sciendo

Who Is Afraid of the Logical Problem in Meta-Ethics?

Ruth Weintraub
Tel-Aviv University

doi: 10.2478/disp-2022-0020 BIBLID: [0873-626X (2022) 67; pp. 411–26]

Abstract

Expressivism, as applied to a certain class of statements, evaluative ones, for instance, is constituted by two doctrines, only the first of which will concern me in this paper. Evaluative statements, according to this doctrine, aren't propositional (susceptible of truth or falsity). In this paper, I will argue that one of the vexing problems (that I label the "logical problem") this doctrine engenders for the expressivist is equally pressing for some cognitivists (who think evaluative statements *do* have a truth-value). I will present the difficulty and argue that some constructivists, who *are* cognitivists, cannot contend with it at all, and others must resort to more complex ways than the one available to other cognitivists.

Keywords

cognitivism; expressivism; Frege-Geach; meta-ethics;

1 Introduction

The meta-ethical "logical problem" involves a challenge for any meta-ethical theory: to account for the way we invoke moral statements in reasoning. Meeting the challenge is straightforward for some theories, and constitutes a very thorny, perhaps intractable, task for others. My aim in this paper is to show that the dividing line between the former theories and the latter does not lie where it is customarily thought to.

The structure of the paper is as follows. I begin with a brief typology of meta-ethical theories (section 2). I then explain more precisely the nature of the logical problem (section

3), and argue that contrary to prevalent opinion, it is not only expressivists that should be perturbed by it: it arises for some cognitivist theories, too, and they succeed in meeting it to different extents (sections 4, 5, 6).

2 A typology of meta-ethical theories

Expressivism about morality I take to be the doctrine that moral statements aren't propositional, susceptible of truth or falsity.[1] *Cognitivism*, its denial, is the doctrine according to which moral judgements *are* propositional, and it comes in (at least) three varieties. According to moral realism (Shafer-Landau [2003]), moral judgements are true or false in virtue of some moral facts, actions possessing moral properties. According to error-theorists (Mackie [1977]), although there are no *facts* in virtue of which moral judgements can be true, moral statements ascribe moral properties to actions, and we (erroneously) take some moral statements, literally construed, to be true. I call these two doctrines *conceptualist*, because they accept that moral terms denote *genuinely* moral concepts, concepts that purport to denote moral properties.[2] The metaphysical difference between these two doctrines, pertaining to whether or not moral concepts are (can be) *instantiated*, does not matter for the purpose of this paper.

The third cognitivist doctrine, which also comes in several varieties—I label, following Shaffer-Landau [2003], *constructivism*. Moral statements, it claims, *are* propositional, but they don't involve genuinely moral concepts,[3] because there aren't any. The "goodness" of an action is "constructed" from facts about what some speaker (possibly ideal) or society desires. The difference between constructivism and the first two versions of cognitivism

[1] I follow customary usage, from which there are divergences.

[2] For an obvious reason, the term 'conceptualist' seems to me apposite. But it may be misleading to those familiar with the discussion of *general terms* in the literature. Locke is there labelled a *conceptualist*, because he thinks there are abstract ideas (concepts): 'dog', 'red', for instance. In contrast, Hume, who is a *nominalist*, thinks general terms function by calling to mind particular instances. A thought about dogs, for instance, requires an idea of some particular instance, Fido, *e.g.*, to be present in the mind, together with a disposition to call to mind other representative ideas of particular dogs when "the occasion requires". Both Hume and Locke think statements involving general terms may be true, which the moral error-theorist and the expressivist deny with respect to moral statements.

[3] For the sake of brevity, I will henceforth drop the qualification 'genuinely', and talk about 'moral concepts', to refer to concepts that purport to denote moral properties.

will turn out to be very significant.

We can further distinguish between *actualist* and *idealising* versions of constructivism. The former (call it A-constructivism) takes goodness to be constructed from *actual* desires (individual or collective); the latter (call it I-constructivism)—from *idealised* desires (typically individual). The difference between the two versions of constructivism will also be seen to be important.

