

The Fictionalist's Attitude Problem

Graham Oddie · Daniel Demetriou

Accepted: 26 June 2007 / Published online: 11 October 2007

© Springer Science + Business Media B.V. 2007

Abstract According to John Mackie, moral talk is representational (the realists go that bit right) but its metaphysical presuppositions are wildly implausible (the non-cognitivists got that bit right). This is the basis of Mackie's now famous error theory: that moral judgments are cognitively meaningful but systematically false. Of course, Mackie went on to recommend various substantive moral judgments, and, in the light of his error theory, that has seemed odd to a lot of folk. Richard Joyce has argued that Mackie's approach can be vindicated by a fictionalist account of moral discourse. And Mark Kalderon has argued that moral fictionalism is attractive quite independently of Mackie's error-theory. Kalderon argues that the Frege–Geach problem shows that we need moral propositions, but that a fictionalist can and should embrace propositional content together with a non-cognitivist account of acceptance of a moral proposition. Indeed, it is clear that any fictionalist is going to have to postulate more than one kind of acceptance attitude. We argue that this double-approach to acceptance generates a new problem – a descendent of Frege–Geach – which we call the acceptance–transfer problem. Although we develop the problem in the context of Kalderon's version of non-cognitivist fictionalism, we show that it is not the non-cognitivist aspect of Kalderon's account that generates the problem. A closely related problem surfaces for the more typical variants of fictionalism according to which accepting a moral proposition is believing some closely related non-moral proposition. Fictionalists of both stripes thus have an attitude problem.

Keywords Cognitivism · Error theory · Expressivism · Factualism · Fictionalism · Frege–Geach problem · Metaethics · Moral fictionalism · Moral realism · Moral acceptance

On the face of it, moral claims characteristically involve reference to various moral properties and relations: not only such thin moral properties and relations as permissibility and obligatoriness, goodness and badness, betterness and worseness, virtuousness and viciousness, but also thick properties such as being compassionate, cruel, selfish, kind, greedy, generous, honest, or wicked. To accept a moral claim is to believe a proposition ascribing such properties, thick or thin, to various entities – persons, acts, states of affairs, dispositions and so on. And to utter a moral claim is to express one's belief in the associated

G. Oddie (✉) · D. Demetriou

Department of Philosophy, University of Colorado at Boulder, Boulder, CO 80309-0275, USA
e-mail: oddie@colorado.edu

proposition. But thin moral properties have seemed problematic to many, for both epistemic and metaphysical reasons, and their thick offspring, to the extent that they implicate the thin properties, inherit those problems. If moral properties and propositions are problematic, how can one rescue moral talk without being burdened with a problematic ontology? For much of the twentieth century the favored solution was a bold and striking denial of the accuracy of the fundamental appearances. If there are no genuine moral properties, thick or thin, the meaningfulness of moral discourse does not involve reference to such properties. To make a moral claim is not to express one's belief in a moral proposition, but to express some other attitude – approval, desire, or some other non-cognitive attitude. Finally, the meaningfulness of a moral claim is somehow to be cashed out in terms of those attitudes. That is to say, non-cognitivism became the ruling paradigm.

John Mackie broke out of the non-cognitivist paradigm (Mackie 1977). According to Mackie, the realists are right that moral claims should be taken at their face value: moral predicates really do refer, or attempt to refer, to thick and thin moral properties; moral claims involve truth-evaluable propositions about such properties; and moral utterances express beliefs in moral propositions. Mackie thought the non-cognitivists were just wrong about all that. What the non-cognitivists got right, according to Mackie, was the extraordinary implausibility of the metaphysics which moral talk presupposes, and hence of the truth of moral claims taken at face value. So, Mackie put forward his now famous error-theory of moral discourse: moral talk is cognitively meaningful, representational talk all right, but it is at worst wildly false, and at best wildly unbelievable.

In his seminal book on the subject, *Ethics: Inventing Right and Wrong*, after arguing for this error theory Mackie went on to formulate, and apparently endorse, various substantive moral claims. This seems odd. If you think moral claims are just plain false, and radically misleading, because the basic properties and relations they presuppose do not exist, then presumably you have good reason not to believe, utter, or endorse them. Even if you might have a reason to appear to endorse them at times – say, while playing backgammon – you shouldn't be overtly endorsing them in the context of a philosophical treatise in which you are arguing for their radical defectiveness on the score of truth. Recently Richard Joyce has suggested that Mackie can be rescued from this awkward predicament by an explicitly *fictionalist* account of moral practice (Joyce 2001).

The fictionalist about a certain discourse acknowledges that, although the characteristic claims of the discourse have propositional content, those claims are false (or likely to be false, or truth-valueless, or rationally unjustifiable), but thinks there is nevertheless a good reason to go on using the discourse, not only making the characteristic utterances, but in some sense *accepting* them. Morality is, strictly speaking, an elaborate fiction, but it is a valuable fiction, in the sense that it is worth our while going about our business as though it were true. Well, perhaps not *valuable*, at least not in any realist sense. But useful, or helpful, or something like that. So, an error theorist about a discourse who thinks that there is reason to continue engaging in the discourse should go fictionalist.

