefs is stark! (2018: 62)

It's not so clear, however, that it's that easy to form a responsible moral belief. Williams is making controversial assumptions about when a moral belief is epistemically rational—he's wading into that issue of moral evidence I raised earlier. Many ethicists hold views that are inconsistent with Williams's position. Consider Elizabeth Harman's view (found in a paper which Williams also references) that ordinary people often hold *epistemically* unjustified false moral beliefs even when they've been given testimony from people they reasonably view as authoritative, they've considered arguments for and against their view, some of those arguments do defeasibly support their view, and they've genuinely thought hard and honestly about the issue.⁴² Another alternative to consider is the view, held by Kieran Setiya (2012) among others, that the "basic" (2012: 40) evidence for any moral views consists entirely of *non-ethical* facts (e.g., facts about harming or treating others as a mere means) (see 2012: 5, 40, 48–55). If Setiya's view is right, then hard as Sally may have tried, her belief in utilitarianism (in the example) will simply be going against the total evidence, and thus her moral beliefs will be left epistemically unjustified. This leaves the contrast between Sally and Suzy very murky, not stark as Williams wants to say.

41. For some discussion of Bentham on this, see Schroeder (in press).

42. Elizabeth Harman (2011: 461–62). The view is also endorsed by Wedgwood (2019).

The complaint I'm making against Williams arises, I believe, because he focused on an illustrative conceptual role for moral concepts that features only an "exit" rule. You make *W* mean wrongness just by moving directly from *W*-judgments to blame. Williams only meant it as an illustrative possibility of what the role *R* could be (the conceptual role that makes a mental token, *W*, the concept of wrongness). But his chosen illustration subtly conceals the possibility for how to combine *Rationality-Max* with the view that rival communities don't disagree: feature in *R* an "intro" rule, a rule for how judgements applying *W* are based on moral evidence *and* give the communities different evidence. (This suggestion makes it very easy for it to determinately be the case that there are no disagreements between rival communities. Elizabeth Harman's and Setiya's views, if combined with *Rationality-Max*, would leave many of us like Suzy, for whom it's indeterminate what she is thinking, and so indeterminate whether there is a disagreement.)

To repeat, though, I don't have a theory of moral evidence that I endorse. What I mean to be showing is only that we need some argument, and we never got any argument, that the Sally case is even possible, much less an abundant and ordinary sort of case. If there are no Sallys around, there's nobody who Williams's argument shows us to be in moral disagreement with.

But the reader may wonder: wasn't there a whole planet of Sallys in the famous Moral Twin Earth thought experiment that inspired Williams in the first place? Didn't Horgan and Timmons, and didn't Mackie and (Gilbert) Harman, directly point us to some real or easy-to-imagine Sallys? I say, no, they didn't. It wasn't people like Sally who featured in all these prominent illustrations of the cases that drive the argument from disagreement.

In Horgan and Timmons's famous Moral Twin Earth thought experiment, we were told to imagine two very similar planets except that different properties regulate the moral terms on each planet. As Williams himself fleshed out Horgan and Timmons's case, "the citizens of Utilitas are disposed to blame other agents for taking hedonically suboptimal options. The citizens of Kantopia, on the other hand, blame other agents when they take actions that they cannot will as a general law" (2018: 44) The main point of the thought experiment was to elicit the intuition that people on the different planets disagree; at least one of them is getting things wrong. But nothing more was ever said, not by Horgan and Timmons or even by Williams, that would suggest to any reader that the mistaken planet is a population of Sallys. Did Horgan and Timmons intend that we imagine that everyone on a whole planet carefully examined the moral question, conferred with "trusted kith and kin", etc., and somehow they all took a wrong turn, like Sally did? I doubt that's what they were intending when they posed the case. But, again, on Williams's explanation of how *Rationality-Max* has them disagree, we'd need here a planet of Sallys. So, I don't think Williams

has explained how *Rationality-Max* secures disagreement in the original thought experiment.

Now think about actual moral disagreements, the kinds that Mackie and Harman pointed to. Do actual people resemble Sally who, again, we know “considered a range of actual and possible cases”, and “consulted and argued the issue through with trusted kith and kin”, and formed a verdict after “weighing her evidence”? Almost nobody has ever lived is like that. Does your average homophobe consider even for one second his background fundamentalist doctrines, entering them as a premises into his “reasoning”, when he “infers” that gay sex is morally wrong? No, I doubt there is any reasoning going on at all. He might bring up the background theory *if* you press him on what his reasoning is, but I really don’t think that it *antecedently* featured at all in his thinking, whatever thinking took place.

