

SEMANTIC AND MORAL LUCK

AXEL ARTURO BARCELÓ ASPEITIA

Abstract: The similarities between the philosophical debates surrounding assessment sensitivity and moral luck run so deep that one can easily adapt almost any argument from one debate, change some terms, adapt the examples, and end up with an argument relevant to the other. This article takes Brian Rosebury's strategy for resisting moral luck in "Moral Responsibility and 'Moral Luck'" (1995) and turns it into a strategy for resisting assessment sensitivity. The article shows that one of Bernard Williams's examples motivating moral luck is very similar to one of the examples John MacFarlane uses to motivate the assessment sensitivity of epistemic modals, and in particular the assessment sensitivity of the auxiliary verb "might." This means that, if Rosebury is right and we do not actually need moral luck to explain Williams's example, we may not need assessment sensitivity to account for the semantic behaviour of the epistemic modal verb "might" either.

Keywords: assessment sensitivity, relativism, luck, apologies, epistemic modals, might, control principle.

Introduction

According to John MacFarlane (2005a, 2005b, 2011), epistemic modalities and other expressions of natural language show what he calls "assessment sensitivity," that is, their basic semantic properties depend not just on features of the context in which they are used but also on features of the context in which they are assessed. Traditional (that is, indexical) contextualism inability to account for this has motivated MacFarlane to develop a new, non-indexical variety of contextualism. This kind of contextualism, also known as relativism, has faced enormous opposition from philosophers and linguists. In particular, they find the very idea of assessment sensitivity unpalatable because it goes against the intuitive principle that our assertions can be normatively regulated only to the extent that the semantic content of our assertions depends on factors under our control.

More than thirty years ago, Bernard Williams (1981) and Thomas Nagel (1979) argued that luck substantially influences our moral judgment of a person and her actions. They based their claim on the consideration of cases where "a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment" (Nagel 1979). A truck driver who

blamelessly runs over a pedestrian, for example, exemplifies what he calls (bad) “moral luck.” In the following decades, many moral philosophers have challenged Williams’s and Nagel’s claim. In particular, they find the very idea of moral luck unpalatable because it goes against the intuitive principle that we are morally assessable only to the extent that what we are assessed for depends on factors under our control.

The present article has two goals. The simpler one is to weaken one of MacFarlane’s main motivations behind semantic relativism. The second but more fundamental aim is to show that there are fundamental similarities between the aforementioned two philosophical debates. Such is the similarity between these two debates that one can easily take practically any argument from the former debate, change some terms, adapt the examples, and end up with an argument relevant to the latter. Actually, that is what I do in the second half of the article: take a well-known argument against moral luck and turn it into an argument against relativism.

The article is structured in two parts. The first presents what I take to be the main similarities between the current debate around relativism in semantics and the metaethical debate surrounding moral luck. I start by presenting one of the semantic puzzles that motivate semantic relativism. Then I introduce one of the most common and general worries many authors have with relativism, and show that it shares some essential features with a similar concern many ethicists have about moral luck. Finally, I give a broad overview of the debate around moral luck, highlighting its similarities with the debate surrounding assessment sensitivity. In the second part of the article, I take Brian Rosebury’s strategy for resisting moral luck in his 1995 article “Moral Responsibility and ‘Moral Luck’” and turn it into a strategy for resisting relativism. I show that one of Williams’s main examples motivating moral luck is very similar to one of the examples MacFarlane (2011) uses to motivate his relativism. This means that, if Rosebury is right, that is, if we do not actually need moral luck to explain Williams’s example, we may not need assessment sensitivity to account for MacFarlane’s example either.

Why Go Relativist?

Relativism in semantics is the claim that there are expressions in our natural language that are contextually dependent on features of the context of assessment, rather than on the usual context of utterance. Accordingly, relativism is motivated by several puzzling semantic phenomena that regular contextualism seemingly cannot explain. Consider the following scenario, from MacFarlane (2011): Sally’s mother comes into Sally’s bedroom to find her looking under the bed. “What is going on?” asks Sally’s mother, “why are you looking under the bed?” “My glasses, they might be there,” replies Sally. After taking a long look under

the bed, Sally finds no glasses under it. So she moves on to look in other places, but not before saying, “Oops, I was wrong.”

Simplifying the example, let us consider Sally asserting the following two sentences, the first one just *before* unsuccessfully looking for her glasses under the bed (call it time t_1) and the second right *after* (call it time t_2):

- (1) My glasses might be under the bed.
- (2) I was wrong.