3 The logical problem

I use the label 'the logical problem' to denote one of three challenges facing every meta-ethical theory, which together constitute the *Frege-Geach problem* (as I understand it). The first challenge is to account for the very *fact* that we use sentences involving moral terms in reasoning. This is straightforward for cognitivists, who can invoke the standard account of validity, in terms of truth preservation. After all, they think moral sentences are truth-evaluable propositions, so arguments in which they figure may be logically valid. By contrast, non-cognitivists, who think moral statements are *not* truth-evaluable, must invoke a non-standard notion of "validity", which eschews invocation of truth and falsity.

The second component of the Frege-Geach challenge is the one most followers of Geach focus on: the need to provide a compositional semantics for sentences involving moral terms. Cognitivists can straightforwardly contend with this challenge, by invoking the standard truth-conditional semantics for sentences involving moral terms. For expressivists, the challenge is much more formidable. The meaning of a sentence, according to them, is the attitude it is used to express. Consider the two sentences 'Telling the truth is good' and 'Getting others to tell the truth is good'. They express a positive attitude towards telling the truth and getting others to tell the truth, respectively. The two sentences can be combined to engender the conditional 'If telling the truth is good, it is good to get others to tell the truth'. In the conditional, neither component is asserted. So the expressivist owes us an account of the way their meanings, *i.e.*, the attitudes they are used to express when asserted, determine the meaning of the compound sentence.[4]

The third challenge, the *logical* one, with which I am concerned in this paper, is to

[4] Geach thinks that the two components do not have the same meaning when asserted on their own and when embedded, and (plausibly) takes this to be a *reductio* of the expressivist construal of ethical statements, indeed, of expressivism in general. But expressivists have since shown that a compositional semantics *can* be given, thereby rebutting Geach's argument for the radical conclusion.

account for the *way* we use moral sentences in our reasoning: the *particular* arguments we affirm and the inferences we invoke. (These are different tasks, so the logical challenge has *two* facets, and some theories can succeed differentially in coping with them. This will become clear as my argument unfolds.) For the expressivist, who is debarred from using the standard notion of validity, the challenge is clearly formidable.[5] The prevalent thought is that *cognitivists* can be sanguine here. Thus, Schaffer-Landau [2003: 39] suggests that "because all forms of constructivism are forms of cognitivism, every constructivist theory has the advantage of being able to sidestep [some of] the difficulties that beset non-cognitivism".[6] But this thought is due to a conflation of two of the challenges that I discerned above: explaining the very *fact* that we endorse some inferences and explaining why we endorse the *particular* inferences that we do. To contend with the latter challenge, but not with the former, we need to explain, for instance, why we (normally) abide by *Modus Ponens* (henceforth, MP), but not with the rule that permits the inference of $p \wedge q$ from p. The latter problem, I will argue in the rest of the paper, is insoluble for some cognitivists. So cognitivism does not ensure success *vis-à-vis* the logical problem.

4 A-constructivism and the logical problem

Consider the following (well-worn) argument (Geach [1965]):

> If telling the truth is good, it is good to get others to tell the truth
> Telling the truth is good
>
> ---
>
> It is good to get others to tell the truth

A natural language argument is valid iff it is an instance of a valid schema, something that cannot just be read off its *grammatical form*. Whether or not Geach's argument is an instance of MP, and therefore valid (as we judge it to be), depends on its *logical form*, which depends, in its turn, on the meanings of its constituent sentences.

Conceptualists and A-constructivists interpret the argument so that it *is* an instance of

[5] The best-known solutions are Blackburn's [1984, 1985], Gibbard's [1990, 2003], Schroeder's [2008] and Silk's [2015].

[6] Shaffer-Landau goes on to point to *other* difficulties constructivism engenders.