Notice that this fictionalist response to Mackie's error theory requires a double-attitude approach to the acceptance of propositions. On the one hand, there is the kind of acceptance of a proposition that is simply belief in that very proposition. There are all manner of propositions that we have reason to believe. If Mackie is right then all the substantive moral claims – that torture is wrong, that abortion is permissible, that cruelty is bad, that courage is a virtue, and so on – are either false, or at least unbelievable. We have no good reason to believe them. But, despite this, we should carry on somehow as though we accept them. There are two possibilities if acceptance is somehow a matter of belief. One is that in accepting a moral proposition one straightforwardly believes some other, related proposition, perhaps the

proposition that *in the moral fiction the moral proposition in question is true*. So to accept that *lying is wrong* is to believe that *in, or according to, the moral fiction it is wrong to lie*. The other is that in accepting a moral proposition one adopts a belief-like attitude to the proposition, perhaps the attitude of *make-belief*. To accept that lying is wrong, for example, is not to believe it is wrong, but to make-believe that it is wrong.

It has to be admitted that *make-belief* sounds a bit feeble for the kind of attitude which morality, on the surface, demands. Make-belief is, after all, a highly overridable attitude. Whenever we make-believe something we are primed to abandon the attitude if reality intrudes in a rude or demanding way, and it is entirely appropriate to be so primed. Suppose make-belief is the appropriate attitude of an audience to the propositions explicitly or implicitly true in a work of fiction – like a play one is attending. Suppose it is true in the play that two people are chatting comfortably on a couch in their home, and that their home is not on fire. As we get into the play, we make-believe that that is true. Now, if smoke starts seeping onto the stage from backstage but it is clearly true, in the play, that there is no smoke in the room, we tend in such circumstances to abandon the make-belief (that there is no smoke in the room) and go with the belief (that there is smoke), and it is entirely reasonable to do so. When push comes to practical shove, make-belief will rightly give way to genuine belief. Suppose I am a moral error theorist, but also a fictionalist. So I make-believe that it is wrong to lie, and generally act as though that were true, conducting myself accordingly, avoiding both large and small lies. By analogy with the fire in the theatre, however, when the consequences of acting as though I really believe that lying is wrong compete with rude self-interest, then why wouldn't I have more reason to be guided by my actual beliefs about lying (i.e., it is not wrong to lie).

The other possibility is for the moral fictionalist to equate acceptance of a moral proposition with belief in some related proposition – an unexceptional non-moral proposition associated with that moral proposition. If morality is an elaborate fiction, then at least it is a fiction according to which, or *in* which, various moral propositions are true. Suppose it is true, in the moral fiction, that torture is morally wrong. To accept that torture is morally wrong is thus to believe the proposition that *according to the moral fiction, torture is morally wrong*. Again, however, this seems too feeble for genuine full-blooded moral acceptance. An error theorist who is also a moral nihilist might well believe that *according to the moral fiction lying is morally wrong* but he does not thereby *accept* that lying is morally wrong. He doesn't accept that anything at all is morally wrong.

Double-attitude fictionalism need not hold that the appropriate attitude to moral propositions is either belief or belief-like. Mark Kalderon, in his recent defense of moral fictionalism, explicitly argues for a double-attitude account (Kalderon 2005a, b). He argues both that the acceptance of a moral proposition is quite different from belief, and also that fictionalism need not be grounded in error-theory. Briefly, Kalderon thinks that the realists and error-theorists are right about the semantics of moral talk, while the non-cognitivists are right about acceptance and utterance of moral claims. Kalderon argues that two different issues have traditionally been systematically conflated within the debate about cognitivism versus non-cognitivism. One issue is this:

- 1 Do the target sentences of a certain discourse have genuine propositional content; are they truth-apt, do they have truth conditions?

The other is this:

- 2 Is acceptance of a sentence in that particular discourse wholly a matter of belief?

Traditionally cognitivism has been thought of as the affirmation of both of these, and non-cognitivism, the denial of both. Assuming that the object of a belief has to be a proposition, if the answer to the second question is YES then the answer to the first question should also be YES. If the acceptance of a moral claim, for example, is a matter of belief, then moral sentences must have truth-apt propositional content. To believe P is to believe that P is true. But the two questions can be decoupled, because answering YES to the *first* question is quite compatible with answering either YES or NO to the second. You might hold that moral claims, say, have propositional content, but go on to deny that acceptance of a moral claim is a matter of belief. There may be contexts in which acceptance of an utterance is not a matter of belief. And that is precisely what Kalderon does say.

A positive answer to the first question Kalderon calls *factualism*, a negative answer *non-factualism*. (This is not perhaps the best terminology: it sounds as though a factualist about morality, say, accepts that there are some moral facts, but this is clearly not Kalderon's intention. So it might be better to call a positive answer to the first question *propositionalism*, and a negative answer *non-propositionalism*.) In any case, Kalderon calls it factualism, and asks us to reserve the term *cognitivism* for the positive answer to the second question, and *non-cognitivism* for the negative answer. So a non-factualist is *ipso facto* a non-cognitivist, but a factualist need not be a cognitivist. The possibility of accepting a proposition without believing it provides logical space for non-cognitivist *factualism* – and that, Kalderon contends, is what fictionalism (or at least the best version of fictionalism) amounts to.

