ARE MORAL FUNCTIONALISM'S MORAL *A PRIORI* COMMITMENTS REALLY *A PRIORI*?

SERÁ QUE OS COMPROMISSOS *A PRIORI* DO FUNCIONALISMO MORAL SÃO, DE FACTO, *A PRIORI*?

Ashley Lane*
ashleylane1999@yahoo.co.uk

Abstract. *Moral functionalism, a metaethical theory developed by Frank Jackson and Philip Pettit, claims that we can attain moral knowledge by ascertaining the commonplaces about morality that are typically accepted by actual agents. It has important a priori commitments; whilst we may discover a posteriori that a particular descriptive property is identical to a particular moral property, it is a priori that the thing that is identical to the moral property, whatever that thing actually is, plays a particular role. Jackson holds a particular metaphysical position, and moral functionalism is a development of that position as it applies to ethics. In this paper I adapt an objection made by D.H. Mellor against Jackson's metaphysics to show that moral functionalism's a priori commitments are actually a posteriori. We can only discover if moral functionalism's purportedly a priori claims are true through a posteriori investigation.*

Keywords: *Moral functionalism, metaethics, metaphysics.*

Sumário. *O funcionalismo moral, uma teoria desenvolvida por Frank Jackson e Philip Petti, afirma que podemos obter conhecimento moral determinando os lugares comuns acerca da moralidade que tipicamente são aceites por agentes concretos. É uma posição que tem compromissos a priori importantes; podemos descobrir a posteriori que uma propriedade descritiva particular é idêntica a uma propriedade moral particular, mas é a priori que a coisa que é idêntica à propriedade moral desempenha um papel moral particular. Jackson defende uma posição metafísica particular e o funcionalismo moral é um desenvolvimento desta posição aplicada à ética. Neste artigo adapto uma objecção feita por D. H. Mellor contra a metafísica de Jackson para mostrar que os compromissos a priori do funcionalismo são, de facto, a posteriori. Só podemos descobrir se as afirmações alegadamente a priori do funcionalismo moral são verdadeiras através de uma investigação a posteriori.*

Palavras-chave: *funcionalismo moral, meta-ética, metafísica.*

---

* PhD student, Birkbeck, University of London.

Moral functionalism, a metaethical theory developed by Frank Jackson and Philip Pettit, claims that we are able to attain moral knowledge by ascertaining the commonplaces about morality that are typically accepted by actual agents. However, whilst we may be able to ascertain these commonplaces through *a posteriori* investigation, moral functionalism has certain purportedly *a priori* commitments regarding the extension of moral terms, which cannot be verified *a posteriori*. Moral functionalism represents the application of Jackson's general metaphysics to ethics, and moral functionalism's *a priori* commitments are a type of commitment that is also found elsewhere in Jackson's general metaphysics.

Even if these commitments are true, though, and contrary to what Jackson claims, they are not *a priori*. We can use an objection made by D.H. Mellor against Jackson's general metaphysics to generate a problem for moral functionalism. Jackson believes that we identify things such as water with what actually plays a particular role; in the case of water, H$_2$O plays a particular role which makes it identical with water. Mellor denies that this is true. As Jackson claims that moral properties are identical with whatever plays particular "moral property" roles, we can adapt Mellor's objection to criticise moral functionalism. The problem does not show that the commitments are false, but it does show that they can only be ascertained through *a posteriori* investigation. They therefore cannot be *a priori*, and it may be possible to show *a posteriori* that they are false. In the first section of this paper, I will set out moral functionalism's commitments. In the second section, I will outline Mellor's objection and adapt it to show that the commitments are not *a priori*.

1. Moral functionalism

Jackson and Pettit (1995) claim that agents use moral terms "in a way that presupposes a large network of connections with other terms, both evaluative and descriptive" (p. 22). We accept commonplaces about the terms we use, and this allows us to identify particular descriptive properties that ordinary moral thinking tells us are moral properties. We are thus able to grasp moral concepts. For example, when we use the term "fairness", it is a commonplace for us that if

an action is fair, that is usually a reason to do it. Agents who use moral terms competently know that the commonplace is true.