MP, and therefore, valid. But the difference between the two kinds of theories does matter in the case of arguments whose validity depends on the *inner structure* of moral statements. Consider the following argument (call it ARGUMENT):

Killing for fun is forbidden
A is a killing for fun

A is forbidden

Conceptualists take the statement 'X is forbidden' to ascribe (possibly falsely) a property, 'is forbidden', to X. So ARGUMENT is straightforwardly valid for them, being an instance of the following valid schema:

$\forall(Px \rightarrow Qx)$
Pa

Qa

ARGUMENT would pose no problem for A-constructivists if they took the moral term 'forbidden' to mean 'green' or some other non-moral predicate. Its logical structure would be identical with its structure when it is conceptualistically construed. But, of course, they don't. And when we consider A-constructivist renditions, things are less auspicious.

Consider, first, Russell's [1952: 49] *speaker-subjectivism* (for brevity—subjectivism), according to which, the good is the satisfaction of the speaker's desire. 'X is forbidden' means, for Russell, 'I desire ¬X'. To see whether ARGUMENT is, for Russell's subjectivist, an instance of a valid schema, we need to know how its premises and conclusion are to be interpreted. In fact, we must consider two possibilities. On the first, *de re*, interpretation, the meaning of the first premise, that killing for fun is forbidden, is that for each action that is a killing for fun, I desire that *it* not be performed, and ARGUMENT is rendered thus:

For every act that is a killing for fun, I desire that it not be performed
A is a killing for fun

I desire that A not be performed

ARGUMENT is valid on this interpretation, but it is very implausible, because whereas *we* think the first premise of ARGUMENT (killing for fun is forbidden) is true, the interpretation renders it false, at least for humans. There are (at least) two reasons for thinking that in general, an open-ended generalisation can't be believed *de re*. (The restriction is required because it seems possible for me to believe *de re* that all the people in this room are tall, for instance.)

First, it is intuitively implausible that anyone has all the (singular) thoughts about all the relevant individuals (the killings for fun, in our case). The implausibility is exacerbated by the fact that some of these individuals will never exist. Nomological generalisations, 'There are no rain-storms without clouds', for instance, hold also for non-actual objects. The counterfactual 'Had there been a storm today, the sky would have been cloudy' is true. Closer to home, the generalisation 'Killing for fun is forbidden', call it KF, also applies to non-actual actions. Suppose I am deliberating whether to perform some action, and find out that it would be a killing for fun. Taking into account my belief KF, I decide not to perform the action. This shows that I (indeed, we all) take KF to apply to it—a (particular) non-actual action.

Perhaps we shouldn't set too much store by our intuitions: after all, the notion of thinking directly about (desiring) an object is highly abstract and likely to be theory-laden. So, arguing more theoretically, I will show that on none of the main views about the conditions required for a singular (*de re*) belief can the belief in KF be *de re*. They all assume, eminently plausibly, that to have a *de re* attitude towards a particular object (as opposed to one that picks out whichever object happens to satisfy its descriptive content), one needs to *pick it out*, so that the attitude is about *it* and not about something else. On these views, I will now show, this cannot happen in the case of KF.

Russell [1911] thinks *de re* thought requires direct acquaintance with the object at which the attitude is directed. Clearly, this condition is not satisfied by KF: we don't have direct acquaintance with most, if any, of the actions that are killings for fun, even actual ones. Bach [1987] claims that having a singular attitude involves standing in *some* kind of acquaintance towards the object: perception, memory or communication chains involving at least one person who has perceived the relevant object. Genone [2014] thinks it requires appropriate, though defeasible, evidence that the object exists, and Jeshion [2010: 136] thinks the subject matter of the singular attitude must be "significant to the agent with respect to her plans, projects, affective states, motivations". Because KF is open-ended, it doesn't satisfy any of these conditions.

On the most liberal view (Harman [1977], Kaplan [1989]), singular thoughts can be

generated simply by means of using a definite description, such as 'the oldest tree in the forest', as a device of direct reference. So even if one has no direct informational or perceptual acquaintance with some individual, one can think singular thoughts about it if one employs a definite description that it uniquely satisfies. KF doesn't pass even this (to many minds too lenient) test, because in the case of open-ended generalisations, there isn't (in our minds) for each individual in the domain a definite description that it uniquely satisfies.

