A Problem for Moral Naturalism: Outsourcing Moral Judgments

Rafael Vogelmann

Federal University of Rio Grande do Sul
Department of Philosophy
Porto Alegre, RS
Brazil
rafael.vog@gmail.com

ABSTRACT

Moral Naturalism is the view that moral judgments aim at describing moral facts and that these are ordinary garden-variety natural facts. Moral Naturalism has trouble accounting for the intuition that we cannot outsource moral judgments, i.e., we cannot ground a moral conviction that p on the fact that a reliable moral adviser holds that p. There have been, however, several attempts to explain this intuition away or to discredit the intuition pumps that bring it forward. I argue that moral naturalists are not in a position to deny this intuition. Moral Naturalism embodies a conception of the minimal conditions for someone to qualify as capable of making moral judgments; among these conditions is the acknowledgment of the supervenience of the moral on the non-moral; given the naturalist's conception of what it takes for someone to be capable of moral judgment, if we allow agents to outsource their moral judgments we come to situations in which the convictions of moral agents do not comply with the acknowledgment of moral supervenience. The naturalist must, therefore, deny the possibility of moral outsourcing. Moral Naturalism, then, faces the problem of making sense of the ban on moral outsourcing.

----

1. Introduction

While we think it is appropriate to ground ordinary beliefs on the testimony of reliable well-informed advisers, there seems to be something fishy about grounding our moral convictions on the testimony of others, even when we

believe them to be morally wise, well-informed and well-meaning. This intuition can be brought to light if we compare cases of ordinary, non-moral testimony with cases of moral testimony.

> Non-Moral Testimony: Peter plans to go on a hike tomorrow. He wonders if it is going to rain. He cannot figure it out by himself so he checks the weather forecast. He is told that tomorrow is going to be a sunny day. He comes to believe that. When asked why he believes that tomorrow is going to be a sunny day he replies: "because that is what the weather forecast says and I trust it".

> Moral Testimony: John wonders if it is morally permissible to spend some money going to the movies when he could instead use it to buy food for a homeless man. He thinks long and hard about what to do, but cannot come to a conclusion. He decides to ask his friend, Sofia, whose moral judgment he trusts (since every time they disagreed about moral questions in the past he came to admit that she was right). She tells him it is morally permissible for him to go to the movies instead of helping the homeless man. He comes to believe that. When asked why he believes that going to the movies was morally permissible he simply answers: "because Sofia told me so and I trust her".

Peter and John ground their beliefs in reasons of the same kind: both of them ground the conviction that *p* on the fact that someone who is reliable in her claims about the topic in question holds that *p*. I will say of someone who takes this kind of reason as a ground for a moral conviction that he has *outsourced* his moral judgment.[1] Someone *genuinely outsources* his moral judgment when this kind of reason provides the *only* ground for his moral conviction.

---

[1] Outsourcing moral judgments is a way of relying on moral testimony, but there is a point in introducing this technical terminology. Testimony may be a source of knowledge in at least two ways: either the fact that the testifier holds that *p* is taken as a reason to hold the same belief (that is what I call "outsourcing") or reliance on testimony is taken to be a reliable mechanism of belief formation, in which case the epistemic agent may, when confronted, be unable to produce any reason in which to ground his conviction while, nevertheless, being justified in holding it. My argument is concerned with the first kind of testimony-based moral conviction. Whether the second kind of testimony-based moral knowledge is possible is irrelevant to my argument. It

   The intuition that the two examples above are suppose to bring forward is this: while there is nothing wrong with genuinely outsourcing ordinary non-moral judgments (this is something we do on a daily basis), there is something wrong with genuinely outsourcing moral judgments. John's moral conviction seems to be deficient in a way in which Peter's conviction is not.

   If this intuition is accepted, a problem immediately arises for Moral Naturalism as defended, for instance, by David Copp (2007, p.230 and 2017, p.31). Moral Naturalism is the view that (1) moral judgments are ascriptions of moral properties, (2) these properties are ordinary garden-variety natural properties and (3) moral properties correspond to properties ascribed by purely non-evaluative predicates.[2] This view entails that ordinary fact-stating judgments and moral judgments are very closely related – whatever difference there is between them must be explained in terms of differences between the features of reality each of them takes as its subject matter. But then we can ask: if moral judgments are nothing but fact-stating judgments, and there is nothing wrong with the outsourcing of ordinary fact-stating judgments, why is the outsourcing of moral judgments problematic?