Prima facie, both assertions seem to be justified. Both times Sally seems to be justified in asserting what she did. Before she looks under the bed, there seems to be nothing wrong in Sally’s affirming that her glasses might be there. Yet she later seems also to be justified in asserting that she was wrong. We are not inclined to say that she said something false, misleading, or in any other way unjustified either time. However, it seems that, in order for Sally to be justified in asserting (2), there must have been something wrong in her assertion (1), and vice versa. Therefore, a puzzle rises from our conflicting intuitions that what Sally asserted in both (1) and (2) was right.

To solve the puzzle, one must determine what was asserted in (1), what was retracted in (2), and whether it was the same proposition or not. If what was asserted in (1) was later retracted in (2), one must explain how it is possible for Sally to be right in asserting it at time t_1 and later in rejecting it at time t_2 . Otherwise, one must explain what *wrong* Sally is recognizing in (2), if not her assertion of (1). Relativism bills itself as a plausible solution to the aforementioned puzzle: one where what Sally says in (1) is true, what she says in (2) is false, and yet both times her utterances express the same proposition!

The relativist’s argument goes as follows. There are just a few salient propositions that might have been asserted in (1) and/or retracted in (2). Without loss of generality, we may consider three kinds of candidate propositions:

- (3) Sally’s glasses were under the bed.
- (4) For all Sally knew *at time* t_1 (the time of her assertion), her glasses could have been under the bed.
- (5) For all Sally would come to know *at time* t_2 (the time of her retraction), her glasses were under the bed.

Let us first check whether any of these candidates could have been the content of Sally’s first assertion or not. According to the relativist, Sally could not have been asserting (3), that is, that Sally’s glasses were under the bed in (1), because Sally could not have known whether or not her glasses were under the bed before she looked for them. Therefore, at that

time Sally could not have been justified in asserting that they were there. However, she was justified in asserting that they *might* have been there. Consequently, the content of her assertion could not have been (3). This means that the word “might” in (1) adds something to the *content* of the assertion, not just force.

The content of Sally’s first assertion (1) could not have been (5) either. Before she looked under the bed, she could not affirm that, for all she would come to know at the time of her retraction, her glasses would be under the bed. Even at the time of her assertion, Sally already knew that after looking under the bed, she would come to know whether her glasses were there or not. What she did not know was what the outcome of her search would be. Therefore, even if she had found her glasses under the bed, she could not have known that before looking. Therefore, she could not have been justified in asserting (5) at t_1 . Consequently, (5) cannot be what Sally asserted in (1) either.

Remaining candidate (4), however, does sound promising. For all Sally knew at the time of her assertion, her glasses could well have been under the bed. Not only is this something true, it is also something Sally knew at the time of her assertion. Therefore, it would have been right for her to assert it at time t_1 .

Let us turn now to the question of what was retracted in (2). This time, (3) seems a promising candidate, for this is something Sally came to know to be false after looking under the bed. What she discovered from looking under the bed was that her glasses were not there. The same holds for (5). It is also something she came to know at the time of uttering (2). Therefore, (3) and (5) could well be what she was retracting in (2). Proposition (4), in contrast, is still true. Even after Sally is done looking under the bed, it remains true that, for all she knew before, her glasses could have been there. Thus, (4) is not something she could justifiably retract at t_2 .

To summarize, (4) fits better our intuitions about what was asserted, but not about what was retracted, while (3) and (5) fit better our intuitions about what was retracted, but not about what was asserted. If what was asserted in (1) was also what was retracted in (2), it must have been something that Sally knew at t_1 (just like (4) above), but also something she could rightly retract at t_2 (like (3) and (5)). Presumably, it must have been something that was true at t_1 (like (4)) but false at t_2 (like (3) and (5)). There cannot be such a proposition within the limits of traditional contextualism. Therefore—relativists conclude—we require a new kind of proposition that is sensitive both to features of the setting or context of assertion (t_1) and to the context of its assessment (t_2). What we require is what they call an assessment-sensitive proposition, something like:

- (6) For all Sally knew in the context of assessment c , her glasses were under the bed.

This is precisely how MacFarlane interprets Sally's puzzle, since he writes: "After Sally learns [from looking under the bed that her glasses are not under it], she occupies a context of assessment relative to which her original claim is false (since she now knows more than she did). So it is proper for her to retract it" (MacFarlane 2011, 161). Thus, MacFarlane's interpretation of Sally's puzzle delivers an argument for assessment sensitivity:

- P1. What is asserted at t_1 is true at t_1 (and Sally knows it at t_1).
- P2. What is retracted at t_2 is false at t_2 (and Sally knows it at t_2 , but not at t_1).
- P3. What is asserted at t_1 is what is retracted at t_2 .
- C. Therefore, what is asserted at t_1 and retracted at t_2 is true at t_1 (and Sally knows this at t_1) and false at t_2 (and Sally knows this at t_2).