> [T]he meaning of relevant moral terms will be fixed by roles which certain commonplaces give them, and so moral thinking is bound to involve the attempt to use commonplaces as a base, and holding on to as much of that base as possible, or at least to the parts considered most secure, to fix opinions on particular questions. (p. 26)

Grasping the commonplaces correctly means that we grasp the concepts relating to the commonplaces correctly, which then means that we can correctly use the terms that denote moral properties. Once we know about the roles that moral properties play and the moral commonplaces that we accept, we can find out which descriptive properties we can identify with the moral properties.

Moral functionalism has a certain type of *a priori* commitment (It may have others, but I will only focus on one type here). Jackson (1998) claims that

> [w]hat is a priori according to moral functionalism is not that rightness is such-and-such a descriptive property, but rather that *A* is right if and only if *A* has whatever property it is that plays the rightness role in common folk morality, and it is an a posteriori matter what the property is. (pp. 150-151)

"Common folk morality" is the moral theory that endorses the moral commonplaces accepted by most people in a particular society (pp. 117-118). Moral functionalism proposes that we look at the commonplaces typically endorsed by people, and that we use that as the basis of a suitable moral theory. We hold on to as many of these commonplaces as possible, and they allow us to identify moral properties with particular descriptive properties. However, to know what moral properties are, we must know *a priori* that moral properties play a particular role. Only after that can we use *a posteriori* evidence to ascertain which descriptive properties are moral properties, since they play that role. To take a non-moral example, we discover *a posteriori* that water is H_2O. However, we have *a priori* knowledge that the substance that is water is the substance that plays a certain "watery" role of being transparent, drinkable, and so on. The *a posteriori* discovery is that H_2O plays the role, and so is identical to water.

Since the moral commonplaces we accept may conflict irreconcilably with each other, we can refine common folk morality into a mature folk morality. As people continue debating moral issues and refining their moral views, we work out which commonplaces to keep and which to drop. Jackson claims that this gets

us closer to a mature folk morality, a "folk morality that has been exposed to debate and rational reflection". It makes "good sense of the raft of sometimes conflicting intuitions about particular cases and general principles that make up current morality" (p. 113).

In his metaethics, as in his general metaphysics, Jackson separates the concepts that actual agents use from what he calls the "essence" of what these concepts relate to. For example, what is essential about water is that it is H_2O, but we do not need to know that to grasp the concept "water" (p. 50). People grasped the concept and could refer to water long before they knew its chemical composition. It is this distinction that makes clear the *a priori* commitments of moral functionalism.

Jackson accepts that the investigation of which commonplaces we endorse is *a posteriori*. The *a priori* element becomes apparent when we talk about the application of terms in the actual world or in counterfactual situations. Let T be a particular term. T may apply to various entities, events, relations, etc., and this defines its extension. It may have different extensions in different possible worlds. The *A-extension* of T in w is the actual extension of T in world w, where w may be our world or another possible world. In our world, the A-extension of "water" is all and only the occurrences of water.

We may ask, given the assumption that we are talking about our actual world, what T would apply to under various counterfactual situations. The answer would give us the *C-extension* of T, which would tell us what T would apply to across possible worlds. For example, the term "water" applies to all the watery occurrences in a world, so in the actual world the A-extension of "water" would be all and only the occurrences of what we call "water", which are occurrences of H_2O. In another world, where substances of chemical composition XYZ are called "water", the A-extension of "water" in that world would be all and only the occurrences of XYZ. But since in our actual world water is H_2O, the C-extension of "water" in the actual world is the occurrences of H_2O in every possible world (p. 49). The C-extension would not extend to any occurrences of XYZ, even though people on other possible worlds may call XYZ "water".