   Of course, the naturalist will not go down so easily. He can try to explain this asymmetry by appeal to some peculiarity of moral facts. And, if that proves too hard, he can try to reject the intuition that there is something wrong with moral outsourcing. After all, this intuition is far from uncontroversial. Many have tried to explain it away or to discredit the intuition pumps that bring it forward. Karen Jones, for instance, has argued that we have reason to adopt a default stance of distrust towards moral testifiers and that this explains why we feel there is something fishy about reliance on moral testimony (Jones, 1999, p.72-73). Driver (2006, p.639) argues in similar fashion that the worry about relying on moral testimony is due to the worry that it is harder to identify a reliable moral testifier than it is to identify a reliable testifier on other matters. And Paulina Sliwa has argued that what causes our intuitive unease with problematic cases of moral testimony is not the fact that there is something

---

will be useful, therefore, to use a technical term to refer to kind of reliance on moral testimony relevant to my argument. I discuss this further in section 4.

[2] Clause (3) clearly excludes Sturgeon's Non-reductive Moral Naturalism, which challenges the assumption that "nothing counts as a natural property unless we have some non-ethical terminology to represent it" (Sturgeon, 2003, p.536). In what follows I use "Moral Naturalism" to refer to reductive versions of this view only.

wrong with moral outsourcing itself, but additional features of the examples that are supposed to bring forth that intuition, such as the fact that the agents presented in these examples ignore basic moral truths or the fact that they gladly settle with an opinion on matters that are highly controversial without any inkling of the reasons that ground it (Sliwa, 2012, p.184-88).

It is, therefore, a hopeless enterprise to try to establish that there is something wrong with moral outsourcing by mere appeal to intuition pumps – we quickly come to a stalemate. And, worst yet, even among those that agree that reliance in moral testimony is inappropriate there is no consensus about how to understand this inappropriateness. On a strong interpretation, moral outsourcing is impossible and this impossibility is a distinctive feature of moral judgments: one cannot take the fact that a reliable moral judge holds the moral opinion that *p* as grounds for the conviction that *p;* if one tries to, one will end up with a moral conviction that has no grounds whatsoever. This strong interpretation gives rise to the aforementioned problem for Moral Naturalism – why is there such a restriction on the possible grounds for a moral conviction if moral judgments are nothing but fact-stating judgments? The majority of those convinced that there is something wrong with reliance on moral testimony, however, prefer a weaker interpretation: although possible, it is morally objectionable to ground moral convictions on moral testimony. Philip Nickel, for instance, argues that in order for our actions to have moral worth we must understand the moral claim on which we base our action. To understand a moral claim one has to grasp its justification basis – one must not only know that the claim is true but grasp the reasons why it is true (Nickel, 2001, p.257-8). Obviously, moral testimony cannot provide moral understanding in this sense (for, even if it provides a ground for moral conviction, it gives no insight into the reasons why the outsourced judgment is true). Moral testimony, therefore, cannot provide a basis for morally good actions (Nickel, 2001, p.260). For this reason genuine moral outsourcing is a morally inferior way of coming to moral beliefs (even if these beliefs are correct) – and that is what is wrong with it. Alison Hills adds that the willingness to rely on moral testimony points to a lack of responsiveness to moral reasons; such responsiveness is typical of virtuous persons and the lack of it points, therefore, to a flaw in character (Hills, 2009, p.112). Robert Howell (2014) also claims that what is wrong with reliance in moral testimony is that it results in judgments that do not stem from the agent's virtue.

In this weaker interpretation, moral outsourcing is a perfectly legitimate way of coming to moral convictions in that it does indeed provides a ground for such convictions, and the intuition that there is something wrong with it stems from the fact that it is morally inappropriate or, at least, sub-optimal to form moral beliefs in this way. If this is correct, it poses no problem for Moral Naturalism. The metaethical claim that moral judgments are nothing but fact-stating judgments is by all means compatible with first-order normative claims about the conditions for morally worth actions or about what it takes for someone to be virtuous.

In order to come up with an objection to Moral Naturalism we cannot settle with the weaker interpretation. In what follows I argue that it is possible to vindicate the strong interpretation given some assumptions that are shared by the defenders of Moral Naturalism