In the second part of this article, I develop a few strategies to challenge this argument. As motivation, in the following section I present a very common worry many regular contextualists have with assessment sensitivity. As I said before, this worry is analogous to a common one many ethicists have with regard to the notion of moral luck, and it will help me ground the analogy between moral luck and assessment sensitivity.

What Might Be Wrong with Relativism?

Debates on relativism commonly address questions of two different kinds. On the one hand, there are foundational questions surrounding the very idea of assessment sensitivity, like: Is it self-undermining? Does it mesh well with the rest of our semantics? Is it tractable, stable? And so on. On the other hand, there are descriptive debates evaluating how well relativism can account for the semantic phenomena that supposedly motivate it. For example, how well does the relativist proposal solve Sally's puzzle above? Are there any other solutions? And if so, what are their relative virtues and flaws?

Recently, Manolo García-Carpintero (2008), following earlier remarks by Jeffrey King (2003) and Gareth Evans (1985), criticized MacFarlane's assessment-sensitivity account of assertion for making its normative stance intractable. According to King and García-Carpintero, by making truth conditions relative to contexts of assessment, which are not in any strong sense dependent on the context of use,¹ MacFarlane has made truth conditions dependent on things (possibly) well beyond the asserter's

¹ So that no information available in the context of assertion may determine all possible contexts or standards of assessment. Contexts of assessment may not be merely the context of the hearers or whoever the assertion may be directed to, but may include any contexts from which the assertion in question may be challenged, including those that go beyond the asserter's lifetime.

information and control. Given the truth-commitment account of assertion that MacFarlane favors, this makes many (probably most) assertions commit the asserter to the truth of a proposition whose truth conditions are intractable from the asserter's standpoint.

We take speech acts like assertions as central cases of intentional, rational action. Rational agents should be at least by default in a position to do whatever is necessary to perform them correctly; they at least should have some measure of control about that. But it is unclear how this could be the case, if truth were relativized in the way suggested by proponents of truth relativism. How can I rationally take responsibility for making correct assertions, if the correctness or otherwise of my assertions depends on parameters set at different contexts of evaluation about which I lack whatever information, in ways on which I have no control? (García-Carpintero 2008, 141)

For both García-Carpintero and MacFarlane, assertion requires some form of commitment to truth on the part of the asserter. This commitment is best satisfied when the speaker *knows* the truth of what she says.² Yet, for King and García-Carpintero, if relativism is right, too little information is available to the speaker at the time of assertion to know under which conditions what she is saying is true or not. This makes assertion too burdensome on the speaker, almost to the point where it is impossible for her to honor her commitment to truth.

MacFarlane agrees that much important information is unavailable to the speaker at the moment of assertion. However, he does not think that this makes assertion too burdensome to be carried out by us, normal human beings. In his proposal, he recognizes that assertions may have *unexpected* contexts of assessment (by eavesdroppers, for example, or people who may find out about the assertion in unexpected circumstances), that is, contexts of assessment that could not be predicted from the information available to the asserter *in the context of use*. In MacFarlane's account, in order for the asserter's commitment to actually compel her to do something (to retract or to defend her assertion, for example), she must know enough about the context from where the truth of her assertion is challenged (or new evidence is considered) to fulfill such commitment. Commitment to truth starts in the context of assertion, but does not end there. New challenges and new evidence can trigger it at any time in the future.

Challenging and providing new evidence are acts whose value and effects are also context sensitive. The truth of (the content of) an assertion is challenged (and the challenge is met) relative to (epistemic standards and shared knowledge at) a context of assessment. Similarly, new evidence is relevant only relative to (the epistemic standards and shared knowledge

² García-Carpintero prefers to talk about "what is said," while MacFarlane uses "the content of the assertion." I find no substantive difference and will use both expressions without distinction.

in) a context of assessment. One is not obliged to defend one's assertion against challenges one does not know of, or from contexts one is not able to situate oneself in. Similarly, one is not committed to retract from new evidence that one is not faced with.³

The main difference between the contextualist and the relativist seems to boil down to determining just *when* the normativity of assertion kicks in. García-Carpintero thinks that whatever determines the extent of one's commitment to truth in assertion must be available to the asserter in the context of use. In other words, with the world of evaluation fixed, the truth of what one asserts must be determined at the moment and in the context of assertion. MacFarlane, in contrast, thinks that the sort of normativity involved in assertion is more forward-looking. Thus, even after the world of evaluation and all elements sensitive to the context of assertion have been fixed, the truth of what one asserts may still depend on circumstances that may drastically change *after* the assertion is made.