For "water", the A-extension and the C-extension in our actual world are the same. "Water" applies to all and only occurrences of H_2O under both extensions.

But, says Jackson, there is an epistemological difference. Before we found out that water was H_2O, we did not know its essence. So we could not determine its C-extension.

> [I]n order to pick out water in a counterfactual world, we need to know something about *relationships* between the counterfactual world and the actual world that we could only know after discovering that in the actual world H_2O plays the watery role. (p. 50)

Before we found out that water was H_2O – an *a posteriori* discovery – we could not say whether the stuff that plays the watery role in a counterfactual world was part of the C-extension of "water". For consider a counterfactual world that is identical with the actual world except that the stuff that plays the watery role in the actual world is H_2O and the stuff that plays the watery role in the counterfactual world is XYZ. Until we find out that water in the actual world is H_2O, how do we work out that the stuff playing the watery role in the counterfactual world does not fall under the C-extension of "water"?

A-extensions (in our actual world) do not need *a posteriori* knowledge of the actual world. "Water" is whatever plays the watery role in the actual world. We do not need to know that water is H_2O in order to know that. And this applies to any counterfactual world as well, because when we work out A-extensions of "water" in a counterfactual world, we do not compare the watery substance in that world to the watery substance in any other world.

Since we can know A-extensions in the actual world without knowing what the actual world is like, Jackson concludes that such knowledge is *a priori* (p. 51).[1] In the actual world, it is *a priori* that the A-extension of "water" is just the watery stuff of our acquaintance. Water is whatever happens to fulfil the watery role. What precisely that stuff is (H_2O) is an *a posteriori* question, but we do not need to know what precisely the watery stuff is in order to know that water is the watery stuff of our acquaintance.

What about the term "fairness"? The A-extension of "fairness" (at our actual world) will be the thing that we are acquainted with because it fulfils a certain "fairnessy" role, and we know this *a priori*. What we must discover *a posteriori* is exactly what descriptive property fulfils the role. So even though we can

---

[1] We cannot analyse all terms like this (such as names and demonstratives), and I will discuss this briefly later.

examine many moral commonplaces *a posteriori*, it is *a priori* that moral properties are simply the things that play the various "moral property" roles, that play certain roles in the eyes of moral agents. We can roughly say that moral properties are those that fulfil certain functions according to our common folk morality, which we can refine to become a mature folk morality. Metaethical theories that analyse moral concepts in ways that are completely alien to our ordinary conception of them are most likely wrong (p. 31). The way to find out what our ordinary conception is is to go out and look at what moral commonplaces most people endorse. What is *a priori* is our knowledge of the A-extensions of moral terms. What is *a posteriori* is our knowledge of the "essence" of moral properties. We investigate *a posteriori* what descriptive properties play the "moral property" roles, and so discover which descriptive properties are identical with which moral properties. This identification reveals the "essence" of moral properties, just as we discover that the "essence" of water is that it is H_2O.

Both A-extension propositions (propositions that express A-extensions) and C-extension propositions can be commonplaces. The proposition "The property of moral rightness is whatever plays the moral-rightness role" is a commonplace, and it expresses an A-extension. However, Jackson and Pettit (1995) also believe that for many people it is a commonplace that saving lives is morally more important than being fair (p. 23). This seems to be an *a posteriori* commonplace, since humans might have come to believe it a commonplace that being fair is morally more important than saving lives instead. We find out such commonplaces *a posteriori*, and in fact we come to identify particular descriptive properties as moral properties *a posteriori*. Such identifications are discoveries of moral C-extensions. Commonplaces are able to express A-extensions or C-extensions, and thus can be either ascertainable *a priori* or *a posteriori*.

Despite the significant *a posteriori* elements in moral functionalism, Jackson tells us that it ultimately rests on an *a priori* foundation of A-extensions. However, this foundation is not actually *a priori*. In the next section, I will examine an objection from D.H. Mellor that Jackson's endorsement of A-extensions in his metaphysics fails. I will use this objection to argue not that moral functionalism's A-extensions fail, but that they are *a posteriori*. They are thus vulnerable to *a posteriori* objections.