A non-factualist is under some pressure to explain what the content of a moral utterance amounts to, if it does not pick out a moral proposition. Some version or other of *expressivism* has been the standard response. That is to say: “the content of a moral sentence consists wholly or partly in the non-cognitive attitudes conveyed by its utterance.” (Kalderon 2005a, p. 53). One of the standard objections to expressivism is, of course, the Frege–Geach problem. Suppose you think that a moral or value claim does not come along with a truth-evaluable proposition. When you utter:

Torture is morally wrong.

you are not attributing a moral property – the property of moral wrongness – to an action type, torturing. Rather you are expressing, say, your disapproval of torturing. The trouble with this is that that sentence can be embedded in other claims such as:

It is simply false that torture is morally wrong.

If it is morally wrong to torture then condoning others torturing is also morally wrong.

If torture is morally wrong then it is morally wrong to condone the torture of suspected terrorists.

Now, in giving utterance to these claims you evidently do not express your disapproval of torture, since each of them is compatible with the rejection of the proposition that torture is morally wrong. Furthermore, moral argumentation makes extensive use of such embedded uses of the original moral sentence. For example:

Torture is morally wrong.

Waterboarding is a form of torture.

If some action is morally wrong then it is morally wrong to condone someone's acting in that way.

Cheney condoned the waterboarding of suspected terrorists.

therefore

Cheney did something morally wrong.

Or consider this argument:

If torture is morally wrong then it is morally wrong to torture a suspected terrorist who may have information about the imminent explosion of a dirty bomb.

It is morally permissible to torture a suspected terrorist who may have information about the imminent explosion of a dirty bomb.

therefore

It is false that torture is morally wrong.

Unless *torture is morally wrong* has a truth-evaluable propositional content it is hard to see how these arguments can be classified as valid. Validity is a matter of the guaranteed transfer of truth and falsity. Validity guarantees a downward transfer of truth. Truth flows from the premises of a valid argument down to the conclusion. For the same reason validity guarantees an upward transfer of falsity. Falsity seeps (rather than flows) up from the conclusion to the premises – not all the premises of a valid argument with a false conclusion can be true. As Kalderon makes clear, the apparent validity of these arguments presents the non-cognitivist with a dilemma (Kalderon 2005a, p. 59).

Either:

The meaning of an utterance of ‘Torture is morally wrong’ is not the same in both embedded and unembedded contexts (and in that case we have a fallacy of equivocation in all arguments involving a moral sentence which is both embedded and unembedded).

Or:

The meaning of the utterance ‘Torture is morally wrong’ is the same in both embedded and unembedded contexts (and in that case the meaning is not given by any expressivist semantics).

This dilemma captures a version of the famous Frege–Geach problem – we will call it the Frege–Geach *truth-transfer* problem. As Kalderon rightly points out, it is a problem for non-factualism, and specifically for expressivism, not for non-cognitivism as he characterizes it. For provided one rejects non-factualism and unabashedly embraces propositional content for moral sentences, then even if one is a non-cognitivist about acceptance, one can happily deny the first horn and embrace the second.

The good arguments for traditional non-cognitivism are arguments to the effect that moral acceptance is not wholly or even partly a matter of belief in a moral proposition. (As the internalists have maintained, it is plausible that there is at least something more than belief involved in accepting a moral judgment.) They are not arguments that moral acceptance cannot be some attitude (other than belief) to a moral *proposition*. Kalderon argues that if acceptance of a moral proposition is not belief in that proposition (or in some other related proposition), but rather some non-cognitive attitude to that proposition, then we can combine the good features of non-cognitivism while avoiding the bad features of a

problematic non-factualist semantics. The Frege–Geach truth-transfer problem, in particular, is not a problem for the non-cognitivist factualist.

Kalderon has an interesting if somewhat controversial view of what acceptance of a moral proposition amounts to: to accept a moral proposition is to *decide* that things are going to appear to one, affectively, in a phenomenologically vivid way – presumably just the way they would appear if that moral proposition were true and you were the sort of being who experienced moral states of affairs appropriately (see, for example, Kalderon 2005a, p. 147). Well, whatever acceptance of a moral proposition is, the important thing for a non-cognitivist like Kalderon is that such acceptance be fundamentally different from belief, which is the attitude of acceptance appropriate to a non-moral proposition. So, following Kalderon, let acceptance be, in effect, a *determinable*, which embraces quite distinct determinate realizing attitudes. Different kinds, or determinates, of acceptance are appropriate to different kinds of propositions. In the case of a non-moral proposition acceptance is the attitude of belief, and in the case of a moral proposition it is some other attitude, involving an affective or conative component – let’s label it *endorsement*. (Analogously, rejection of a proposition is a determinable of which disbelief is the cognitive determinate. We can label the non-cognitive determinate of rejection appropriate to a moral proposition *repudiation*.)

So we have:

The two-attitude account of acceptance

Endorsement/repudiation and belief/disbelief are distinct determinates of the determinable, acceptance/rejection. If P is a non-moral proposition, to accept P is to believe P (and to reject P is to disbelieve P); if P is a moral proposition, to accept P is to endorse P (and to reject P is to repudiate P).

Belief and endorsement are thus two fundamentally different attitudes – one cognitive, the other non-cognitive – appropriate to these two quite different types of propositions.