This kind of difference in perspectives is well known in metaethics, where there exists a long ethical tradition for which moral responsibility and judgment must be based on elements available and assignable to the agent *at the time of the action*. However, this tradition has been strongly challenged by consequentialists and others, for whom the moral value of an action substantially depends on its consequences, even if those consequences could not be determined at the time. This parallelism between the semantic and the ethical cases must come as no surprise to anyone who notices that García-Carpintero's and King's criticism of relativism is grounded in the assertoric analogue of the so-called Control Principle in Ethics:

(Control Principle in Ethics) We are morally assessable only to the extent that what we are assessed for depends on factors under our control.

Williams (1981) and Nagel (1979) argued against this principle in their seminal pair of articles. There, they tried to show that our everyday moral judgments and practices commit us to the existence of moral luck, that is, cases "where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment" (Nagel 1979, 175). Thus, according to Nagel and Williams, "it is impossible to morally assess anyone for anything if we adhere to the Control Principle" (Nelkin 2008). Now, if we just remove the word "morally" from the above formulation of the control principle, and change it so that it is no longer

³ This last point is central to understanding the subjective nature of matters like taste (and how they are different from other assessment-sensitive expressions, such as epistemic ones). For, in matters of taste, it is very difficult for challengers to make their standards of evaluation clear and explicit enough for the speaker whose assertion is being challenged.

about moral value but about the truth conditions of assertions, what we get is an analogue principle for assertion:

(Control Principle for Assertion) One cannot rationally take responsibility for making true assertions, if the truth conditions of one's assertions depend on what is the case in contexts about which one lacks any control.

According to relativists, on the other hand, our everyday semantic judgment and practices commit us to the existence of what I am calling "semantic luck," that is, cases where a significant aspect of what someone asserts (that is, its truth conditions) depends on factors beyond his control, and yet we continue to treat him in that respect as committed to the truth of what he asserted. Relativism, therefore, entails the claim that it is impossible to be responsible for many kinds of assertions if we adhere to the Control Principle for Assertion.

To further illustrate the similarities between moral and semantic luck, consider the following characterization of moral luck, by Andrew Latus (2001):

The problem of moral luck traps us between an intuition and a fact:

- 1) the intuition that luck must not make moral differences (e.g., that luck must not affect a person's moral worth, that luck must not affect what a person is morally responsible for).
- 2) the fact that luck does seem to make moral differences (e.g., we blame the unfortunate driver more than the fortunate driver).

Responses to the problem have been of two broad sorts. Some claim that the intuition is mistaken, that there is nothing wrong with luck making a moral difference. Others claim that we have our facts wrong, that luck never does make a moral difference. The first sort of response has been the least popular. (Latus 2001)

For the relativist, semantic luck can be characterized by a similar tension, since it also seems to trap us between an intuition and a fact:

- 1) the intuition that luck must not make semantic differences, for example that luck must not affect what a person is responsible for when he makes an assertion
- 2) the fact that luck does seem to make semantic differences (e.g., we apologize from unfortunate falsehood).

Just as in the case of moral luck, responses to the problem of semantic luck have been of two broad sorts. Some (relativists) claim that the intuition is

mistaken, that there is nothing wrong with luck making a difference in truth conditions. Others (antirelativists) claim that we have our semantic facts wrong, that luck never makes a difference in an assertion's truth conditions.

Before I develop the analogy further, I must emphasize that, despite the strong parallelisms between moral and semantic luck, they are still two different phenomena. The motivations behind the postulation of these two kinds of luck are very similar, but not quite the same. The assertoric version of the control principle is not a direct application of the control principle in ethics. Consequently, semantic and moral luck do not need to stand or fall together. Williams and Nagel may be right about the role of luck in morality, and MacFarlane wrong about its role in semantics, or vice versa. However, if the similarity between these two debates is as robust as I claim, semantic theorists debating assessment sensitivity may still have something to learn from the ethical debate surrounding moral luck. I will try to demonstrate this here by way of example. In particular, I will borrow an old argument against moral luck from Brian Rosebury's influential 1995 article "Moral Responsibility and 'Moral Luck'" and try to adapt it to work against assessment sensitivity.

Luck and Responsibility

Rosebury's article is aimed against one of Williams's original examples motivating moral luck. As such, it might not be cogent against Nagel's motivations or his version of moral luck. Consequently, from now on, when talking about moral luck, I will follow Williams's account. In particular, Rosebury's article is aimed at Williams's example in which "a truck driver accidentally kills a child; despite being innocent even of negligence, he will feel worse than any spectator, and though people will rightly seek to move him away from this feeling, [and] it is important that this is seen as something that should need to be done and indeed some doubt would be felt about a driver who too blandly or readily moved to that position" (Rosebury 1995, 514–15). So described, this is in Williams's view a clear case of moral luck. On his interpretation, (i) the truck driver is innocent, but (ii) he is unfortunate. He is innocent because the fatal accident cannot be traced back to any negligence on his part. Furthermore, (iii) he is justified in feeling bad because (iv) he actually did something awful: he killed a child. For Williams, this is a clear case of moral luck, because the truck driver is both an agent of evil and completely innocent.