So much for the *de re* interpretation of KF. The second, *de dicto*, interpretation construes the object of a desire, in analogy with that of *de dicto* belief, as a (possible) state of affairs. KF is true iff the speaker desires the obtaining of the state of affairs in which no killings for fun occur. This interpretation (of subjectivism), much more plausibly than the first, makes the first premise of ARGUMENT true. But it renders the argument *invalid*: the speaker may desire the state of affairs 'Killing for fun never occurs' without desiring the state 'This killing for fun does not occur'. Our desires, unlike the (psychological) facts about them, aren't *logically cogent*. They may contradict one another, and they are not deductively closed.

If this diagnosis is correct, we would expect the same kind of validity failure to occur in some *non-moral* contexts involving subjectivist construals. And indeed, it does. Subjectivism about non-moral judgements will engender the same failure of validity, because, like desires, our beliefs are logically imperfect. Thus, if the statements 'Dogs are brown' and 'Fido is a dog' meant 'I believe dogs are brown' and 'I believe Fido is a dog' respectively, the inference from 'Dogs are brown' and 'Fido is a dog' to 'Fido is brown' would be invalid. A person could believe the two premises without believing the conclusion.[7]

I conclude that the first interpretation of Russell's subjectivism is extremely implausible, and the second cannot contend with the logical problem, at least not in the ordinary way, available to conceptualists. (I will consider a different strategy in the next section.)

A similar diagnosis will show that *general subjectivism* cannot straightforwardly solve the logical problem. This doctrine has two versions. According to Wong's [2008] "collectivist" version, moral truths are determined by the *complete* agreements of some groups within specific traditions. Clearly, there is no reason to suppose that (complete) agreements will

[7] This kind of subjectivism must be restricted. If *every* descriptive statement, p, were to be construed as 'I believe p', the definition would engender a vicious infinite regress: the meaning of 'p' would be 'I believe p', which would mean, in its turn, 'I believe 'I believe p", and so on *ad infinitum*. But we needn't concern ourselves with the details of this (implausible) subjectivism. I am only using it to show that (undesirable) failures of validity may occur outside the moral realm.

engender logical cogency. Suppose everybody agrees about all moral judgements. Since no individual is logically cogent, neither is the collective: it inherits the logical imperfections of the individuals. So it may desire that killings for fun do not occur while failing to desire that some particular killing for fun not occur.[8,9]

According to the second, majoritarian, version of general subjectivism, 'A is good' means that A is desired by *most* people. It, too, engenders violations of logical cogency, because "majority-rule" desires coincide with "consensus-rule" desires if everyone has the same desires, which we have seen to violate logical cogency. And there are other desiderative profiles in the community for which the majority may desire that killings for fun do not occur while failing to desire that some particular killing for fun not occur.

I haven't (almost impossibly) surveyed all versions of A-constructivism, but such a survey is not required to establish the conclusion that several of them render ARGUMENT (and other arguments of its ilk) invalid, their cognitivist nature notwithstanding.

5 An A-constructivist strategy for contending with the problematic arguments?

The existence of conceptualistically valid arguments that are invalid according to constructivist interpretations does not, by itself, show that the A-constructivist cannot solve the logical problem. The conceptualist explains why we affirm the arguments that we do in terms of their soundness (for which logical validity is a necessary condition). But this is not the only explanation. A natural language argument doesn't wear its validity on its sleeve. When ARGUMENT is adduced in ordinary contexts, the words 'deductively valid' do not appear. So we cannot simply assume that we take it to be deductively valid. The

[8] Suppose *a* and *b*, but not *c*, desire that killings for fun not occur, and only *a* desires that A does not occur. The majority desires that killings for fun not occur, but does not desire that A not occur, thereby violating logical cogency.

