

MORAL BEHAVIOR AND MORAL REALISM: THE CASE FOR ARGUMENTS IN ETHICS

Fabrizio Pontin
PUC-RS

RESUMO: Meu objetivo aqui é mostrar o quão fácil é lidar com o utilitarismo de preferências e maximizar a utilidade pessoal (ou animal) quando somos confrontados com um pequeno conjunto de ações possíveis. Uma vez que a reduzimos o escopo das ações possíveis, ou seja, ao que é factível, nós também nos tornamos facilmente capazes de identificar o comportamento adequado ou inadequado, ao menos em termos da maximização racional das preferências desejadas.

ABSTRACT: My intention here is to show how easy it is to deal with preference utilitarianism and to maximize personal (or animal) utility when we are faced with a small set of possible actions. Once we reduce the scope of possible actions, that is, what it is *feasible* to be done, we are also easily able to identify the proper or improper conduct, at least in terms of rational maximization of desired preferences.

Consider the case of a cat. I have a couple of these creatures, and my interactions with them are usually borderline asinine. A typical interaction between me and my cats, however, usually involve a peculiar set of words: words of encouragement (good cat!/nicely done Fiddle MacFee!/who dat nice kitty? You dat nice kitty, yes you are!) or discouragement (bad kitty!/no kitty, no!/ Stinky Ol'Katt, you, sir, are in deep trouble!). While using such words I attribute some value to actions my cats are performing, and I am hoping that they are able to respond to my input by reproducing action that I

find positive (actions peculiar of “good” kitties) and avoiding actions that I regard as negative (let’s call those “bad” kitties actions).

Now, it is relatively easy to consider the scope of “bad” and “good” actions in kitty behavior. For example, messing up a litter box is easily recognizable as a “bad” action, as, inversely, bringing back a hair presser that I’ve just thrown away is a “good” action.

If Stinky decides to press Fiddle’s neck too harshly, or, as it was once the case, biting her ear hard enough to drill a hole into it, then, he is decidedly a “bad” kitty. But the scope of possible actions for cats is quite limited (cats cannot, after all, build atomic bombs, and let’s all now consider our luck that such is not the case), and consequently so are the possibilities of attributing the adjective “good” or “bad” to a cat action. Moreover, we would be in danger of anthropomorphizing the animal if we were to presume it is able to understand that it is “good” or “bad”. The set of pleasures that my little kitties are capable of activating is rather peculiar to kitties, that is, they respond to their favorite treats (and lack thereof) and the dreaded noise of the vacuum cleaner. So, I know that a very effective way of enforcing the Rule of Right towards my cats is to approach the closet door where I keep the vacuum. Of course, my kitties do not know that there is a vacuum cleaner in the closet. All they know is that I am able to produce something capable of creating a terrible, scary, noise that makes them seek for immediate comfort under the bed. So, the set of priorities for my kitties could be so summarized:

Pleasure priorities for Fiddle MacFee (in Mill’s terms, “preferences”):

- 1 – Canned food
- 2 – Star shaped treats
- 3 – Sleep
- 4 – Have its belly rubbed

Pleasure priorities for Stinky Ol’Katt:

- 1 – Canned food
- 2 – Sleep
- 3 – Torment Fiddle MacFee
- 4 – Hunt for flies

Avoidance priorities for both Fiddle MacFee and Stinky Ol’Katt:

- 1 – The Vacuum Cleaner
- 2 – Lightning
- 3 – Fireworks
- 4 – The travel box

What have we learned from this little exercise?

My intention here is to show how easy it is to deal with preference utilitarianism and to maximize personal (or animal) utility when we are faced with a small set of possible actions. Once we reduce the scope of possible actions, that is, what it is *feasible* to be done, we are also easily able to identify the proper or improper conduct, at least in terms of rational maximization of desired preferences. An analogy here could be done with newborns. Newborns have a very small subset of priorities, almost all of them related to the universe of “Mom” – they also have a relative small subset of possible “bad” actions. As a matter of fact, the form of being in the world of a newborn is so restricted, that we take a while to attribute the adjectivation of “bad” to baby actions.

While we are speaking of my kitties, it is relatively easy to be a realist, that is, to attribute moral significance to certain actions as if they really were “good” or “bad”, a very real action is connected to a very real consequence which really is “good” or “bad”. I know that my kitties are terrified of the Vacuum cleaner, so I can use this to discourage bad behavior. It is equally as

efficient to use their preferred preferences (such as canned food) as a way to enforce positive behavior. I am also aware that it is actually a good thing that my kitty is not devastating its own litter box (meaning that it brings real benefits for my kitty to be clean) and that I should not allow Mr. Stinky Ol'Katt to eat as much as it wants, as it will likely get sick and fat. Again, there are very real benefits in avoiding certain priorities – but this is much easier done when we are dealing with small sets of possibilities, that is, with a small moral universe.