In his article, Rosebury challenges Williams's interpretation of the truck driver case. He questions both the truck driver's innocence and the reason why he feels bad about what happened. Rosebury does not try to demonstrate that Williams's interpretation is wrong, only to weaken the intuition behind it, showing that other interpretations are equally plausible. Thus, he offers three other plausible interpretations where there

is no moral luck at play, that is, where the truck driver is either not innocent or has done nothing reprehensible.

(a) First, Rosebury offers an interpretation where the truck driver is not innocent but negligent. The main insight behind this interpretation is that, given what was at stake (a child's life), the truck driver should have been more careful. If he had noticed that a child's life was at stake, he probably would have been more careful. Yet, he did not, and so he was not. For Rosebury, care and negligence are stake-relative notions (1995, 515). One may be negligent or careful, depending on how much is at stake. Higher stakes demand higher standards of care. In this case, the truck driver might not have realized how high the stakes were until the fatal accident happened. That might explain why he feels so bad and deserves our sympathy. In this scenario, even though (iv) is true, that is, the truck driver feels bad for doing something awful, there is no moral luck because conditions (i) and (ii) for moral luck are not met. First of all, (i) is false, that is, the truck driver is not innocent of negligence but guilty. Second, since the truck driver is not innocent, he is justified in feeling bad, not because of the unfortunate consequences of his actions, but because of an irresponsible fault of his own.

(b) Moreover, we could also interpret the example the other way around. It might have been the case that, even considering that a child's life was at stake, the truck driver was appropriately careful, yet unlucky enough to kill the child. That is not enough to make this a case of moral luck, yet, argues Rosebury, because what the truck driver has done may not be morally wrong, (iv) above may be false, precisely because (i) the driver is *completely* innocent. Therefore, if he did nothing wrong, there must be other reasons why he feels bad. Rosebury has us consider the possibility that the truck driver perhaps "feels awful, because he has been *the instrument* of another's death. . . . For who could possibly bear being part of an accident that ends with a killed child?" (Rosebury 1995, 516; emphasis in the original). Thus, the truck driver's (iii) being justified in feeling bad does not entail (iv) any wrongdoing on his part, because what he feels is neither *guilt* nor *remorse*.

(c) Finally, Rosebury shows that it is possible that (iii) be false too: that is, the truck driver may not be justified in feeling bad after all. For example, this could be the case if he is unduly *blaming himself* for something he is not responsible for. In this case, in contrast to the previous two, it would not be adequate to show sympathy. Instead, we must "dissuade him from doing so by pointing to the absence of agency on his part" (Rosebury 1995, 516). Saying something along the lines of "Don't feel bad, it was not your fault" may be the proper thing to say to the truck driver here.

According to Rosebury, any of these alternative interpretations may be appropriate in the situation. We may not know which one is the right one, because Williams's example is underspecified. Williams is especially vague

when he talks about the truck driver *feeling bad*, for there are certainly many ways of feeling bad. So, we do not know, for example, if the truck driver is feeling *grief, sorrow, horror, self-reproach, regret, guilt, embarrassment, or shame*. All these are ways of feeling bad, yet only the last few are “moral sentiments,” that is, only the last few are associated with some particular moral judgment. One can easily be horrified at things one does not consider morally reprehensible. Grief comes from loss, not from evil; and sorrow is broad enough to cover both moral and nonmoral distress. This means that there are many kinds of bad feelings the truck driver may be experiencing, independently of (his own appraisal of) the moral status of his actions. Furthermore, even if the truck driver’s feeling was moral—that is, some sort of reproach, regret, guilt, embarrassment, or shame—it is unclear what he is *feeling bad about*. He may be feeling bad about something he has done. But he may also be feeling bad about something someone else did. Perhaps from the truck driver’s perspective it was someone else’s actions that were responsible for the child’s death, and the moral judgment that triggers his moral emotions is not directed against himself.

The Semantic Turn

After this brief account of Rosebury’s main point, it is time to see how well it transfers from the ethical to the semantic arena. Thus, it is necessary, first, to determine how much MacFarlane’s interpretation of Sally’s case is similar to Williams’s interpretation of his truck driver example. My contention will be that the truck driver is as justified in apologizing and feeling bad about what happened, despite not having done anything *bad*, as Sally is justified in apologizing and feeling bad about what happened, despite not having said anything *false*.