2. Mellor's objection

Return to Jackson's metaphysics. Consider the following argument:

a)   Sixty percent of the earth is covered by H_2O.

b)   Water is the stuff that plays the watery role.

c) H_2O is the stuff that plays the watery role.

Therefore

d)   Sixty percent of the earth is covered by water.

(a) and (c) are *a posteriori*. Before we could know (c), we had to check that H_2O actually plays the watery role. According to Jackson (2003), (b) is *a priori* (p. 87). It is something we can know *a priori* about water.

Consider now:

e)   Ben performed an act that maximised utility.

f)   The property of moral rightness is whatever plays the moral-rightness role.

g)   The property of maximising utility plays the moral-rightness role.

Therefore

h)   Ben performed a morally right act.

(e) and (g) are *a posteriori*, as (g) acts as a commonplace that we must discover *a posteriori*. (f) is *a priori*, which is what we should expect, since it is an A-extension proposition. Jackson's justification for (b) and (f) being *a priori* is that they are about reference fixers. As far as (b) goes, "water" refers to what plays the watery role, the stuff that plays the role of a colourless drinkable liquid that makes up the ocean and so on. It does not indicate what actually plays the watery role or what the watery role consists of.

The same happens with (f). (f) may be *a priori*, but note how little information it contains. It does not tell us a great deal about moral properties or what their characteristics are. But without it, Jackson claims, we could not

discover which descriptive property is identical with the moral-rightness property. The *a posteriori* discovery rests on *a priori* assumptions about moral properties.

Currently, then, we have *a priori* A-extension propositions that the moral functionalist uses to justify moral C-extension propositions. The critic's aim is to show that this type of justification always fails, and that *a priori* A-extension propositions cannot justify any C-extension propositions. There are at least three strategies what we can use.

a) Moral A-extension propositions are false, and so cannot help justify any C-extension propositions.

b) Moral A-extension propositions are really *a posteriori*, contrary to what moral functionalism claims.

c) Moral A-extension propositions cannot be used to justify any moral C-extension propositions, even if the A-extension propositions are true.

If option (c) is right, then moral functionalism has *a priori* commitments, but these commitments are effectively redundant in moral arguments. We cannot use them to justify any moral C-extension propositions, and so their role in a folk morality is very small. Moral C-extension propositions must be justified on other grounds which have nothing to do with A-extensions. However, as my focus is on options (a) and (b), I will not develop option (c) here.

Option (a) looks attractive. The proposition "X is whatever plays the X-role" (whether X is a moral property or something else) may be wrong in several ways. Most obviously, the X-role may not exist, although this is unlikely to be a truly damning objection to Jackson's metaphysics. For example, "X" must not be a name; London exists, but it is not the case that the thing called "London" must play a "Londony" role. People just applied the name "London" to a particular entity. However, all Jackson has to do is exclude names from his metaphysics, and he is back in business. It seems hard to claim that there is no watery role for water to play, because the watery role is just a collection of properties that the thing that is water has to have. "Water" is not a name. The same point holds for demonstratives such as "this" and "these". The proposition "This is whatever

plays the thisy role" is false, because there is no thisy role, but Jackson has no reason to claim otherwise.

More interestingly, it may be that we are mistaken about the nature of X, and hence of what the X-role should be. Suppose a scientist finds a new chemical that appears to dissolve plastic. For something to be this chemical, he believes, it has to play a role that includes the ability to dissolve plastic. It comes to be generally accepted that if something cannot dissolve plastic, it cannot be the chemical. But then it is discovered that the scientist made a mistake, and the chemical does not dissolve plastic. The chemical role includes the ability to dissolve plastic, the chemical itself cannot dissolve plastic, and so the chemical is not identical with whatever plays the chemical role.