The perennial question “Am I a good person?” is rather more difficult to answer than the question “Is Stinky Ol'Katt a good cat?”. What if more than just difficult, the question is totally loaded? What if it is a catch question?

Moral epistemologists are usually divided on this issue. Moral skeptics, most famously J.L. Mackie, tend to follow Wittgenstein’s famous proviso that Ethics cannot be the case, and hence propositions about values can be neither right nor wrong. Moral realists, inversely, tend to scoff at this proviso, pointing at all the circumstances in which, indeed, we follow moral rules at least *as if* they were the case. Moral constructionists tend to try for a mid-field, either recognizing that the myriad of cultural facts disavows the possibility of stable moral facts, while, at the same time, pointing at general intuitions for a cultural norm that “x” rather than “y”, or, at some cases, going at a general moral relativist “commando” attitude that says that moral attitudes are constructed and real in the reality they are constructed as such, but that a norm “x” can not be really better than a norm “y”. Interestingly, we can find these different moral epistemological attitudes within different styles of moral philosophy, an utilitarian may be a moral skeptic (as it was the case of Mackie) or a realist (Mill comes to mind), a deontologist may be a constructionist

(Habermas comes to mind) and a relativist may defend cultural supremacy (Rorty seems to have defended this position in his last papers).

Some moral realists, however, seem to have created what one could describe as an *empirically informed* theory of moral facts. Jesse Prinz summarizes this project when he writes that “[P]sychology and neuroscience are being used to determine what goes on in people’s heads when they decide that something is morally right or morally wrong (...) [e]mpirical methods can actually shed light on what we ought to do” (Prinz, 2007(b):271). Prinz’s description seems to fit Paul Churchland’s project in *Neurophilosophy at Work* (2007), though Prinz’s own account seems to rely a bit more heavily on cultural claims than Churchland.

In this paper, however, I want to provide a small summary of the arguments in what I will call “moral realism qua moral physiology” in Paul Churchland. Ultimately, I want to see how Churchland’s argument might help us answer the question “Am I a good person?”.

First, I think it is important to denote what Churchland is *not* trying to do. It is quite clear, from start, that Churchland has no interest in providing us with a moral psychology, nor is he interested in giving us an account of the inflation of arguments within our moral universe by culture. As we will see, it is not that Churchland thinks that these issues are not important – he thinks that they are already operating within a very peculiar notion of representation. Clarifying Churchland’s notion of representation, then, becomes paramount in understanding what he is indeed trying to do.

Representation as mental activity

I will not insult anyone’s intelligence by pretending I understand Paul Churchland’s account of the chemistry and physics behind mental representation. For the most part, I will assume that Professor Churchland

knows what he is talking about, and I will, perhaps with a degree of temerity, take his elegant description of equilibrium and nonequilibrium states within the brain to be the case.

But, given the narrative of brain functionality we obtain in Churchland, we can conclude some things about the way in which he understands representations:

- Our physiology is subject to the same laws of physics and thermodynamics any other physical body on this Earth is subject to
- Hence, the behavior of our physiology *must* follow certain rules on a molar level
- Such rules indicate the limits of what we are physiologically able and unable to do
- Before “cognitive phenomena” may occur, some instances for the possibility of cognitive phenomena are given
- Cognitive phenomena hence occurs under a set of possible representations which are already inducing cognitive phenomena
- We, as physical identities “me”, “you”, “us”, “ourselves”, are in fact literal epistemic engines – our physiology *induces us* to produce cognitive relations in the same way a propulsion engine induces a car to move (Churchland, 2007:30)

So far, I have said nothing about representations, right? Apparently, all this talk about non-cognitive instances of physiological operation which induce future cognitive production seem to be previous to any in-fact representation, correct?

Allow me to try to interpret this a bit differently. In fact, Churchland seems to break up representation into two categories. First, we will have non-cognitive, non-linguistic, representations. These representations establish a set of primary “stuff” in which all future linguistic representations will feed. Linguistic representation, in that sense, is conditioned by non-cognitive representations of physical interactions – and brain physiology can, within limits, already be identified as operating properly or improperly.

This second proviso seems particularly important for many moral realists, and I want to spend some time breaking it up before moving forward. What do we mean by a “properly functioning brain”? This is definitely a complicated question to ask philosophers, as many of us are still reminded, with good reasons, of Pavlovian and Lambrosian nightmares when discourses about “proper mental functioning” are argued. Churchland was one of the precursors in the current thesis that we should not hesitate in arguing for proper mental functioning, and his point seems to be that if there are proper, that is, ideal conditions for physical facts, for chemical facts and for biological facts in general, then it is also true that there are ideal conditions for the operation of the brain. Will that mean that all properly functioning brains will think the same way? I don’t think so. But they will not suffer certain strains that will disavow the possibility of certain acts of empathy, for example.