In order to see how similar these two cases are, notice that Sally’s situation fits Williams’s diagnosis of the truck driver’s situation perfectly. According to Williams, the truck driver was innocent, yet justified in feeling bad about the accident he was involved in. In Sally’s case, according to MacFarlane, she was also innocent, yet justified in her retraction. In the case of the truck driver, his bad feeling was significantly stronger than any bad feeling any innocent bystander may justifiably feel. Similarly, Sally’s retraction cannot be mistaken for any sympathetic utterance of “I’m sorry” that any innocent bystander may make. Furthermore, just as other people may rightly seek to move the truck driver from his woeful reaction to the accident, people may also rightly seek to move Sally away from her retraction. Finally, in both cases, Sally’s and the truck driver’s reaction is seen as something that was not only justified but morally mandatory or even virtuous.

Sally’s case fits Williams’s diagnosis so well that it would certainly be surprising if it did not fit Rosebury’s prognosis as well. It would be

surprising if Rosebury's criticisms against Williams's interpretation would not work against MacFarlane's interpretation too. Thus, if we adapt them to the semantic case, we may have a way to weaken MacFarlane's reading of Sally's retraction as a case of genuine assessment sensitivity. That is what I will try to do in the remainder of the article.

Remember that, according to Williams, the truck driver's accident exemplified a genuine case of moral luck because it satisfied four clear conditions: the truck driver was (i) innocent, (ii) unlucky, and (iii) justifiably feeling bad after the fact because (iv) he did something awful. Thus, he did both nothing bad (so he is innocent) and something awful (so he is unlucky). In other words, the truck driver had the bad (moral) luck to do something awful. In a very similar fashion, Sally's case is assessment sensitive because it satisfies four similar conditions: Sally is (i) innocent (of asserting that the glasses might be under the bed), (ii) unlucky (that the glasses were not there), and (iii) justified in her post facto retraction because (iv) she actually said something false. Rosebury has given us a way to reinterpret the truck driver example so that conditions (i) to (iv) do not hold. In particular, he has shown us ways to interpret the truck driver case so that either

- (a) the truck driver did something wrong, so he is not innocent after all (not (i)),
- (b) he is innocent, so he did nothing wrong and is justified in feeling bad for some other reason (not (ii) or (iv)), or, finally,
- (c) he is innocent, but not justified in feeling bad afterward (not (iii)).

In none of these three plausible interpretations did the truck driver innocently do something awful and, therefore, there is no moral luck. The purpose of the rest of this article is to give similar new interpretations (a) to (c) for Sally's case, where assessment sensitivity does not occur, because at least one of the four conditions (i) to (iv) is not satisfied.

Case (a): Sally Is not Innocent

Let us start with case (a), where Sally could actually have been negligent. Remember that Rosebury suggests that the difference between being negligent or being careful depends on what the stakes are in a given context. So, when we think that the truck driver is not negligent, we may be using the wrong standard of safety. Once we realize, because of the accident, that a child's life was at stake, we realize that those standards were actually higher than we thought, and that according to them, the truck driver was actually negligent.

Something similar may be said about our appreciation of Sally's justification for her assertion. We may say that Sally's use of "might" makes her assertion sensitive to a contextually determined epistemic standard, so

that whether she was justified in asserting that her glasses *might* be under the bed depends on those standards. Sally could have been right, under a certain *low* epistemic standard, in asserting that her glasses might be under the bed (since she did not know that they were not there). But she might also have been wrong under a *higher* epistemic standard (since she did not know that her glasses were there). Whether or not she was actually justified depends on what the relevant standard actually was. If the actual standard had been low enough, she would have been justified at t_1 , but would have nothing to retract at t_2 . If the relevant standard had been high enough, however, Sally would not have been justified at t_1 , and this is what she could have acknowledged at t_2 . On this interpretation, after looking under the bed, Sally realized not only that her glasses were not there but also that the epistemic standard relevant for (1) were higher than she originally thought. In other words, she realized that she had not been justified in asserting (1) and, instead, should have been more careful. If this is the case, in uttering “I am sorry” Sally might have retracted from having used too low an epistemic standard (just as the truck driver had felt bad for driving with too low a standard of safety).

The temptation to think of this as a case of semantic luck may well be explained by the fact that our appreciation of the epistemic standards also changed between t_1 and t_2 . At t_1 , we thought the standards were low enough for Sally to justifiably assert (1), but by t_2 they had risen enough to make the assertion unjustified. That is why we first thought she was justified and then realized she was not. Just like Sally, we were wrong at t_1 , but did not find out until t_2 . Once we recognize this, the motivation to postulate any assessment-sensitive proposition disappears. We can interpret Sally’s utterances so that she is no longer innocent, no semantic luck is involved, and no assessment-sensitive propositions are required.