Consider the simplest possible argument that involves both moral and non-moral propositions. Consider the following trivial inference:

- C *Cheney condoned torture.*
 T *Condoning torture is morally wrong.*
 therefore
 W *In condoning torture, Cheney did something morally wrong.*

The argument is surely valid. C is (let’s say) a non-moral proposition. Perhaps it is shorthand for some purely non-moral report such as: *Cheney said it was a ‘no-brainer’ to torture a suspected terrorist by waterboarding.* T is clearly a moral proposition, attributing the property of moral wrongness to the condoning of torture. The conclusion, we may suppose, is simply tantamount to the conjunction of the premises. (If you are skeptical of that then you can substitute C&T for the conclusion as stated.) If we embrace non-factualism then we have no problem with the validity – the guarantee of a downward transfer of truth – of this argument.

But what exactly is the *point* of offering or uttering a valid argument? Typically the point of arguing is to get someone (perhaps oneself) who *accepts* the premises to *accept* the conclusion, (or alternatively, to get someone who rejects the conclusion to reject the conjunction of the premises). C, the first premise of our argument is, we stipulated, a

purely descriptive, non-moral proposition. Assume that you believe C. You accept that Cheney condoned torture. Suppose you also endorse the moral proposition T. You accept that condoning torture is morally wrong. So, you accept both the premises, and the conclusion follows from the premises, because it is simply logically equivalent to their conjunction. Since the argument is valid, the conclusion is a demonstrated consequence of the premises, so...surely you are rationally obliged to accept the conclusion. Why? Because acceptance, like truth, should transfer from premises to conclusion of a demonstrably valid argument.

In the case where we have non-moral propositions, transfer of acceptance is unproblematic, because presumably we have a principle of transfer for belief. As a first stab at this, let's take a formulation that is often cited in elementary logic courses.

Naive transfer principle for belief/disbelief

A rational person who believes (i.e. accepts) some propositions is rationally obliged to believe (accept) any proposition which is a demonstrated logical consequence of those propositions.

A rational person who disbelieves (rejects) some proposition is rationally obliged to disbelieve (reject) the conjunction of any propositions of which it is a demonstrated logical consequence.

Harman (2002) has criticized principles like this one for failing to take into account the distinction between principles of logic, which describe logical relations between propositions, and the principles of reasoning, which tell us how to proceed epistemically. Harman argues that the fact that some of your beliefs can be shown to entail some proposition does not mean that you are irrational to disbelieve that proposition (Harman 2002, pp. 173–174.) Of course, you might be just as rational to disbelieve the conjunction of the premises instead, and the transfer principle articulated above embraces that in its second clause. However even if you continue to believe the premises, and acknowledge the entailment, you may still not be *obliged* to adopt the conclusion as an additional belief. It may be just too costly, in some sense, to add that proposition to your stock of beliefs. For example, each of us has limited epistemic storage resources, and one may not be obliged to clutter up one's belief-box with the trivial or uninteresting consequences of other elements in the belief-box. Or, if you find the conclusion unacceptable, you may find it too difficult either to dislodge the antecedent beliefs or to accept their implication. You might have to forget about it in the meantime while you pursue some other task.

Our little argument above illustrates Harman's point. Clearly this argument would never be used for the purposes of exerting rational pressure on someone to accept the conclusion on the basis of their acceptance of the premises. Since the conclusion of this argument is the conjunction of the two premises, it is just too obvious. But we have a reason for keeping our example simple, and the simplicity of this particular argument does not undermine the following observation: that one of the main points of classifying arguments as valid or invalid is to enable us to reason soundly, and one of the main points of reasoning, at least in the theoretical realm, is to improve or maintain the rationality of what one accepts. Valid arguments can exert rational pressure on us as regards which propositions to accept and which to reject, and they do so by channeling reasoned acceptance down from the premises of a valid argument to their conclusion, or by channeling reasoned rejection from the conclusion of a valid argument up to the premise set.

Another kind of objection to naive belief transfer trades on the familiar lottery paradox. There are a million tickets in the lottery, and as it happens you believe each of the following propositions:

Ticket #1 will not win the lottery.

Ticket #2 will not win the lottery.

...

Ticket #1,000,000 will not win the lottery.

It follows from these propositions that *no ticket will win the lottery*. But you are certainly not rationally obliged to believe *that*, even if you do believe each and every one of the premises.

Finally, it does seem odd to say that someone who happens to believe some outlandish or highly improbable propositions is thereby rationally obliged to believe some even more outlandish and improbable consequence of their conjunction.

What this last objection suggests is that what is at issue is not so much the transfer of brute acceptance, as the transfer of *good reasons* for acceptance. If you have good reason to believe some propositions, and some other proposition is a demonstrable logical consequence of them, then you have a good reason to believe that consequence.

We can state a more nuanced Transfer Principle, to make it compatible with these criticisms. We need to soften the rational pressure involved.

Transfer principle for belief/disbelief

A rational person who has a good reason to believe some propositions has a good reason to believe any proposition that is a demonstrated logical consequence of those propositions.

A rational person who has a good reason to disbelieve some proposition has a good reason to disbelieve the conjunction of any propositions of which it is a demonstrated logical consequence.

Does this principle deal with the objections to the naive principle? First, it deals with Harman's objections. Having a good reason to accept a proposition does not entail that one is rationally obliged to accept it. There might be other reasons in the offing. Second, clearly it blocks the rational transfer of acceptance from crazy beliefs to their consequences – unless you have some good reason to believe the premises, the transfer principle alone does not guarantee you have a good reason to believe the conclusion. Third, the second half of the principle applies straightforwardly to the lottery paradox even if the first does not. Anyone who rejects the proposition that no ticket will win the lottery has a good reason to reject the *conjunction* of the premises of the paradox – and that seems right.