[9] A reviewer pointed out that it is possible that most people F and that most people G without most people (F & G). But this is *not* an example of failure of logical cogency. It is in general not true that if A is good and B is good, A & B is good. Suppose (plausibly) that 'A is good' means that everything else being equal, A's occurrence is better than its absence; for every two worlds the only difference between which is that A occurs in one (and everything that the difference logically entails), that world is better than the other. Now, suppose I can buy an expensive life-saving medication in one of two places. It is good that I buy it in one of the two, but not good that I buy it in both (needlessly spending money I can ill afford).

incontestable phenomenon is our taking it (and others of its ilk) as *normatively binding* (Hale [1993: 338]). And the construal of this normative force in terms of deductive validity is going beyond the evident. The A-constructivist can claim that some of the moral arguments that we affirm are deductively *invalid* when interpreted (as they should be) A-constructivistically, and we affirm them because they are good *inductive* arguments: if the premises are true, the conclusion is *likely* to be true.

Consider, first, the application of the inductivist strategy to ARGUMENT considered subjectivistically:

I desire the state in which killings for fun do not occur
A is a killing for fun

I desire that A not be performed

The premises of the argument provide a good (albeit inconclusive) reason for the conclusion, at least when background assumptions are brought to bear. If a person desires that no killings for fun should occur, and believes that a particular action, A, is an instance of killing for fun, he will (typically) desire that A not be performed: desires *typically* respect simple relations of the kind ARGUMENT involves. So I can apply this (statistical) assumption to my own case, and infer the (probable) truth of the conclusion from the premises.

But even if ARGUMENT is a good inductive argument, its invocation is *pointless*. (And since we *do* invoke such arguments, this constitutes strong evidence against subjectivism.) An argument is useful only if its premises can be more easily known than its conclusion: it can then be invoked to gain new knowledge. But the truth-value of the conclusion of ARGUMENT ('I desire that A not be performed') seems epistemically much more accessible than the two premises ('For every act that is a killing for fun, I desire that it not be performed', 'A is a killing for fun'.) For conceptualists (and other constructivists), the epistemic situation is (quite often) the reverse: the premises are often more easily known than the conclusion. So these theories can better contend with this challenge (the second aspect of the logical challenge).

So much for subjectivism. According to general subjectivism, ARGUMENT may well have a point: its conclusion is sometimes less easily known that its premises. But it is not a good inductive argument. I defend this claim for the first (consensual) version, but it is easy to see that it holds for the majoritarian one as well.

Everyone desire the state in which killings for fun do not occur
A is a killing for fun

Everyone desires that A not be performed

The two premises constitute a very poor (if any) reason for the conclusion. What is required by way of strong support for the conclusion is that everyone *believe* the second premise; its truth is irrelevant.

I conclude that no A-constructivist theory can contend with both aspects of the logical problem: accounting for our endorsements of certain arguments (albeit in a somewhat surprising way) and for our invoking them.

6 I-constructivism

In this section, I will consider the logical challenge for constructivist theories that involve an *idealisation*, the good being determined by the desires of some ideal agent (or group). For Smith [1994], the judgement that it is desirable to do A in circumstances C is identical with the (non-moral) judgement that if I were perfectly rational, I would desire that I perform A in circumstances C. An agent must satisfy several conditions in order to be rational (Smith [1994: 156–61]).[10] I will consider them in turn, and show that although none of them engenders desiderative logical cogency, there is a fix, but it comes at a cost.

First, a person must have no false non-moral beliefs and all relevant true ones. Copp [1997: 44] objects that this is not the ordinary meaning of the term 'rational'. A rational person isn't required to be omniscient: if he correctly responds to misleading evidence, he will have justified false beliefs. But from our (logical) point of view, the objection can be conceded and shrugged aside. We will consider the view according to which 'A is good' means that were I rational$_{Smith}$, I would desire that I perform A in circumstances C. Smith is free to use the word 'rational' as he wishes, so long as his usage is clear. And because I am only interested in the logical problem, I need not concern myself with the plausibility of his proposal *qua* analysis of the word 'good'. (For ease of exposition, I will continue to use the term 'rational'.)