Case (b): Sally Did Nothing Wrong

Now we can move on to case (b), where Sally is justified in saying “I was wrong,” not because what she asserted was false, but for some other justified reason. Contextualists like Price (1983), Barker (2009), and Barker and Taranto (2003) have all held some version or other of this interpretation for cases like Sally’s. According to Barker, for example, even though Sally was justified in apologizing for being *wrong*, she could not have been equally justified in asserting that what she said was *false*. Consider the following four possible sentences Sally could have uttered after looking under the bed:

- (2) Oops, I was wrong!
- (2′) Oops, I am sorry!
- (2′′) Oops, they were not there!
- (2′′′) Oops, it was false!

Notice that (2), (2'), and (2'') seem fine, but not (2'''). According to Baker, this is further evidence that falsehood is in no way involved in Sally's apology.

Once the burden of finding a false proposition is removed, we can look for other plausible explanations for Sally's apology. For example, it is possible to find a scenario where, by uttering (2) at t_2 , Sally retracts her assertion of (1) at t_1 but recognizes no fault of her own. After all, it makes perfect sense to retract an assertion of something because of its negative consequences, even if those consequences were unintended or unforeseeable. Thus, it makes sense for Sally to retract her assertion that her glasses were under the bed, not because she believes that she was unjustified in so asserting, but because her assertion had negative unintended consequences (like wasting some valuable time looking under the bed).

To further ground this latter possibility, suppose Sally was not on her own searching for her glasses under the bed but was leading a search group, looking for a fugitive, in a dangerous combat zone. In this context, her apologetic stance seems even more natural. She not only thought that the fugitive was in that zone, she also *acted on it*. She took the search group into the danger zone because she thought the fugitive *might* be there. She took a risk and, in the end, it did not pay off. Like the glasses, the fugitive just wasn't there. Her conjecture had a price, and it was paid not only by her but by everyone who risked it in the combat zone. It is in solidarity with those who followed her that she offers the apology "Oops, I was wrong."

Following Rosebury, we can note that saying "Oops, I was wrong" is very uninformative, for it may express a wide variety of feelings, many of which may not entail any recognition of negligence. Just as the truck driver could have felt awful about his involvement in the child's death without assigning himself any blame, so could Sally have felt bad about wasting her time looking fruitlessly under the bed, for example, without recognizing any negligence of her own. As a matter of fact, we frequently apologize for things we are innocent of. For example, it is not rare to apologize for so-called innocent mistakes. In doing so, we are apologizing without recognizing any negligence. One is also justified in feeling sorry for the unfortunate consequences of one's actions, without recognizing that such actions were wrong or that one was negligent or unjustified in doing them. Either way, we have an alternate explanation for Sally's retraction (and the truck driver's pain) that does not involve moral or semantic luck.

In general, one commonly feels bad about getting things *wrong*, even when one could not do better; and this is what Sally might be expressing in (2). So, when Sally says, "I was wrong," she may not be asserting that she was unjustified in asserting (1), only reiterating that her search was fruitless. Maybe Sally is recognizing that it was wrong to look under the bed. Not because she did not have good reasons to look there (after all, for

all she knew, her glasses could have been there), but because she did not find them there. After all, it is common to say that a search was wrong, not because it was unjustified, but because it was fruitless. Searching is a goal-oriented activity. As such, it aims at finding. The aim of Sally's search under her bed was to find her glasses. Since they were not there, her search was fruitless and as such wrong. However, the way it was wrong is completely compatible with Sally's being as careful as was required.

There are at least three ways in which a goal-oriented activity like a search may be *wrong*. First, it may be *unjustified*. We may not have good enough reasons to look for something somewhere. Thus, to look for something in a place where we have no good reason to think it might be is wrong in this first way. In a second way, a search may be *so badly performed* that, in the end, we may still be at loss as to whether the object we were looking for was there or not. Finally, a search may be *fruitless*, that is, the object may not be found where we look. Unlike the previous two senses, this third one entails no negligence on the part of the agent. We may have good reasons to search in some particular place for one particular object, perform the search with due diligence, and yet not find the desired object. Thus, the activity may be justified and careful and yet be unfruitful. In any of the three cases, one may justifiably say that something was wrong with the search. But the relativist interpretation takes only the first one into account. It is only the first one that motivates the relativist puzzle and calls for semantic luck. For the other two, the puzzle is a nonstarter. What went wrong in Sally's search had nothing to do with the truth of (1). Therefore, we have a plausible scenario where Sally is innocent (she was justified in asserting what she did), yet there is no need to postulate any assessment-sensitive proposition, because her assertion remains true even at t_2 .