This is not the only way in which we might be mistaken about the relation of X to the X-role. Take the example of water. We say that water is transparent, drinkable, found in oceans, and so on, and to Jackson these properties constitute the watery role. It was not necessarily the case that anything has all these properties, though, or that there was only one thing that did. We might have examined oceans and found out that what we call "water" is actually a mixture of H_2O and XYZ.

Another possibility is that the thing that actually plays the X-role also has other properties that X itself does not have, and so the thing that plays the X-role cannot be X. D.H. Mellor (2003) argues that H_2O is not identical to water, despite playing the watery role. Even if we allow that ice and steam are water, "no single H_2O molecule can be water, since it instantiates hardly any of water's laws, having no solvent powers, density, freezing or boiling points, or latent heats" (p. 224). Of course, Mellor does not deny that water exists or that it is made up of H_2O. What he denies is that H_2O is identical to water, despite it playing the watery role.

It is tempting to think that we can make the same sort of objection about moral A-extensions, and thus show that option (a) is correct. Actually, it is difficult to defend option (a), and it is much better to use option (b) to criticise moral functionalism. I shall outline a couple of ways in which one might try to use option (a) before moving on to Mellor's objection. Initially, it appears that Mellor's objection can be used to show that option (a) is right. However, it really ends up showing that option (b) is correct.

If option (a) is correct, then the following moral A-extension proposition is false:

MRP: The property of moral rightness is whatever plays the moral-rightness role

MRP can be questioned in several ways. First, it may be the case that it is impossible for anything to play the moral-rightness role. This could happen, for example, if the role is logically inconsistent. If the property of moral rightness exists, MRP cannot be true. Not only that, but the moral functionalist is left trying to find another way in which moral agents can recognise moral properties. MRP offered a straightforward way to do so – check what descriptive properties play a certain role, and those properties are moral properties. Now another method must be found.

So the objection throws up two problems for the moral functionalist. The first problem is metaphysical – if moral properties are not the properties that play a particular moral role, what are they? The second is epistemological – how can agents recognise moral properties if the properties do not play a particular role? However, the moral functionalist need not worry yet, because it first has to be shown that it is impossible for anything to play moral roles. The moral functionalist will of course deny that it is impossible, because we seem to identify moral roles all the time with the aid of the moral commonplaces we accept. The "fairnessy" role is the role of being even-handed, impartial, unaffected by biases, and so on. Where does the logical inconsistency lie in this? And if we can identify an inconsistency, why can we not just alter our conception of the role to exclude it? Jackson explicitly states that we can refine our moral commonplaces to exclude inconsistencies and so develop a mature folk morality.

Another way that MRP may be false is if there is no connection between the property of moral rightness and the moral-rightness role. Suppose we discover a new non-moral property, Y, and we ascribe a number of characteristics to it that come to define the Y role. We then discover that Y has near to none of those characteristics. The property of Y is not what plays the Y role. Could this not be the case for the property of moral rightness? The trouble with this objection is that it once again does not allow for correction and refinement. As we find out more about Y, we refine our understanding of what the Y role is, so once we have

realised our mistake we change our conception of the Y role to reflect characteristics that Y actually has. Our conception of the Y role changes so that whatever plays the Y role is indeed identical to Y. This does not seem impossible, so the moral functionalist can say the same thing about moral properties. We may be currently mistaken about what the "moral rightness" role is, but by examining the property of moral rightness we can correct our beliefs about what the "moral rightness" role is, so that whatever plays that role is identical to the property of moral rightness.

The third objection is a variation of Mellor's metaphysical objection. Something plays the moral-rightness role, and there may be such a thing as moral rightness, but the property that plays the moral-rightness role is not identical to the property of moral rightness. Take MRP. Moral-Mellor may say about it:

> Assume that there is a property that plays the moral-rightness role. Let us say that the property of maximising welfare pays it. That certainly does not mean that such a property is the property of moral rightness, because it may not play that role at all times and under all conditions (H_2O plays the watery role, but it does not always do so.). And we certainly cannot assume ahead of investigation that there is any single descriptive property that is identical to the property of moral rightness, or that there is a collection of descriptive properties that constitute the property of moral rightness.