Of course, it’s an empirical question whether people reach false moral views directly (like Suzy) or indirectly (like Sally), but empirical matters are often common sense, and I think common sense is on my side here. I think Mackie (1977) was on my side too!

[P]eople judge that some things are good or right, and others are bad or wrong, not because—or at any rate not only because—they exemplify some general principle for which widespread implicit acceptance could be claimed, but because something about those things arouses certain responses immediately in them, though they would arouse radically and irresolvably different responses in others. ‘Moral sense’ or ‘intuition’ is an initially more plausible description of what supplies many of our basic moral judgements than ‘reason’. With regard to all these starting points of moral thinking the argument from relativity remains in full force. (1977: 37–38; see also his claim, 1977: 36–37, that our moral codes aren’t based on “inferences”, as in science, but simply “reflect ways of life”, partly quoted above.)

One last example: does a carnivore look at a pork chop, think it’s okay to eat it, and think that *partly because* he has a background theory that permits it? Of course not! The background theory never enters his thinking. And I don’t think it’s much different for the fundamentalist homophobes, or for me, or for other people in stock examples of apparent disagreement.

So, for all these reasons, my conclusion about *Rationality-Max* is that it doesn’t necessarily lead to a threatening moral lottery. To get from *Rationality-Max* to the lottery, you’d need to add on controversial extra assumptions that make the Sally case possible, including assumptions about the nature of moral evidence. And even then, you’d still have the uphill task of arguing, empirically, that the Sally

case is at all common in real life. Since many of the most worrisome versions of the moral lottery appeal to disagreement in actual cases (Mackie, Harman), or in easily possible cases (Street, at least as Bogardus 2016 and others flesh out her argument), the result is that the gap from *Rationality-Max* to any strong argument for moral revisionism is a large gap. That's good enough news to save moral realism, which enjoys enough pre-theoretical plausibility that we'd need a powerful argument *for* the lottery to give realism any trouble.

4. Conclusion

Throughout this paper, in the background of my thinking has been the ordinary language response to skepticism. I'm also guided by Philippa Foot's (1958) famous argument that, as I'd summarize it, an outsider who resists reasoning in the ways we would reason to normative conclusions from non-normative premises simply doesn't grasp our normative concepts.⁴³ Whether the connection with that literature is apparent or not, I hope my approach in this paper has helped to further show how sensible views in metasemantics and in metaethics can reinforce each another.

References

- Björnsson, Gunnar and Stephen Finlay (2010). Metaethical Contextualism Defended. *Ethics*.
- Bogardus, Tomas (2016). Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument. *Ethics*.
- Boghossian, Paul (2006). What Is Relativism? In Patrick Greenough and Michael Lynch (Eds.), *Truth and Realism* (13–37). Oxford University Press.
- Boghossian, Paul (2011). Three Kinds of Relativism. In Steven Hales (Ed.), *A Companion to Relativism* (53–69). Wiley-Blackwell.
- Boyd, Richard (1988). How to Be a Moral Realist. In Geoffrey Sayre-McCord (Ed.), *Essays on Moral Realism* (181–228). Cornell University Press.
- Brogaard, Berit (2008). Moral Contextualism and Moral Relativism. *Philosophical Quarterly*.
- Craig, Edward (1990). *Knowledge and the State of Nature*. Oxford University Press.
- Dasgupta, Shamik (2018). Realism and the Absence of Value. *Philosophical Review*.
- Davidson, Donald (1984). *Inquiries into Truth and Interpretation*. Oxford University Press.
- Dowell, Janice (2016). The Metaethical Insignificance of Moral Twin Earth. *Oxford Studies in Metaethics*.
- Eklund, Matti (2017). *Choosing Normative Concepts*. Oxford University Press.
- Foot, Philippa (1958). Moral Arguments. *Mind*.