Most arguments for the impossibility of moral outsourcing appeal to the close connection between moral judgments and action, to the practical importance of moral decision or to the connection between an ideal of autonomy and morality. Such arguments have been thoroughly criticized and I do not intend to insist on them.[3] My argumentative strategy will be different: Moral Naturalism embodies a particular conception of moral competence (that is, a conception of the conditions one must meet in order to qualify as capable of making moral judgments, acquiring moral knowledge and engaging in meaningful moral debate); I argue, in section 2, that this imposes a constraint on the kind of capacity the naturalist can presuppose in every agent capable of moral judgment: all such capacities must have a role in the regulation of the agent's moral beliefs. I call this the "Doxastic Relevance Constraint". In section 3 I argue that the naturalist has independent reasons to accept that the acknowledgment of moral supervenience is a condition of moral competence. Finally, in section 4, I argue that, given the naturalistic picture of moral competence, genuine moral outsourcing conflicts with the acknowledgement of supervenience. My point is that, given the admissibility of genuine moral outsourcing, competent moral agents may

---

[3] For an exposition of such arguments and objections to them see Jones (1999), Driver (2006, p.635-636), and Hopkins (2007, p.615-26) and Jones & Schroeter (2012).

find themselves in circumstances in which their moral convictions are not compatible with the acknowledgement of moral supervenience. Ascribing a certain mode of reasoning to competent moral agents could prevent this conflict, but the ascription of the mode of reasoning in question violates the Doxastic Relevance Constraint. The naturalist has no option, therefore, but to admit that the ban on moral outsourcing is a distinctive feature of moral judgments. This poses a serious problem for Moral Naturalism, as I argue in section 5. The problem is that while the naturalist must deny the admissibility of moral outsourcing, this denial makes no sense under the supposition that Moral Naturalism is correct.

2. Moral Naturalism and Moral Competence

Even if moral predicates do ascribe natural properties, most competent moral agents are quite capable of making moral judgments and engaging in meaningful moral debate without having any idea how to express the truth-conditions of those judgments in naturalistic vocabulary. If we are lucky enough not to be systematically mistaken in our moral judgments, then the naturalist must admit that the use of moral predicates by competent moral agents is guided by heuristic mechanisms that at least tend to reliably track the relevant natural properties (even when the agent is not aware that his judgments track these properties).

Moral Naturalists usually acknowledge this. Take, for instance, those who, like Boyd (1988, p.195) and Sayre-McCord (1997, p.313), propose a *causal semantics* for moral terms. This semantic theory claims that there are causal mechanisms which govern the use of moral predicates in such a way that our use of a given moral predicate tends to track some natural property (or so the moral naturalist hopes). These mechanisms may include a sense of similarity to paradigms of the moral concept in question (such as paradigms of courage, honesty or cruelty we are exposed to when learning to use those moral concepts), a moral sensibility (conceived of as a set of dispositions to immediately make certain moral judgments in certain circumstances), a set of heuristic modes of moral reasoning etc. Very roughly, if a moral predicate "M" refers to a natural property N,

then our use of M is governed by causal mechanisms in such a way that we tend to apply "M" in response to, and only to, instances of N. Whether or not there is for any given moral predicate a single natural property that causally regulates its use is a question to be answer *a posteriori.*

How exactly the causal relation between instances of a natural property and our use of a moral predicate fixes the reference of the predicate need not concern us here. For our purposes, what is important is that this general picture of a causal moral semantics provides us with a naturalistic image of moral competence. According to Moral Naturalism, moral judgments consist in the ascription of certain natural properties to the object of evaluation. All it takes for an agent to be capable of making moral judgments, therefore, is for her to be capable of applying moral predicates in response to the relevant properties. The causal moral semantics entails that what allow us to make moral judgments in response to the relevant natural properties is the fact that our use of moral predicates is governed by certain causal mechanisms. Implicit in the causal moral semantics is, then, the image of the competent moral agent as someone whose use of moral predicates is governed by these causal mechanisms in such a way that it *could* track, more or less reliably, some natural property, regardless of what the agent thinks about the truth conditions of her moral judgments. In order to vindicate Moral Naturalism, it would have to be shown that our moral judgments and beliefs, guided by these mechanisms, do track certain natural properties.

This picture of moral competence is shared by naturalists who seek to distance themselves from the causal semantics theory as put forward by Boyd. Take Copp, for instance. He defends a semantics for moral predicates according to which the property a given moral predicate ascribes is determined jointly by the speakers' semantic intentions, the speakers' interests and the nature of the stuff the speakers use in ostensive explanations of the meaning of the predicate (Copp, 2007, p.232-33).[4] He notices, however, that our moral practice is at odds with Moral Naturalism. We usually take ourselves to have moral knowledge that p (a moral proposition that Moral Naturalism claims to have truth-conditions q) even though our belief that p is not based on the belief

---

[4] See also Copp (2000). For a similar semantic theory for moral terms see Brink (2001, p.170-76).

that *q* or even when we do not believe that *q* (Copp, 2007, 71-2). His answer to this apparent problem is to adopt a reliabilist epistemology in which a belief produced by a reliable belief-generating mechanism counts as knowledge, even if we have no inkling what its truth-conditions are and whether they obtain (Copp, 2007, 74-5). The reliable moral-belief-generating mechanism consists in the set of behavioral and judgmental dispositions (that is, a moral sensibility) that we come to have as a result of the process of moral education and in the general moral rules (such as "torture is wrong") we inculcate and can guide our moral reasoning (Copp, 2007, 76-9).