Call this the Moral-Mellor objection. This objection concedes a certain amount to the moral functionalist. It allows that agents can recognise moral properties by examining what plays moral roles (Even if water is not identical to what plays the watery role, examining what plays the watery role will help us identify examples of water.). It is consistent with the claims that there are moral properties, that there are moral roles that descriptive properties play, and that we can know about moral properties and roles. There may also still be a metaphysical connection between a moral property and a moral role. If something does not play the watery role, then it cannot ever be an occurrence of water. Similarly, if something does not play the role of a moral property, then it can never be identical to the moral property. The only claim that is rejected is the moral functionalist's claim that moral A-extension propositions are correct. It may be possible to know the proposition "Maximising the welfare of the homeless is morally right", even though the property of maximising welfare is not identified with any particular moral property.

At this point, the moral functionalist can reasonably object that we cannot just assume that moral A-extension propositions are false. Perhaps when we investigate further, we will discover that the propositions are true. This is a fair comment, but we can now show that even if they are true, they are not *a priori*. H_2O is identical to water, says Jackson. Not so, says Mellor. Mellor's objection appears to be an *a posteriori* objection. H_2O does not always physically act like water; for example, occurrences of H_2O need not have any boiling point. By observation we attribute various features to occurrences of water and of H_2O, and it is a matter of *a posteriori* fact that these features are not always the same.

So it appears that we discover whether the proposition "Water is the stuff that plays the watery role" is true or false by *a posteriori* observation. H_2O plays the watery role, but that does not necessarily mean that it is identical to water. We cannot discover *a priori* that it is water, and so we cannot discover *a priori* that water is identical to the stuff that plays the watery role. By observation we attribute various features to occurrences of water and of H_2O, and it is an *a posteriori* matter whether these features are always the same. The proposition is an *a posteriori* proposition, and we can argue analogously about moral A-extension propositions. We can only discover if, say, MRP is true by investigating *a posteriori* what people take the moral-rightness role to be, seeing what plays that role, and then checking whether it conforms exactly with the content of "the property of moral rightness". If the proposition is true, we can only discover that it is true by *a posteriori* investigation. If MRP can only be shown to be true *a posteriori*, though, that means that it is not *a priori* and therefore option (b) is correct. Moral A-extension propositions, if they are true, can only be verified *a posteriori*, not *a priori*.

A possible disanalogy between Mellor's original objection and the Moral-Mellor objection is that whilst we discover *a posteriori* that water is not identical to what plays the watery role, this is not the case for moral properties. With water, we make the identification of water with whatever plays the watery role *a priori*, and we make the identification of water with H_2O *a posteriori*. However, the objection claims, we make the identification of a moral property with whatever plays the "moral property" role *a priori, and also* we make the identification of

the moral property with the descriptive property that plays that role *a priori*. Moral C-extension propositions are thus *a priori* themselves, not *a posteriori*.

The moral functionalist cannot make this claim. The less important reason for this is that in Jackson's general metaphysics, he states that C-extensions are knowable *a posteriori*. The moral functionalist may choose to drop this requirement for moral C-extensions, and so detach moral functionalism from the rest of Jackson's metaphysics, but this brings us to the main reason. We ascertain moral C-extensions by examining certain moral commonplaces, and those commonplaces are discovered *a posteriori*, by observing how actual agents behave and what moral judgements they make. Moral functionalism is committed to "holding on to as [many of these commonplaces] as possible", so the moral functionalist is committed to saying that C-extension propositions are *a posteriori*. It follows that if our discovery of moral C-extensions is based on *a posteriori* discovery of moral commonplaces, and if we can only ascertain moral A-extensions by ascertaining moral C-extensions, then the discovery of moral A-extensions is based on *a posteriori* discovery as well.