This requirement doesn't make for logical cogency. If an agent desires that actions

[10] Some Smith adopts from Williams [1981].

of type A be performed, then even if he recognises that some action is of type A (as he may not if he isn't "rational"), the requirement does not guarantee that he desire that it be performed.

Second, a rational agent deliberates in an unemotional way. But the desires of an unemotional agent may well be logically imperfect. Third, a rational person exercises his imagination in deliberating: he tries to imagine what the outcome of an action would be, and takes it into account. But, again, this doesn't make for desiderative logical cogency. Let us suppose that a rational agent can imagine "correctly": he can infallibly have (if he chooses) the phenomenology which will come about as a result of performing every possible action in any circumstance. (This makes it very implausible that we have a human being here, but that doesn't matter for our purpose.) Here is how such a rational person will deliberate regarding promises. He will imagine what keeping all promises and what failing to keep some promise feel like. Suppose he desires that all promises be kept. This means that keeping all promises feels better, *ceteris paribus*, than each of the alternatives: failing to keep some non-empty set of promises. So choosing by reference to the phenomenology, he will desire for each promise that it be kept. But this will make for logical cogency only if he cares solely about phenomenology. And this means that the constraint will engender logical cogency only for hedonistic agents, those whose rational selves would care only about phenomenology (and about promise keeping, for instance, only insofar as it makes for pleasant feelings).

Fourth, the rational agent must be *coherent*: he desires to do what he believes is desirable that he do. But there can be logically imperfect agents with a "coherent" desiderative profile. Consider someone who desires and believes to be desirable that all promises be kept, and neither desires that a particular promise, P, be kept nor believes that it is desirable that it be kept. His desires perfectly conform to his moral beliefs, but they are both logically defective.

Fifth, the desires of a rational person must be "systematically justifiable". They must form a "unified" whole. For instance, there must be more general desires to "explain" particular ones. For instance, if someone likes to eat chocolate, fruit and peanuts, his desiderative profile is more "unified" (and therefore, more systematically justifiable) if he

also likes to eat tasty food. This clearly doesn't make for logical cogency.[11]

Finally, a rational agent must abide by the means-ends principle. But an agent may desire that all promises be kept and take the means to bring this about without desiring that some particular promise be kept. Of course, he will have a true belief about the means (if any) for bringing about the keeping of this promise. But the keeping of this promise is not a means to the satisfaction of his desire that actions of type A be performed. It is a logically necessary condition for its satisfaction. This problem can easily be solved by requiring that the rational person also desires logically necessary conditions for the satisfaction of his desires.

We can seemingly declare Smith's proposal to be successful in contending with the logical problem. But in fact, we inferred the validity of ARGUMENT as Smith renders it from the fact that if its premises are true, so is the conclusion. And this inference is not legitimate in the case of Smith's proposal. Here is why.

Smith thinks [1994: 173] rational agents have the same desires (in the same circumstances). In particular, for a person with a (logically imperfect) desiderative profile, there is one desiderative profile he would have were he rational. I think he is wrong. It seems to me there may be more than one way to render a (logically imperfect) person's desires logically cogent. For instance, if I desire that all promises be kept but do not desire that promise P be kept, my desires will improve their logical cogency if I give up the first desire or come to possess the second. Sometimes there will be a uniquely superior modification: one that requires least change, for instance. But this won't always be the case. When it isn't, the desires of my rational self are not fully *determinate*. Supposing that if I were rational, I would desire that all promises be kept gives undue weight to one of the modifications. Supposing I wouldn't desire it unduly favours the other.