Copp agrees that agents are usually capable of acquiring moral knowledge and engaging in meaningful moral reasoning without having any inkling of the truth conditions of moral judgments and relying only on a proper moral sensibility and suitable patterns of heuristic reasoning. He agrees, therefore, that ordinary competent moral agent is just someone whose use of moral predicates is governed by appropriate moral-belief-generating mechanisms.

And this is not a mere peculiarity of Copp's theory. That is how naturalists *must* conceive of the competent moral agent. Whatever natural property *N is* specified as the referent of a given moral predicate "*M*", a moral agent may be a competent user of "*M*" while ignoring the fact that an object of evaluation is *M* if, and only if, it is *N* (if Moral Naturalism is correct, then this is the case for most of us). But, for an agent to be able to make *M*-judgments in such a way as to have a shot at moral knowledge, she must exhibit at least a tendency to apply the predicate "*M*" in response to *N*. And how is that possible if she does not know that "*M*" refers to *N?* The answer must be that it is possible only if her use of "*M*" is governed by suitable mechanisms (like a set of judgmental dispositions and modes of heuristic reasoning) in such a way that it tends to track *N*.

Not even those who, like Frank Jackson, defend an analytic version of Moral Naturalism would deny this. Jackson claims we can, by means of conceptual analysis, specify purely descriptive truth-conditions for moral judgments. His approach starts by positing what he calls "folk morality": "[...] the network of moral opinions, intuitions, principles, and concepts whose mastery is part and parcel of having a sense of what is right and wrong, and of being able to engage in meaningful debate about what ought to be done" (Jackson, 1998, p.130). Mastery of folk morality is here described as a condition of moral competence. Competent moral agents share, then, a great deal of moral opinions, principles and intuitions about particular cases, and understand

a number of connections between these convictions and between different moral concepts. The meaning of moral terms, Jackson claims, is given by their place on this network (Jackson, 1998, p.131). This meaning may be analyzed in purely descriptive terms by applying the following method of defining theoretical terms. Let "mature folk morality" be the network of moral opinions, concepts etc. on which current folk morality will converge under critical reflection, that is, if subjected to usual methods of moral reasoning and critique (these, we can presume, are just the modes of reasoning and debate the capacity to engage with is a condition for moral competence – such as the capacity for subsuming cases under general principles or to engage in exercises of reflective equilibrium). Jackson claims the mature folk morality can be represented as a very long conjunction of moral claims. The first step on the analysis is to write every moral predicate in this conjunction in "property name style" (Jackson, 1998, p.140). For example, "Killing someone is typically wrong" becomes "Killing typically has the property of being wrong". The conjunction of all this claims captures the network of relations between moral properties. We can represent this network as a relational predicate true of all the moral properties in question: **M**[r w t …], where "r" stands for the property of rightness, "w" for wrongness and so on. The second step is to replace each distinct moral property predicate by a distinct variable: **M**[x y z …]. The property of being right, for instance, can now be defined as the property that fits the x-role in **M**[x y z …]. To be sure, any property that fits that role will be the property of being right, so we should add a clause of uniqueness to this definition. But, leaving these details aside, the idea is quite clear – we can explain, for instance, what is for an action A to be right as follow: A is right if, and only if, there is a property x such that **M**[x y z …] an A is x. And, since all occurrences of moral predicates on the right-hand side of this bi-conditional have been replaced with variables, we have thus provided, by means of mere analysis, a purely descriptive account of the truth-conditions of the moral judgment in question.

Jackson adds that, given the supervenience of the ethical on the descriptive, we can be sure there are natural properties that fill the roles that define moral predicates. For our purposes, however, there is no need to reconstruct further Jackson's argument. It should be clear that, for Jackson, what characterizes a competent moral agent is the mastery of folk morality, that is, the sharing of certain moral convictions, of dispositions to have certain intuitions in particular cases and an inkling of the complex relations between moral concepts, moral opinions and other non-moral convictions (which underlies the capacity for

moral reasoning). If it does turn out that there is a single natural property that fits the role a given moral concept has in mature folk morality, that is because our use of the moral concept in question happens to respond to that property. Folk morality (the network of moral concepts) is what, in Jackson's approach, pinpoints the natural property a given moral predicate designates; and folk morality is just a function of how we use moral concepts and think about morality.[5]

Naturalists share, therefore, a conception of moral competence: the competent moral agent is someone that, through the usual methods of moral teaching and training, has inherited a number of moral convictions, a reasonable moral sensibility (which leads to reasonable moral intuitions) and suitable modes of moral reasoning.