43. Shafer-Landau (2012: 11–12) also applies Foot's point to the debunking debate.

- Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.
- Hare, R. M. (1952) *The Language of Morals*. Oxford University Press.
- Harman, Elizabeth (2011). Does Moral Ignorance Exculpate? *Ratio*.
- Harman, Gilbert (2000). Moral Relativism Defended. In *Explaining Value* (3–19). Oxford University Press.
- Harman, Gilbert (1991). Moral Diversity as an Argument for Moral Relativism. In Douglas Odegard and Carole Stewart (Eds.), *Perspectives on Moral Relativism* (13–31). Agathon.
- Harman, Gilbert (1996). Moral Relativism. In Gilbert Harman and Judith Jarvis Thomson, *Moral Relativism and Moral Objectivity* (1–64, 157–87). Blackwell.
- Harman, Gilbert (2019). Moral Realism is Moral Relativism. Retrieved from <http://www.princeton.edu/~harman/Papers/Published.html>. Published in modified form as “Moral Relativism is Moral Realism” (2015). *Philosophical Studies*.
- Hilbert, David (1987). *Color and Color Perception*. CSLI.
- Hills, Allison (2009). Moral Testimony and Moral Epistemology. *Ethics*.
- Horgan, Terence and Mark Timmons (1990–91). New Wave Realism Meets Moral Twin Earth. *Journal of Philosophical Research*.
- Horgan, Terence and Mark Timmons (1996). From Moral Realism to Moral Relativism in One Easy Step. *Critica*.
- Horwich, Paul (1998). *Truth* (2nd ed.). Oxford University Press.
- Khoo, Justin and Joshua Knobe (2018). Moral Disagreement and Moral Semantics. *Noûs*.
- Lewis, David (1974). Radical Interpretation. *Synthese*.
- Mackie, J. L. (1977). *Ethics*. Penguin.
- McGlynn, Aidan (2012). Interpretation and Knowledge Maximization. *Philosophical Studies*.
- McGrath, Sarah (2011). Skepticism about Moral Expertise. *Journal of Philosophy*.
- Neander, Karen (2017). *A Mark of the Mental*. MIT Press.
- Pautz, Adam (2013) Does Phenomenology Ground Mental Content? In Uriah Kriegel (Ed.), *Phenomenal Intentionality* (194–234). Oxford University Press.
- Plunkett, David and Tim Sundell (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers’ Imprint*.
- Quine, W. V. O. (1960). *Word and Object*. MIT Press.
- Sarkissian, Hagop, John Park, David Tien, Jennifer Cole Wright, and Joshua Knobe (2011). Folk Moral Relativism. *Mind & Language*.
- Schroeder, Mark (in press). The Common Subject Problem for Ethics. *Mind*. Available at <https://markschroeder.net/research>
- Schroeter, Laura and François Schroeter (2005). Is Gibbard a Realist? *Journal of Ethics and Social Philosophy*.
- Schroeter, Laura and François Schroeter (2013). Normative Realism: Co-Reference without Convergence? *Philosophers’ Imprint*.
- Schroeter, Laura and François Schroeter (2019). The Generalized Integration Challenge in Metaethics. *Noûs*.
- Schwarz, Wolfgang (2014). Against Magnetism. *Australasian Journal of Philosophy*.
- Setiya, Kieran (2012). *Knowing Right from Wrong*. Oxford University Press.
- Shafer-Landau, Russ (2012). Evolutionary Debunking, Moral Realism, and Moral Knowledge. *Journal of Ethics and Social Philosophy*.
- Stevenson, Charles Leslie (1944). *Ethics and Language*. Yale University Press.
- Street, Sharon (2006). A Darwinian Dilemma for Moral Realism. *Philosophical Studies*.

- Street, Sharon (2016). Objectivity and Truth: You'd Better Rethink It. *Oxford Studies in Metaethics*.
- Street, Sharon (2019). How to Be a Relativist about Normativity. Retrieved from <https://nyu.academia.edu/SharonStreet>. Accessed February 2019. Cited/quoted with permission.
- Thomson, Judith Jarvis (1996). Moral Objectivity. In Gilbert Harman and Judith Jarvis Thomson, *Moral Relativism and Moral Objectivity* (65–154, 188–217). Blackwell.
- Velleman, J. David (2013). *Foundations for Moral Relativism*. Open Book.
- van Roojen, Mark (2017). Evolutionary Debunking, Realism and Anthropocentric Metasemantics. In Diego Machuyca (Ed.), *Moral Skepticism: New Essays* (163–81). Routledge.
- Wedgwood, Ralph (2019). Moral Disagreement and Inexcusable Irrationality. *American Philosophical Quarterly*.
- Williams, Robert (2018). Normative Reference Magnets. *Philosophical Review*.
- Williams, Robert (2020). *The Metaphysics of Representation*. Oxford University Press.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Blackwell.
- Wilson, Neil (1959). Substances without Substrata. *Review of Metaphysics*.
- Wong, David (2006). *Natural Moralities*. Oxford University Press.