In any case, assume that this or some suitably refined transfer principle for belief links logic and argument to good reasons for acceptance and rejection of propositions. There must be some such principle and it must work something like this.

Now consider the case where we have a bunch of purely moral propositions, and acceptance is a matter of (non-cognitive) endorsement rather than belief. For example, suppose you endorse the proposition that torture is wrong, and you also endorse the proposition that condoning someone else's wrongdoing is itself wrong. Then it seems that you have a good reason to endorse the consequential proposition that condoning someone else's torturing is wrong. So, what we apparently require here is a parallel principle of transfer for the non-cognitive attitudes of endorsement and repudiation.

Transfer principle for endorsement/repudiation

A rational person who has good reason to endorse some propositions has a good reason to endorse any proposition that is a demonstrated logical consequence of those propositions.

A rational person who good reason to repudiate some proposition has a good reason to repudiate the conjunction of any propositions of which it is a demonstrated logical consequence.

At first blush this transfer principle looks as sound as the parallel principle for belief. But on reflection, it is not at all obvious why a double-attitude fictionalist about morality is entitled to it. In the case of belief, truth-seeking is built in. We aim at having true beliefs, and a necessary condition for the truth of our beliefs is their joint consistency. If our beliefs are inconsistent and we know this then we also know they cannot all be true. So if we come across an apparently unbelievable consequence C of some our beliefs, the goal of truth provides us with a reason either to believe C after all, or else to reject the conjunction of those antecedent beliefs. One cannot maintain consistency in one's beliefs if one continues to believe the premises but believes the negation of the conclusion. But why, if you are a fictionalist about some domain, should you aim at consistency of what you accept? As a fictionalist about morality, for example, you are not concerned that the moral propositions you endorse be true. So it is not obvious that the consistency of the moral propositions one endorses is required. And if it is not required, then we lack the corresponding argument for endorsement transfer.

One could argue that the propositions we endorse are intended to be action guiding, and a set of endorsements that are logically incompatible might well imply contradictory injunctions, and that would leave us without any guidance how to act. Hence the desirability of consistency follows from the action-guiding nature of our endorsements. But of course, this presupposes that if our endorsements can be shown to entail some moral injunction, then we have a reason to endorse that moral injunction. Absent endorsement transfer, it is not clear that endorsements that jointly entail some contradiction will actually yield confusing injunctions.

In any case, let us grant to the double-attitude fictionalist this endorsement-transfer principle in addition to the belief-transfer principle. Suppose we have these two transfer principles for the two determinates of acceptance. Does that guarantee what we really need: namely, a transfer principle for acceptance itself? This, or something like it, is the principle the double-attitude fictionalist has to end up with:

Generalized transfer principle for acceptance/rejection

A rational person who has good reason to accept some propositions has a good reason to accept a demonstrated logical consequence of those propositions.

A rational person who has good reason to reject some proposition has a good reason to reject the conjunction of propositions of which it is a demonstrated logical consequence.

For the factualist cognitivist about morality, the belief-transfer principle guarantees a generalized acceptance-transfer principle, whether the propositions involved are moral or non-moral. But for the non-cognitivist, it is not clear that the conjunction of the two specific principles guarantees generalized acceptance-transfer. Consider our little argument

concerning torture. Does either of our two initial transfer principles apply given the double-attitude thesis? Belief transfer applies when the premises are all believed. Endorsement transfer applies when all the premises are endorsed. The first premise is believed but not endorsed, while the second premise is endorsed but not believed. So neither principle taken by itself can give one a good reason to move from rational acceptance of the premises to rational acceptance of the conclusion. We need some additional transfer principles.

Note that this is not the standard Frege–Geach problem, which is a problem for the transfer of truth and falsity. Rather it is a Frege–Geach-like problem – the problem of transfer of rational acceptance and rejection.

Why not just *postulate* a generalized principle of acceptance/rejection transfer in addition to belief and endorsement transfer? Even disregarding the rather ad hoc nature of such a postulation, we still face a problem: what does it take to *accept* the conclusion of the little argument? The conclusion, since it is tantamount to the conjunction of two perfectly good propositions which are logically compatible, is itself a proposition – but is it either a fit object of belief, or a fit object of endorsement, or is it rather a fit object of some other attitude? Remember that belief and endorsement are, according to the non-cognitivist, quite distinct attitudes – one of them a cognitive attitude appropriate to a non-moral proposition, the other a non-cognitive attitude appropriate to a moral proposition. So the attitude appropriate to acceptance of the conclusion will depend on what kind of proposition the conclusion is. Is it a moral proposition or is it a non-moral proposition, or is it something else? Since it is equivalent to the conjunction of two propositions, one moral the other non-moral, it is not clear what it is. Call any proposition that is equivalent to the conjunction of some moral and non-moral propositions (like W and the logically equivalent proposition S&T) a *fusion*. What does acceptance of a fusion amount to?

There are three possibilities. Acceptance of a fusion is either a matter of belief, or a matter of endorsement, or it is some other attitude. If it is some third attitude, that attitude might be a combination of endorsement and belief (or, more accurately, it may supervene on endorsement and belief), or else it may be a distinct determinate of acceptance that does not supervene on endorsement and belief.