This does not mean that my "rationalised" desires can't be logically cogent.[12] But if they are, this will require a departure from classical logic, which caters only for two truth-values. To see this, suppose some agent desires that no killings for fun occur, but doesn't desire that a particular killing for fun does not occur (or even desires that it does occur). Suppose, further, that both ways of rendering his desires cogent (his rational self doesn't desire that

---

[11] Sayre-McCord [1997] objects that such "unity" doesn't make for greater "justification" in the case of desires, even if it is perfectly reasonable when it comes to beliefs. The objection is plausible, but irrelevant from the logical point of view. The only question *we* need to ask is whether satisfaction of the requirement makes for logical cogency.

[12] I am grateful to the editor for pointing this out to me.

no killings for fun occur, his rational self desires that a particular killing for fun does not occur) are tied: both are equally "near" his actual profile of desires. So while the second premise of ARGUMENT is true, the first premise and the conclusion are both neither true nor false. Now, in a two-valued logic, an argument is valid iff it preserves truth from premises to conclusion; if whenever the premises are true, so is the conclusion. But this is not an adequate characterisation of validity when there are, as Smith's proposal requires, *three* values: true, false, indeterminate, because it renders valid an argument whose premises are indeterminate and whose conclusion is false: the definition of validity is (vacuously) satisfied. And this seems wrong. So we need to assess ARGUMENT's validity by reference to a *non-classical* notion of validity, catering for arguments with components (premises or conclusion) that may be neither true nor false.[13,14] This means that the ability of Smith's proposal to contend with the logical challenge comes at a serious cost.

According to the second version of I-constructivism, Firth's [1952], 'A is good' means that an *ideal observer*, one who is impartial and perfectly informed about non-moral facts, would desire A. But plausibly, as Brandt suggests [1955], different ideal observers may desire very different things. There are here three suggestions that we need to consider by way of a response.[15] First, 'A is good' means that *some* ideal observer would approve of A. But this makes it possible for A to be good and not to be good, one ideal observer desiring it, another not desiring it. According to the second suggestion, 'A is good' means that *every* ideal observer would desire A. According to the third, it means that most ideal observers would desire A. The third suggestion makes for fewer indeterminacies than the second, but like Smith's proposal, both make for logical cogency only relative to a non-classical logic. (Unlike the analogous subjectivist suggestion, the majority rule doesn't engender violations of logical cogency, because each ideal observer is logically cogent. If most ideal observers desire that murder for fun not be committed, they will all desire that a particular murder for fun not be committed.)

---

[13] We may adopt a three-valued logic, as do some theories of vagueness. Williamson [1994: 101] suggests that in Körner's [1966] three-valued logic, the appropriate notion of validity is preservation of non-falsity. It is easy to see that it renders ARGUMENT valid.

[14] So while Smith's requirements straightforwardly comply with classical logic, the requirement that my rational self's desires are logically cogent, which we might suppose can simply be added to them, does not.

[15] I am grateful to the editor for suggesting them to me.

7 Conclusion

Both constructivists and expressivists deny the existence of moral concepts in our language. But because the former espouse a cognitivist theory about moral discourse, they seem to have two advantages over the latter. First, when construed cognitivistically, moral statements have standard truth-conditions, so constructivism is a much simpler semantic theory. Second, constructivism seems to have a straightforward solution to the logical problem, which is thought to arise when it is supposed, as it is by expressivists, that moral statements are not propositional. I have argued that not all constructivist theories partake of the second advantage, and those that do must suppose moral discourse is subject to a non-classical logic. I haven't considered the *overall* credibility of these constructivist theories, only their ability to contend with the logical problem. But of course, if they fail, that is a weighty reason for rejecting them.[16,17]

Ruth Weintraub
Tel-Aviv University
Ramat-Aviv 69978
Israel
weintrat@tauex.tau.ac.il

References

Austin, John Langshaw [1961]. "Other Minds". In *Philosophical Papers*. Oxford: Clarendon Press (1979).

Bach, Kent [1987]. *Thought and Reference*. Oxford: Oxford University Press.