Case (c): Sally Is Blaming Herself

Finally, we should also consider the possibility that Sally's retraction at t_2 was unjustified. In the original description of the puzzle, it seemed that Sally was justified in saying "I was wrong." But this may also be mistaken. Depending on how Sally delivered her retraction, it is also possible that Sally was *blaming herself*, that is, taking the blame for the falsity of something she did not assert. In other words, she may be wrong in saying she was wrong. Just as the truck driver may overdo his woe, Sally may also be overdoing her reaction. And just like the truck driver, if Sally showed undue penitence in her retraction (instead of light sorrow or grief), it would be appropriate for us to conclude that she was overreacting. Thus, we would conclude that Sally was blaming herself for something she was not responsible for, and we would try to dissuade her from doing so by pointing to her innocence. Responses like "It's O.K., it wasn't your fault," "You could not have known," "For all you knew, they might have been

there,” and so on, all seem completely in order. But if Sally’s utterance of (2) was not justified, the puzzle dissolves. It is no longer necessary to explain how her assertion (1) became unjustified, or what else Sally could be wrong about. This leaves no ground for the postulation of an assessment-sensitive proposition to account for her linguistic behavior. The content of her original utterance (1) could well be (4). Thus, semantic luck and relativism are both avoided.

In the end, all of Rosebury’s options remain open for whoever wants to challenge relativism by disputing one of its main motivations, that is, the existence of semantic luck. Burrowing from Rosebury’s interpretations of Williams’s example, I have developed several alternative interpretations of Sally’s puzzle, none of which requires the postulation of an assessment-sensitive proposition. I pretend neither to have shown that any one of them is the right interpretation nor to have shown that they are preferable to relativism. To do so would require broader considerations and detailed comparisons. All I wanted to do was to raise a challenge to the relativist interpretation of a particular semantic phenomenon. This challenge is not based on any general concern one might have about assessment sensitivity. It is based on the existence of alternative explanations of the phenomenon. As such, it weakens the motivation behind relativism, at least for the epistemic modality “might.”

*Instituto de Investigaciones Filosóficas
Circuito Mario de la Cueva s/n
Ciudad Universitaria, Coyoacán
Mexico DF 04000
abarcelo@filosoficas.unam.mx*

Acknowledgments

I would like to thank the audiences at the Third Annual University of Texas-UNAM Conference and the Jean Nicod “Communication, et Signification” workshop, Leonard Clapp, and an anonymous *Metaphilosophy* referee for their valuable comments on earlier drafts. I would also like to acknowledge the valuable support of UNAM PAPIIT IN401909, “Significado y Comunicación,” and CONACYT 083004, “Lenguaje y Cognición.”

References

- Barker, Chris. 2009. “Clarity and the Grammar of Skepticism.” *Mind and Language* 24, no. 3:253–73.
- Barker, Chris, and Gina Taranto. 2003. “The Paradox of Asserting Clarity.” In *Proceedings of the Western Conference in Linguistics*

- (*WECOL*) 2002, edited by Brian Agbayani, Paivi Koskinen, and Vida Samiian, 14:10–21. Fresno: California State University.
- Evans, Gareth. 1985. “Does Tense Logic Rest upon a Mistake?” In his *Collected Papers*, 343–63. Oxford: Clarendon Press.
- García-Carpintero, Manolo. 2008. “Relativism, Vagueness and ‘What Is Said.’” In *Relative Truth*, edited by Manolo García-Carpintero and Max Kölbel, 129–54. Oxford: Oxford University Press.
- King, Jeffrey. 2003. “Tense, Modality, and Semantic Value.” *Philosophical Perspectives* 17, no. 1:195–246.
- Latus, Andrew. 2001. “Moral Luck.” In *The Internet Encyclopedia of Philosophy*, edited by James Feiser. URL = <http://www.utm.edu/research/iepl/>.
- MacFarlane, John. 2005a. “Making Sense of Relative Truth.” *Proceedings of the Aristotelian Society* 105:321–39.
- . 2005b. “The Assessment Sensitivity of Knowledge Attributions.” In *Oxford Studies in Epistemology Volume 1*, edited by Tamar Szabo Gendler and John Hawthorne, 197–233. Oxford: Oxford University Press.
- . 2011. “Epistemic Modals Are Assessment-Sensitive.” In *Epistemic Modality*, edited by Brian Weatherson and Andy Egan, 144–78. Oxford: Oxford University Press.
- Nagel, Thomas. 1979. *Mortal Questions*. New York: Cambridge University Press.
- Nelkin, Dana. 2008. “Moral Luck.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. URL = <http://plato.stanford.edu/archives/fall2008/entries/moral-luck/>.
- Price, Huw. 1983. “Does ‘Probably’ Modify Sense?” *Australasian Journal of Philosophy* 61, no. 4 (December): 396–408.
- Rosebury, Brian. 1995. “Moral Responsibility and ‘Moral Luck.’” *Philosophical Review* 104, no. 4:499–524.
- Williams, Bernard. 1981. *Moral Luck*. New York: Cambridge University Press.