Suppose, first, that acceptance of a fusion is a matter of endorsement. Then, anyone who has a good reason to accept C and T would have a good reason to accept W. So (given that acceptance of a fusion = endorsement) such a person would have a good reason to endorse W. But W entails C. So, by the transfer principle for endorsement, upon grasping a very simple demonstration, she would have a good reason to *endorse a non-moral proposition* (C). But if it is not entirely clear what the non-cognitive attitude of endorsement of a *moral* proposition is, it is even less clear what it would be to take that attitude to a *non-moral* proposition like *Cheney condoned torture*. Recall that this is a non-moral proposition, a report of Cheney’s saying in an interview that it was a ‘no brainer’ to waterboard a suspected terrorist. Can one really *endorse* that non-moral proposition (*viz.* non-cognitively) in exactly the way in which one – non-cognitively – endorses the proposition that *torture is morally wrong*?

Of course, one can certainly *accept* (i.e. believe) that *Cheney condoned torture* without – in the colloquial sense – *endorsing* Cheney’s condoning of torture. For, in the colloquial sense, to endorse Cheney’s condoning of torture is to deem it morally acceptable, or even morally obligatory. In our technical sense, that would be tantamount to accepting (i.e. endorsing), not the proposition that *Cheney condoned torture*, but the proposition that *it is morally permissible (or obligatory) for Cheney to condone torture*. But these are, of course, two logically distinct acceptances. In general, it is hard to see how one could endorse a purely non-moral proposition in exactly the same way that the non-cognitivist thinks one endorses a moral proposition, and at the very least the non-cognitivist would owe us an account of that.

A corollary of this observation is that the non-cognitivist is in trouble accepting endorsement transfer quite independently of the problem of the correct acceptance attitude for fusions. Suppose one endorses the proposition that *torture is morally wrong*. Then, by endorsement transfer, one has a good reason to endorse the tautology. And the tautology is, presumably, a non-moral proposition. There is a good reason to accept the tautology (it is, after all, necessarily true) but that only applies if acceptance is belief. If endorsement transfer is in force, and one accepts any moral proposition at all, then one has a good reason to endorse the tautology as well. So one would have a good reason to believe the tautology, but also, if one endorses any moral proposition at all, one also has a good reason to endorse the tautology. This would constitute an odd case where two distinct determinates (belief and endorsement) of the one determinable (acceptance) apply to one and the same entity. Perhaps the non-cognitivist could solve this particular anomaly by restricting the domain of the moral to non-tautologous propositions. That is not totally implausible. Even so, we are left with the problem that for any given proposition, one might still bear any one of four pairs of acceptance/rejection attitudes to it.

Suppose, then, that acceptance of a fusion is a matter of simple belief. Anyone who accepts (believes) C and accepts (endorses) T would (by the generalized transfer principle) have good reason to accept W, and so would have a good reason to believe W. Then, by the transfer principle for belief alone, upon grasping a very simple demonstration, she would have a good reason to believe T: that torture is morally wrong. But the moral fictionalist clearly does not want one's acceptance of a moral proposition together with acceptance of some non-moral propositions to endow one with a good reason to believe that same moral proposition. The whole point of non-cognitivism is to avoid having to say that our acceptance of moral propositions creates rational pressure to believe those moral propositions.

This result can be generalized – there is nothing special about S, T and W. Given the two particular transfer principles and the generalized transfer principle, if acceptance of a fusion is a matter of belief, then a person will have good reason to believe all the moral propositions she endorses. And if acceptance of a fusion is endorsement, then a person will have a good to endorse all the propositions she believes. (Suppose you have a good reason to endorse moral proposition M. Suppose you also have a good reason to believe some non-moral proposition N. Then you have a good reason to accept both N and M, and so by the generalized transfer principle for acceptance, you have a good reason to accept N&M. If acceptance of the fusion N&M is belief, then you have a good reason to believe N&M. By the transfer principle for belief, you have a good reason to believe the moral proposition M. If acceptance of N&M is endorsement, then by the transfer principle for endorsement you have a good reason to endorse the non-moral proposition N.) So a rational being with a minimal grasp of conjunction-elimination and conjunction-introduction would have good reason to believe all and only the propositions she has good reason to endorse. Clearly this is not a result that would please the non-cognitivist, or indeed any two-attitude theorist – it obviously contradicts the thesis that there are two quite distinct attitudes appropriate to the two distinct kinds of propositions at issue.

To avoid this acceptance-transfer problem, the fictionalist is going to have to extend the notion of acceptance to fusions without making them the direct objects of either belief or endorsement. This is actually not too hard. A fusion P can be divided into two parts. The non-moral content of P, N(P), is the strongest non-moral proposition following from P. The moral content of P, M(P), is the strongest moral proposition following from P. And since P is a fusion, P is equivalent to N(P)&M(P). (That's what a fusion is defined as.) Now we can characterize an extended notion of acceptance for fusions which is itself a mixture of belief and endorsement.

Acceptance of fusions

A person accepts a fusion P if and only if she believes N(P) and endorses M(P).

Notice that we have two limiting cases. Where P is equivalent to its natural component N(P) (that is to say, P is *wholly* natural), acceptance of P will simply amount to acceptance of its natural component, N(P), and since N(P) is a non-moral proposition, acceptance of P will be a matter of belief that P. Analogously, in the case where P is equivalent to its moral component, M(P) (that is to say, where P is a wholly moral proposition) acceptance will just be endorsement. On this account, there is no problem with accepting W: that *in condoning torture Cheney did something wrong*. The non-moral content of W is C. The moral content is T. So you accept W just in case you believe C and endorse T. So far so good.