In that sense, we cannot hope that someone with extensive front lobe damage will be able to interact, and have the same sort of possibilities for linguistic mental representations that those of us without extensive front lobe damage have (Churchland, 2007:48).

However, front lobe damage refers to a sort of problem in what I would call the *first level* of mental representations. That is, we are dealing with individuals whose noncognitive structure of the brain is operating in a way that does not allow for cognition regarding social interaction in the ways we would expect. And we tend to adapt our social interactions with individuals suffering from these ailments accordingly – in the same way we adapt our social interactions with babies, and, with less of a social weight, with cats, dogs, horses and ants (and let us not forget, these are all creatures capable of non-cognitive representations in the sense Churchland is describing).

What about our second level representation? I mean those representations that we are able to indicate and express through the use of

language, memory and culture (and the jury is out on whether these are different ways of saying the same thing for Churchland).

Here we start to enter the field of moral realism in Churchland. For Churchland, we have no reason to believe that our progress regarding our understanding of facts of science cannot find a correlate in an understanding about facts of morals. Here, I think, the keyword is what he calls “worthy cognitive achievements”.

For example, our understanding of gravity, of the laws of physics, our modern suspicion against folk science, our *cognitive progress* which is deeply connected to our pursuit of proper scientific analysis will eventually allow us to set foot on the Moon. This is no small feat for individuals that just a while ago were stepping down, fearfully, from trees only to be smashed by the nearest hungry beast.

Churchland points that our basic mental representations were already there when we went down the trees. The stuff that was *not* there were the material conditions for the development of such mental representations (those second level materials). Now, if we are able to recognize that there has been *moral* progress during this same time, on what basis can we address such progress?

First of all, moral discrimination (Churchland, 2007:43) participates within our matrix of synaptic connections in the same way that, say, our ability to walk, and, on a more cognitive level, our ability for discriminating laws of physics from folk science. In that sense, our ability to make moral discriminations will trust a subset of moral judgments, memories and skills – in the same way that our ability to understand the chemistry in *Neurophilosophy at work* presupposes a previous set of memories and skill regarding organic chemistry.

Consider, again, the case of cats. The set of possible moral skills we attribute to cats is as small as what we perceive as their ability to interact with the world. Our attribution of “good” and “bad” kitty behavior, then, is connected to what we identify as their possible “skills of goodness”, that is, what they may or may not even conceivably develop as behavior. If we now jump to our own secondary cognitive abilities, we see that the picture is more complex.

First of all, what we are and are not to conceivably develop as behavior is still open. We are not *fully* aware of the possibilities of our brain regarding, for example, future progress in our understanding of physics. If we take Churchland seriously, we should move with the same caution regarding our future progress regarding our understanding of morality.

If, however, we accept this, in my view very modest, request, we can might feel that we have reason to trust some points regarding the possibility of *learning* morality in terms of right and wrong in the same way we learn physics in terms of right and wrong:

When we are taught $F=m.a$, are we taught that $F=m.a$ because our western culture says that $F=m.a$? Or do we say that $F=m.a$ because we have reason to say that $F=m.a$ in any circumstance within this reality which is, until now, the only one we know of?

If we opt for the latter, and I hope we do, then we are recognizing some important things:

- We educate people that truth matters
- Truth matters because it says something about the world in which we live in
- Our understanding of the world makes us better in understanding how things are and are not

Now, all these assumptions rely on a notion of cognitive progress, whether we like it or not. Without modern physics, we would not know that $F=m.a$. We would not know what sort of body would take to enter into an area of high-level atmospheric pressure without causing an explosion, and we would not know, as we now know for sure, that it would be absolutely impossible for a body made of flesh to reach higher atmospheric pressure without exploding. All of these things are facts.

We are hence able to point at things that are worth pursuing, from a cognitive point of view and things that are not. For example, pursuing a theory of how flesh might enter atmospheric pressure, without protection, and without exploding, is cognitively absurd. Anyone with a basic understanding of physics will know that such idea is absurd. Pursuing a “Flat-Earth” theory of geology is equally absurd, for the same reasons.

Churchland will then ask us to look at moral cognition in the same way we look at scientific cognition (Churchland, 2007:54;57). This will entail, necessarily, in recognizing that the representational character of morality is given to us, in noncognitive mental representations, in the same way that physics have been given to us.

This is to recognize that there *is* a proper way to instigate moral behavior, and to educate moral behavior, and the patters for this instigation will follow the same patterns of activation that have allowed us a better understanding of physics.