This conception of what it takes for one to be a competent moral agent establishes a restriction on the kind of capacity the naturalist may put forward as a condition of moral competence. According to this conception, being morally competent is just a matter of being equipped with modes of response and heuristic reasoning that determine one's moral judgments and beliefs. Because these moral-belief-generating mechanisms determine (or at least figure in the process of determination of) one's judgments, we can explain by appeal to them how an agent's moral judgments and beliefs may tend to track a given natural property regardless of what the agent thinks about the truth-conditions of such judgments. If there are some capacities the possession of which is a condition of moral competence, these capacities must then consist in the kind of modes of response and reasoning which can determine one's moral judgments and beliefs. That being the case, we cannot challenge the claim that a moral agent is capable of making moral judgments and engaging in meaningful moral debate on the grounds that she lacks a particular mode of response or reasoning which plays *no role* in the determination of one's moral judgments and beliefs. If a mode of response or reasoning plays no role in the processes by means of which we come to acquire, modify or abandon moral convictions, then, given the naturalistic conception of moral competence, lacking this particular mode of response or reasoning cannot be an obstacle for moral competence. Because this particular mode of response or reasoning is

---

[5] Jackson (1998, p. 137) admits that there may be no convergence in moral opinion, in such a way that there will not be a single mature folk. This is a scenario in which differences in use establish different folk moralities.

idle with respect to the determination of one's moral convictions, one can have all the moral-belief-generating mechanisms necessary for one to qualify as a competent moral agent even if one lacks it. We can be sure, for instance, that the fact one lacks a disposition to cry when exposed to an instance of cruelty provides no ground for denying one's moral competence because this disposition has no impact on your moral convictions; the disposition to make a negative moral judgment when confronted with cruelty, on the other hand, has an obvious impact on our moral convictions and one could perhaps argue that to have such a disposition is a condition of moral competence. The naturalist must comply, then, with the following constraint:

> DOXASTIC RELEVANCE CONSTRAINT: any mode of response or mode of heuristic reasoning the naturalist ascribes to every competent moral agent (as a condition of moral competence) must play a role in the determination of the agent's moral convictions.

The point here is merely this: if moral competence is a matter of being equipped with a suitable set of belief-generating mechanisms, one's moral competence cannot be precluded by the lacking of a mode of response or reasoning which has no role in one's doxastic economy. Any disposition or mode of heuristic moral reasoning the naturalist may want to include in the common repertoire of every competent user of a given moral predicate (as a condition of moral competence) must better not be idle with respect to the production or modification of our moral convictions.

3. The Supervenience Constraint

Consider the following case:

> Supervenience Violation: John occupies a politic position in a government agency aimed at providing financial support to families in a very poor region of his country. He is reasonably rich. He uses his position to steal from the agency's funds. Sarah claims that what John has done is morally wrong. She is aware that Peter occupies a similar politic position, in a similar agency, is equally rich and also uses his position to steal from the agency's funds. Sarah believes, however, that John's actions are more reprehensible than Peter's actions. Even after we

call her attention to the fact that there is no morally significant difference between the cases, she still holds that they should be judge differently. If we insist on asking what the difference between these cases is, she claims there is no difference besides the fact that one is, from a moral point of view, more reprehensible than the other.

Sarah's behavior may indicate that she is biased in her judgment: perhaps she has some kind of personal relation with Peter or just happens to be more sympathetic to him. But suppose she is not acquainted with John nor with Peter and that she is not aware of any fact about any of them that could influence her judgment in an inappropriate, perhaps subliminal, way. If that is the case, then there is no option but to conclude that Sarah does not realize what she is saying or that she is simply morally incompetent: incapable of making moral judgments and certainly incapable of engaging in meaningful moral debate.

The point of this example is to show that it is a condition of moral competence to acknowledge that every moral distinction must be based on a non-evaluative distinction or, to put it another way, that the moral difference between two objects of evaluation cannot be the only difference between them. This feature of moral value is usually referred to as the Supervenience of the Moral on the Non-moral.