Unfortunately this move is not sufficient to solve the acceptance transfer problem. Notice that accepting a fusion is neither belief nor endorsement. It is a third attitude, distinct from the other two. Maybe we should call acceptance of a fusion, *beldorsement*, a *third* determinate of the determinable of acceptance. The two-attitude account has thus become a three-attitude account. Since beldorsement supervenes on belief and endorsement this is not such a huge cost if it is any cost at all. More importantly, not all propositions that are neither purely moral nor purely non-moral are fusions. All propositions have natural and moral components, but some propositions are not equivalent to the conjunction of their moral and non-moral components. These are what we might call *organic hybrids* – propositions that have moral and non-moral components but are not the conjunction of those two components.

Consider the proposition $C \Leftrightarrow T$: *Cheney condoned torture if and only if condoning torture is morally wrong*. What is the moral content of $C \Leftrightarrow T$? $C \Leftrightarrow T$ does not entail either C or $\sim C$. So it doesn't entail this purely non-moral proposition or indeed any other non-tautologous non-moral proposition. Nor does it entail either T or $\sim T$, or any other non-tautologous purely moral proposition. Both its moral and non-moral contents are thus tautologous. But $C \Leftrightarrow T$ itself is not a tautology. So $C \Leftrightarrow T$ is not equivalent to the conjunction of its moral and non-moral components. It is not a fusion but a *hybrid*, and so we clearly cannot cash out acceptance of $C \Leftrightarrow T$ in terms of endorsement of its moral content and belief in its non-moral content.

A factualist can think of propositions as picking out classes of worlds, where each world is a combination of a maximal natural component *n* and a maximal moral component *m*. This suggests a matrix with the columns representing maximally specific natural propositions and the rows maximally specific moral propositions. A purely natural proposition is any collection (disjunction) of maximal natural components – it is a 'vertical' proposition in the matrix. A purely moral proposition is any collection, or disjunction, of maximally specific moral components – it is a 'horizontal' proposition in the matrix. The two-attitude theories holds that belief is the determinate of acceptance appropriate to vertical propositions, and endorsement the determinate of acceptance appropriate to horizontal propositions. Fusions are the intersections of the vertical and horizontal propositions: they are rectangles in the matrix, and so we can define acceptance of those in terms of belief in the vertical component and endorsement of the horizontal component. But hybrids cut across the vertical and horizontal lines. They are propositions on the diagonal. Confining the logical space to just one non-moral proposition (C) and one moral proposition (T) apiece, the situation looks like this

	C: Cheney condoned torture	~C: Cheney did not condone torture
T: Condoning torture is wrong	C&T	~C&T
~T: Condoning torture is not wrong	C&~T	~C&~T

The moral propositions are horizontal rows (T and \sim T); the non-moral propositions are vertical columns (C and \sim C). Fusions are the intersections of the moral and non-moral propositions (C&T, \sim C&T, C& \sim T, \sim C& \sim T). Hybrids are everything else – any propositions on the diagonal: like $C\leftrightarrow T$, but also $C\rightarrow T$, $T\rightarrow C$, $\sim C\rightarrow T$, and $C\rightarrow\sim T$.¹

What, then, is the acceptance attitude appropriate to hybrid propositions? It is neither belief, nor endorsement, nor beldorsement. Let's call it *enlief*. Now, the postulation of enliefing is more radical than the embrace of beldorsement. The latter supervenes on, and is definable in terms of, belief and endorsement. Not so for enlief. Recall the nature of supervenience. A property Q supervenes on a base of other properties R just in case, sameness of distribution of R guarantees sameness of Q. So, is it the case that sameness in beliefs and endorsements guarantees sameness of enliefs? Can two people share the very same beliefs and endorsements but accept different hybrid propositions?

Consider *Cheney Hater* – a person who is agnostic both about the morality of torture and about whether or not Cheney condoned it. He doesn't know what to think about torture or the condoning of it, and also he doesn't know what Cheney is reported to have said, but accepts that whatever Cheney does on this issue, he is bound to be wrong, to unerringly choose the morally wrong thing to do. So Cheney Hater accepts the biconditional $C\leftrightarrow T$: *Cheney condoned torture if and only if condoning torture is morally wrong*. Now consider *Cheney Lover*, who as it happens is also agnostic about both the morality of condoning torture and about Cheney's actual behavior. He hasn't seen the TV interview and he distrusts the reports in the liberal media. But he accepts that whatever Cheney does, he unerringly chooses the morally right thing to do. So Cheney Lover accepts $C\leftrightarrow\sim T$: *Cheney condoned torture if and only if it is not the case that condoning torture is wrong*. Now $C\leftrightarrow\sim T$ is tantamount to $\sim(C\leftrightarrow T)$. Like $C\leftrightarrow T$, $\sim(C\leftrightarrow T)$ is a hybrid (the negation of a hybrid is always a hybrid) and one which, like $C\leftrightarrow T$, has tautologous moral and natural components. Cheney Lover and Cheney Hater share the same beliefs and endorsements (and *ipso facto* they accept the same fusions) but they differ with respect to the hybrids that they accept. They differ in their enliefs.