---

[16] Strandberg [2015] argues, in a somewhat similar vein, that the embedding problem, usually taken to arise only for non-cognitivists, can be generalised to some cognitivist meta-ethical theories, for instance, the theory according to which the meaning of a moral statement is a mental state that is, unlike a belief, both representational and motivating. Our projects differ in two important respects. First, whereas he is concerned with the embedding problem (the second challenge I discern), I focus on the logical problem (the third challenge I discern). Second, Strandberg argues that the embedding challenge for some cognitivist theories is serious, whereas I claim that the logical challenge is *insurmountable* for some cognitivist theories.

[17] I am grateful to the anonymous referees, the editor, and David Kovacs for very helpful comments.

Baz, Avner [2010]. "Geach's 'Refutation' of Austin revisited". *Canadian Journal of Philosophy* **40** 41–62.

Blackburn, Simon [1984]. *Spreading the Word*. Oxford: Clarendon Press.

Blackburn, Simon [1985]. Errors and the phenomenology of value". In *Blackburn, S. Essays in Quasi-Realism*. Oxford: Oxford University Press (1993).

Brandt, Robert [1955]. "The definition of an 'ideal observer theory' in ethics". *Philosophy and Phenomenological Research* **15**: 407–13.

Copp, David [1997]. "Belief, reason and motivation: Michael Smith's *The Moral Problem*". *Ethics 108*: 33–54.

Dreier, Jamie [1990]. "Internalism and speaker relativism". *Ethics* **101**: 6–26.

Firth, Robert [1952]. "Ethical Absolutism and the Ideal Observer". *Philosophy and Phenomenological Research* **12**: 317–45.

Geach, Peter [1965]. "Assertion". *Philosophical Review* **74**: 449–65.

Genone, James [2014]. "Evidential constraints on singular thought". *Mind & Language* **29**: 1–25.

Gibbard, Allan [1990]. *Wise Choices, Apt Feelings*. Cambridge, Mass: Harvard University Press.

Gibbard, Allan [2003]. *Thinking how to Live*. Cambridge, Mass: Harvard University Press.

Hale, Bob [1993]. "Can there be a logic of attitudes?". In *Reality, Representation, and Projection*, edited by J. Haldane and C. Wright. Oxford: Oxford University Press.

Harman, Gilbert [1977]. "How to use propositions". *American Philosophical Quarterly* **14**: 173–6.

Jeshion, Robin [2010]. "Singular thought: Acquaintance, semantic instrumentalism and cognitivism". In *New Essays on Singular Thought*, edited by R. Jeshion. Oxford: Oxford University Press: 105–40.

Kaplan, David [1989]. "Demonstratives". In *Themes from Kaplan*, edited by J. Almog, J. Perry and H. Wettstein. New York: Oxford University Press: 481–564.

Körner, Stephen [1966]. *Experience and Theory*. London: Routledge.

Mackie, John Leslie [1977]. *Ethics*. Pelican Books, 2nd edition (1990).

Russell, Bertrand [1911]. "Knowledge by acquaintance and knowledge by description". *Proceedings of the Aristotelian Society* **11**: 108–28.

Russell, Bertrand [1952]. *Human Society in Ethics and Politics*. New York: Simon and Schuster.

Sayre-McCord, Geoffrey [1997]. "The metaethical problem". *Ethics* **108**: 55–83.

Schroeder, Mark [2008]. *Being For*. Oxford: Oxford University Press.

Shafer-Landau, Russ [2003]. *Moral Realism*. Oxford: Clarendon Press.

Silk, Alex [2015]. "How to be an ethical expressivist?". *Philosophy and Phenomenological Research* **91**: 47–81.

Smith, Michael [1994]. *The Moral Problem*. Oxford: Blackwell.

Strandberg, Caj [2015]. "Can the embedding problem be generalized?'. *Acta Analytica* **30**: 1–15.

Williams, Bernard [1981]. "Internal and external reasons". In *Moral Luck*. Cambridge: Cambridge University Press.

Williamson, Timothy [1994]. *Vagueness*. London: Routledge.

Wong, David [1984]. *Moral Relativity*. Berkeley, CA: University of California Press.