Since the acknowledgement of Moral Supervenience is a condition of moral competence it imposes a constraint on any proper moral judgment. From the point of view of the moral agent this constraint can be formulated as follows:

> SUPERVENIENCE CONSTRAINT: every moral judgment must be grounded on non-moral considerations about the object of evaluation; if the same non-moral considerations are available in other circumstances, then, *ceteris paribus*, they ground an equivalent moral judgment.[6]

---

[6] The *ceteris paribus* clause is necessary for the non-moral considerations which ground a moral judgment may be context sensitive, i.e., they may be available in other circumstances but fail to provide grounds for a similar moral judgment in virtue of other features of the context. The Supervenience Constraint still entails, however, that if the agent believes he is faced with two exactly similar situations, then his moral judgments about both of them must be equivalent, and grounded on equivalent reasons.

In order to comply with this requirement the moral agent need not be able to make explicit the reasons that ground any moral distinction she is willing to make. Sarah can insist on drawing a moral distinction between the actions of John and Peter if she admits there is a non-moral difference between them, even if she cannot quite put her finger on what it is.

Now, this constraint fits perfectly well with Moral Naturalism. Indeed, Moral Naturalism entails a stronger claim than Moral Supervenience. Given Moral Supervenience, it follows that every moral judgment must be grounded on some non-moral feature of the object of evaluation. It does not follow that every application of the same moral concept (every judgment to the effect that an action is morally wrong, for instance) must be based on a non-moral feature that is shared by all the instances of the concept. As far as Moral Supervenience alone is concerned, it could be the case that there is no common feature on which every correct judgment to the effect that an action is morally wrong could be grounded. But, if Moral Naturalism is correct, it entails that there is such a common feature, namely, the natural property the moral predicate in question tracks. Moral Naturalism, then, entails Moral Supervenience and more.

One must not forget, however, that the competent moral agent, according to the naturalist's conception of moral competence, is not someone who has an insight into the nature of moral judgment or, for that matter, any inkling what the truth-conditions of moral judgments are. Even if moral judgments are ascriptions of natural properties, that does not explain, by itself, why competent moral agents comply with the Supervenience Constraint. That must be explained in terms of the capacities that, according to the naturalist, allow us to make moral judgments that, as matter of fact, track natural features of reality.

That explanation is not hard to come by: the competent moral agent is someone equipped with a moral sensibility and heuristic modes of moral reasoning; she is, thus, someone who reacts to certain non-moral features of situations by coming to a certain moral judgment and that is able to reason from these moral judgments to other moral conclusions; if two objects of evaluation present the same non-moral features and no morally significant difference in context, then these features will trigger the same judgmental dispositions and the same modes of reasoning; therefore, the moral agent will come in both cases to the same conclusion. Of course, one can, for instance, judge an action A to be M (where M is a moral predicate) and, when confronted with another action B, fail to see that it shares the same the features of A that led to the judgment that "A is M" and judge B differently – but this is

a case of moral error due to a failure of perception, not a case of violation of the Supervenience Constraint (at least if the agent is willing to revise her judgment once her attention is drawn to the fact that there is no morally relevant difference between *A* and *B*).

This explanation fits well with the Doxastic Relevance Constraint, for it makes reference only to dispositions and modes of heuristic reasoning that have an impact on the agent's moral convictions. Furthermore, dispositions and modes of reasoning in question are exactly the kind that can explain how the agent's use of a moral predicate may happen to reliably track some natural property. If the judgmental dispositions and heuristic modes of reasoning that constitute moral competence are to ensure that our moral judgments respond to certain natural facts, they must be regular, in the sense that they are likewise triggered in likewise conditions. If that were not the case, the processes that lead to the application of a given moral predicate could be triggered in random conditions or fail to be triggered in similar conditions – and that would make it very unlikely that different applications of the same moral predicate were, for the most part, responses to one and the same natural property. Our moral sensibility and our modes of moral reasoning must, then, be regular if they are to allow us to track moral facts and it is this very regularity that explains why competent moral agents comply with the Supervenience Constraint.

This explanation does not involve ascribing to the competent moral agent an explicit formulation of the Supervenience Constraint nor the willing to comply with it. Instead, the agent's acknowledgement of this constraint is embodied in her moral sensibility and in her capacity for moral reasoning, i.e., in that which underlies her capacity for moral judgment. This is an advantage for Moral Naturalism for it does not presuppose that moral competence involves insight into (and some level of explicit formulation of) the nature of moral judgment, its truth-conditions or the formal constraints that govern it – thus avoiding an overly intellectualized conception of moral competence.