Enlief then, is a separate fact, over and above belief and endorsement. So to fully describe your acceptances you would have to list the non-moral propositions you believe, the moral propositions you endorse, (which jointly fix the propositions you beldorse), but as well, and independently, the hybrid propositions you enlief. But although these are separate attitudes, they have to mesh together appropriately in a rational being. Thus multi-attitude fictionalism is a bit too much to accept.

A fictionalist might well be tempted here to think that the source of these problems with fictionalism is the intrusion of the non-cognitive element, and be tempted back to some variant of cognitivism.² After all, if acceptance is a matter of belief, and belief transfer is unproblematic, then presumably we get acceptance transfer for free.

There are two options for the cognitivist fictionalist. One option is that to accept a proposition from the fictional discourse is to believe in the literal truth of that very proposition. However, since the fictionalist holds that there is no good reason to believe the

¹ A referee from the Journal made the following observation which is well worth clarifying: "Consider 'Cheney condoned torture if and only if condoning torture is wrong'. Isn't this simply a fusion of two fusions, namely: 'If Cheney condoned torture condoning torture is wrong and 'If condoning torture is wrong, Cheney condoned torture'? And isn't a fusion of two fusions simply a fusion, instead of a hybrid that requires yet another different attitude of enlief?" It is true that the conjunction of two fusions is a fusion, but the two conditionals involved here, $C\rightarrow T$, $T\rightarrow C$, are not fusions but hybrids. They both cut across the vertical and horizontal divides. Of course, this does not guarantee that their conjunction is a hybrid. $T\rightarrow C$ is a hybrid, as is $\sim T\rightarrow C$ but their conjunction is equivalent to C, neither a fusion nor a hybrid.

² This was suggested to us by Michael Tooley, and he cited van Fraassen's constructive empiricism as a promising model.

fictional propositions, the belief transfer principle has no purchase on arguments involving fictional premises. If there is no good reason to believe fictional propositions, the fact that they demonstrably entail some consequence does not provide one with a good reason to believe that consequence. The much more popular option is that to accept a proposition *A* about some fictional domain is to believe not *A* itself, but to believe some other proposition, $f(A)$, which we have called the *factual* or *real content* of *A*. The factual content of a proposition is that proposition which, according to this kind of fictionalist, one has to believe in order to accept *A*. Typically $f(A)$ is the proposition that *in the relevant fiction, A is true*. For example, to accept the moral proposition that *torture is wrong* is to believe that *in the moral fiction torture is wrong*.

It will be instructive here to briefly consider a fictionalist account completely outside the moral realm – a fictionalist account of theoretical entities. This is a simplified version of van Fraassen's constructive empiricism, which holds that to accept a theory *A* is simply to believe that *A* is observationally adequate – i.e. that all its observational consequences are true. In this case $f(A)$ is the strongest purely observational consequence of *A* itself. Does acceptance transfer fall out of normal belief transfer? Let *P* be any purely theoretical claim (like *there goes a proton*) and *O* any purely observational claim (like *there is such-and-such a streak in the cloud chamber*). Then since *P* is purely theoretical its observational content $f(P)$ is empty. $f(P)$ is logically true. Consider the conditional claim: $P \Rightarrow O$. This also has no observational content, as evidenced by the fact that its ramsification is a second-order logical truth. So $f(P \Rightarrow O)$ is also logically true. Assume, plausibly, that one has good reason to believe any proposition that is demonstrably a logical truth. So one has good reason to believe both $f(P)$ and $f(P \Rightarrow O)$. To accept *P* is to believe $f(P)$. So one has good reason to accept both *T* and $(T \Rightarrow O)$. These entail *O* by a single application of modus ponens. So if transfer acceptance holds then one also has good reason to accept *O*. To accept *O* is to believe $f(O)$, and since $f(O) = O$, that means that one has good reason to *believe* *O* itself. So, simply in virtue of believing demonstrable logical truths one would have good reason to believe any observational proposition at all. And that is clearly absurd. Quite generally, acceptance transfer will be problematic for any cognitive fictionalist whenever there are what we might call *f-hybrids*: that is, where there are propositions *A* and *B* such that the $f(A \& B)$ is logically stronger than $f(A) \& f(B)$.

The virtue of fictionalism over its anti-realist rival, expressivism, is that it avoids the Frege–Geach truth-transfer problem. That's a real advantage, however, only if the fictionalist can sustain some reasonable version of acceptance-transfer. What we have shown is that fictionalism almost invariably requires a double-attitude or multi-attitude account of acceptance, and that generates problems for acceptance transfer. Fictionalists of both the cognitivist and non-cognitivist stripes thus have an attitude problem.

Acknowledgements The authors would like to thank Eric Chwang, Richard Joyce, Simon Kirchin, Charles Pigden and Michael Tooley, for their valuable criticisms and comments.

References

- Harman G (2002) Internal critique: a logic is not a theory of reasoning and a theory of reasoning is not a logic. *Studies in Logic and Practical Reasoning* 1:171–186
- Joyce R (2001) *The myth of morality*. Cambridge University Press, Cambridge
- Kalderon ME (ed) (2005a) *Fictionalism in metaphysics*. Clarendon Press, Oxford
- Kalderon ME (2005b) *Moral fictionalism*. Oxford University Press, Oxford
- Mackie J (1977) *Ethics: inventing right and wrong*. Penguin Books, New York