I think there are two ways of interpreting this in terms of anomic behavior for Churchland. I will try my hand at both, as I think that to some extent Churchland himself defends each of these visions at times, for different reasons. First, we will see moral anomy as cognitive deficiency (Churchland, 2007:67). If “a nondiscursive form of [moral] cognition underlies all of the more advanced forms”(Churchland, 2007:66), this means that any possible

expression of a moral preference refers back to a set of “stuff for morality” within our primary, non-linguistic, cognition (which I called non-cognitive, as there is no performative activity from any “self” going on here). In that sense, any anomie discursive behavior regarding basic emotions can be empirically traced to some degree of failure within the synaptic connection allowing for this behavior. Such behavior may be a case of psychotic behavior, due to very real accidents in brain physiology, and in this case there is little hope for re-education as a tool, or it may be the result of a badly wired cognitive process.

Say, an impoverished childhood in a warzone might produce a set of “cognitions” about expected moral behavior and response which create patterns of activation within brain chemistry – what folk psychology used to call “trauma”, in what Churchland would certainly identify as the pre-history of neurosciences. Such form of moral anomaly is *still* a cognitive deficiency in the terms developed by Churchland. But what is interesting about this sort of cognitive deficiency is that cultural pollution undermines the *proper* development of the potential of moral cognition. This entails us with a responsibility to infer that it is cognitively unacceptable to defend the uses of resources in activities that lead to the corrosion of the structure that allows for the proper development of moral cognition – in that sense, investing more public resources in tools of aggression rather than tools of education is completely asinine from a cognitive point of view. More importantly: it is a moral *fact*, in the terms developed by Churchland, which means that defending otherwise is just as wrong as saying that $a.F=m$.

Second, we see anomaly as cultural deficiency. Anomaly is also a measurement of social and moral progress for Churchland. In the same way that sectors in society might defend Homeopathic “medicine”, and we *know* these sectors are downright wrong, we also know that defending certain moral behavior is equally wrong. At least, we *should know better*. Churchland will point

that certain cultural standards are just repetition for the sake of comfort. People will repeat cultural norms that have become sediment for a myriad of reasons, and it is then necessary to implement the sort of public policy that will not see such manifestations of public dissent as a cultural value to be reckoned and treasured, but rather as bizarre attempts at moral cognition from a bygone era (in the same way that the Flat-Earth theory is a bizarre attempt at cosmology from a bygone era, and a wrong theory at that).

Consequently, we will need a regulative machinery that conducts our collective affairs regarding what we are already able to identify as the minimal parameters for acceptable moral cognition – disallowing the repetition of parameters we have found to be unacceptable.

In that sense, we already have the means to identify what it is to be a “good person”. A good person is someone who is operating within parameters of moral behavior that can be expressed and transitively organized in a way that help us pursue move from merely moral behavior to moral wisdom – something that can only be achieved if we look at the patterns of activations our minds follow when they are operating properly.

In guise of conclusion, allow me to once again refer Prinz’s fantastic essay from 2007(b): we must look at experimental psychology, at behavioral attitudes under pressure, and learn from these attitudes. The famous Milgram experiments are fundamental here: if we know that we tend to have certain responses under stress, and that these responses sometimes motivate us to act against all our moral cognitions, what can we do to reprogram ourselves? How can we activate a different set of behavioral attitudes when confronted with authority so we are not so willing to obey? A normative reconstruction of social reality and, more importantly, moral reality, will have to trust that facts *must* overcome transcendental Hocus Pocus and discursive over-elaborations that have no support in empirical data. Value inconsistencies, then, are a result

of our inability to understand that proper structure of our social and mental interactions as they actually are, and we have all the tools to understand these as they are, instead of as we wish they were.

The Boss had said it so well that:

You been hurt and you're all cried out you say
You walk down the street pushin' people outta your way
You packed your bags and all alone you wanna ride,
You don't want nothin', don't need
no one by your side You're walkin'
tough baby, but you're walkin' blind
to the ties that bind

REFERENCES

CHURCHLAND, P.M. *Neurophilosophy at Work*. Cambridge, MA: Cambridge University Press, 2007.

HAIDT, J. "The Moral Emotions." In *Handbook of affective sciences*, eds. R. J. Davidson, K. R. Scherer, and H. H. Goldsmith, 852–70. Oxford: Oxford University Press, 2003.

MOLL, J., de Oliviera-Souza, R., Eslinger, P. J., Bramati, I. E., Mourao-Miranda, J., Andreiuolo, P.A., and Pessoa, L. "The Neural Correlates of Moral Sensitivity: A Functional Magnetic Resonance Imaging Investigation of Basic and Moral Emotions." *Journal of Neuroscience* 22, 2002, 2730–36

PRINZ, J. J. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press, 2002

———. *The Emotional Construction of Morals*. Oxford: Oxford University Press, 2007(a).

———. "Can moral obligations be empirically discovered?" In H. Wettstein (ed.), *Midwest Studies in Philosophy* 31, 2007(b).