4. Moral Outsourcing and the Supervenience Constraint

We can now show that the admissibility of moral outsourcing conflicts with the Supervenience Constraint, at least on the framework of Moral Naturalism.

To outsource a moral judgment is to ground a moral *conviction* on a specific kind of reason, namely, it is to ground a moral conviction that *p* on the fact that

someone whose moral judgments are reliable holds that *p*. What exactly it takes for someone to be a reliable moral judge need not concern us here – certainly we recognize some people as moral sages (that is how Christians see Jesus, for example) and believe them to be much more likely to make correct moral judgments (so much so that we take the fact that we disagree with them as indication that we are probably mistaken in our judgment). I am going to argue that even if we can identify moral sages we cannot take the fact that they hold that *p* as a ground for holding the same moral conviction.

Controversies surrounding the admissibility of moral testimony usually focus on the question whether moral testimony can lead to moral knowledge.[7] By formulating my claim in terms of outsourcing I sidestep this debate. The claim that moral outsourcing is impossible is a normative claim about what kind of reason can be offered as a legitimate ground for a moral conviction. As such, it is perfectly compatible with the claim that we can obtain moral knowledge by means of moral testimony. Surely we simply inherit a lot of moral convictions from those responsible for our moral education and many of these convictions amount to moral knowledge. We certainly know that killing innocent people is morally wrong but many people have simply inherit this belief and may be incapable of producing a reason on which to ground it – they simply take it to be a manifest truth. This provides a good example of a piece of moral knowledge that is obtained by means of moral testimony alone, but it does not undermine the claim that moral outsourcing is impossible. We may, as Copp does, adopt a reliabilist moral epistemology and hold that someone can have moral knowledge that *p* even if he is unable to produce any reason on which to ground his conviction that *p,* given that the conviction that *p* is the product of a reliable mechanism for belief formation. If moral testimony (at least when certain conditions are met) is a reliable way to form moral beliefs, then we can admit that someone who simply inherits the conviction that killing innocent people is morally wrong knows that to be the case and yet refuse to accept the

---

[7] See Jones (1999), Driver (2006) and Hopkins (2007). Hopkins concedes that moral testimony may lead to moral knowledge, but argues that forming moral beliefs on the basis of moral testimony violates some other, non-epistemic norm of moral beliefs formation (Hopkins, 2007, 629). The claim I am going to defend is not about the formation of moral beliefs. I am going to defend that a particular kind of reason cannot ground moral convictions, regardless of how we come to such convictions. This claim is compatible with Hopkins's thesis, but independent of it.

fact that this conviction is shared by a moral sage (or by everyone) as providing a reason that could legitimately ground it.

The ban on moral outsourcing is also compatible with the outsourcing of the non-moral information on which a given moral judgment is grounded. For instance, I can ground the judgment that "eating animals is wrong" on the fact that animals suffer under modern farming methods; the fact that I ground the conviction that modern farming methods impose suffering on the fact that someone else holds it does not make it inappropriate to ground the moral judgment in question.

That said, it is indeed true that moral outsourcing is a way of relying on moral testimony. And when people give moral testimony they usually provide some explanation of the reasons why they believe the moral judgment in question is correct. For instance, if I ask someone whether eating animals is wrong, the answer is not usually a mere "yes" or "no", but something like "yes, it is wrong because animals suffer under modern farming methods". The ordinary moral testimony comprises not just moral information ("eating animals is wrong"), but non-moral information ("animals suffer under modern farming methods") and a claim to the effect that the provided non-moral information explains why the moral judgment in question is correct. Surely, if I trust the person I am consulting, I can outsource the non-moral information she provides. If I am to *genuinely* outsource the moral conviction that eating animals is wrong, however, the grounding relation between the outsourced non-moral information and this moral judgment cannot be intelligible for me. If I understood how the fact that animals suffer under modern farming methods explains why the claim that it is wrong to eat animals is correct, then, given that I came to believe that animals suffer under modern farming methods, I could ground my moral conviction directly on that piece of information, and not on the fact that the person I am consulting holds that eating animals is wrong (or, at least, not exclusively on that fact). And if I do that, I have not genuinely outsourced my moral judgment – I have been, to some extent, convinced by moral argument.