Even if there are true moral A-extension propositions, this does not mean that option (b) is wrong. The moral functionalist hopes to identify moral properties with particular non-moral descriptive properties by moving from the *a priori* to the *a posteriori*. He starts with *a priori* moral A-extensions, and uses them to justify identifying moral C-extensions, which is done *a posteriori*. But there are other metaethicists, such as Jesse Prinz (2007), who in effect work from the other direction. They attempt first to work out *a posteriori* that particular non-moral descriptive properties are moral properties. In Prinz's case, he identifies moral properties with properties that dispose moral agents to have particular emotions.[2] After doing so, such metaethicists may if they want say that moral properties are whatever plays certain roles. They thus move from C-extensions ("Moral property X is identical to non-moral property Y.") to A-extension propositions ("Moral property X is whatever plays the X-role.")

---

[2] Prinz and the moral functionalist actually agree on a lot. They are both reductionist moral realists who believe that we must investigate human behaviour and attitudes *a posteriori* to discover which non-moral properties are identical with moral properties. However, Prinz gives much more importance to emotions such as guilt in his metaethics, and he certainly would not accept *a priori* moral A-extension propositions.

However, the A-extension propositions are justified by the C-extension propositions. The X-role is defined by the content of the non-moral property Y. In this case, the metaethicist can say that there are true moral A-extension propositions. However, they are not *a priori*, because they are justified by *a posteriori* C-extension propositions.

3. Conclusion

Moral functionalism claims that there are moral A-extension propositions that can be ascertained *a priori*. These propositions provide justification for *a posteriori* discovery of moral C-extensions. Two types of objection to moral functionalism's claims here are that a) moral A-extension propositions are untrue (and hence cannot justify anything), and b) moral A-extension propositions are not *a priori*.

D.H. Mellor objects that A-extension propositions are untrue – it is not the case, for example, that waster is identical with whatever plays the watery role. An analogous argument can be made that moral A-extension propositions are untrue for the same reason, but this goes too fast. We cannot just declare automatically that moral A-extension propositions are false. However, we can use the objection to point out that we can only ascertain *a posteriori* that moral A-extension propositions are true. So moral functionalism's commitment to *a priori* A-extension propositions becomes a real problem. If moral A-extension propositions are true, they cannot be *a priori*.

I must point out here that I have been concentrating only on the A-extension propositions that moral functionalism relies on. Moral functionalism may have other *a priori* commitments, and I say nothing about them. However, A-extension propositions are particularly important for moral functionalism, since to Jackson they are prior to C-extension propositions. The moral functionalist thinks that even if we discover *a posteriori* that descriptive property X is moral property Y, we do so because X plays the Y role, and the identification of Y with what plays the Y role is discovered *a priori*. But the moral functionalist is wrong

on that point. We have to go out and look, just as we have to go out and look at whether water is whatever plays the watery role.[3]

References

Jackson, F. (1998). *From metaphysics to ethics*. New York, NY: Oxford University Press.

Jackson, F. (2003). From H_2O to water: The relevance to a priori passage. In H. Lillehammer & G. Rodriguez-Pereyra (Eds.), *Real Metaphysics: Essays in Honour of D.H. Mellor* (84-97). London: Routledge.

Jackson, F., & Pettit, P. (1995). Moral functionalism and moral motivation. *Philosophical Quarterly*, *45*(178), 20-40.

Mellor, D.H. (2003). *Real metaphysics*: Replies. In H. Lillehammer & G. Rodriguez-Pereyra (Eds.), *Real metaphysics: Essays in honour of D.H. Mellor* (217-238). London: Routledge.

Prinz, J. (2007). *The emotional construction of morals*. New York, NY: Oxford University Press.

---

[3] An earlier draft of this paper was presented at the VIII Meetings on Ethics and Political Philosophy at the University of Minho. I would also like to thank Hallvard Lillehammer and Cristian Constantinescu for their helpful comments.