In genuine cases of moral outsourcing, therefore, the grounding relation between the reasons the moral testifier offers and the outsourced moral claim must remain unintelligible for the moral agent.[8] Nevertheless, if the moral

---

[8] In this sense, genuine cases of moral outsourcing are what McGrath calls "pure cases of moral deference". See McGrath (2011, p.114).

testifier has offered some explanation of the reasons why his moral judgment is correct, the moral agent may come to believe that the outsourced moral belief is correct for just those reasons. That is, he may come to believe not just that "*p"* but that "*p* is the case because such and such is the case". And if one admits that genuine moral outsourcing is possible, it is hard to see how one could deny that we can also outsource judgments about why a given moral claim is correct. A moral agent can, then, have not only a suitably grounded conviction that eating animals is wrong, for instance, but a suitably grounded conviction that it is wrong because such and such is the case. Nevertheless, if he has genuinely outsourced the moral judgment and its explanation, then this suitably grounded conviction about *why* eating animals is wrong does not amount to an *understanding of why* it is wrong.[9] That is, while he can enumerate what it is about eating animals that makes it wrong, he does not understand the relation between the judgment and the reasons that make it correct (otherwise, this would not be a genuine case of moral outsourcing).

Hence, if genuine moral outsourcing is possible, one can have a suitably grounded moral conviction that *p*, a suitably grounded conviction about why *p* is the case and yet fail to understand why *p* is the case. In genuine cases of moral outsourcing the agent's answers to the questions "why do you believe that *p?*" and "why is *p* the case?" come apart.

How does the admissibility of genuine moral outsourcing conflicts with the Supervenience Constraint? Suppose David outsources the moral judgment that "Joseph is charitable". The testimony he receives includes an explanation of why that judgment is correct: "because Joseph is always willing to help those in need, even at his own expense, and does that in virtue of a genuine concern for the well-being of others". Since this is a genuine case of moral outsourcing, the grounding relation between these considerations and the judgment that Joseph is charitable are not intelligible to David.

---

[9] Notice that this claim does not depend on any supposition about what is required for moral understanding, but simply follows from the supposition that the agent has genuinely outsourced his moral judgment. In particular, this claim does not depend on the previous assumption that understanding why *p* is not reducible to knowing why *p*, as defended by Hills (2009, p.100-04) - this thesis has been recently challenged by Sliwa (2017). The point is rather that, given the supposition that the agent has genuinely outsourced his judgment, whatever he knows is not enough for him to understand why the judgment is correct.

Now suppose David comes across Eleanor. So far as he can see, Joseph and Eleanor are exactly alike with respect to every feature relevant to the claim that they are charitable: Eleanor is always willing to help those in need, even at her own expenses, and does that in virtue of a genuine concern for the well-being of others. And there is no contextual difference between the cases that could prevent those features from providing a reason for making an equivalent judgment: Eleanor does not have any morally spurious motives to help others (such as a desire for social prestige), she is willing to sacrifice as much time and resources as Joseph in helping others, she treats the beneficiaries of her charity with respect etc.

Given the Supervenience Constraint, and given that David is a competent moral agent, since he can see both cases are exactly similar with respect to every morally relevant feature, he must admit that if his conviction that "Joseph is charitable" is suitably grounded, then the conviction that "Eleanor is charitable" is likewise grounded.

Now, supposing that genuine moral outsourcing is possible, David's conviction that "Joseph is charitable" is suitably grounded. What grounds could he offer for the equivalent conviction that "Eleanor is charitable"? Surely, he cannot ground it on the same fact on which he grounded the conviction about Joseph, namely, the fact that a reliable moral testifier holds that Joseph is charitable – for that is completely unrelated to the claim that Eleanor is charitable. And while he can enumerate which features of Joseph's character explain why Joseph is charitable, he cannot ground the conviction that Eleanor is charitable on these same features since for that to be possible the grounding relation between these features and the moral judgment in question would have to be intelligible for him (and, given that he has genuinely outsourced his judgment this cannot be the case).

Any reason he can find on which to ground the claim that "Eleanor is charitable" will have to be based on an inference: (i) Joseph is charitable; (ii) there is no morally significant difference between Joseph's and Eleanor's character; (iii) therefore, Eleanor is charitable.

This strategy to reconcile the admission of moral outsourcing with the respect for the Supervenience Constraint is open to the moral naturalist. Ascribing this heuristic mode of reasoning to all competent moral agents does not violate the Doxastic Relevance Constraint, for a disposition to reason from non-moral similarities to similar moral conclusions can have an important role in regulating our moral convictions. And the pressure for coherence this

disposition generates may help explain how our judgments tend to track a particular natural property. However, this mode of reasoning is not by itself sufficient to guarantee compliance with the Supervenience Constraint once moral outsourcing is admitted.

The Supervenience Constraint entails not only that moral agents must judge exactly similar cases the same way, but also that competent moral agents must acknowledge that if the same non-moral considerations are available in two cases, then, other things being equal, they ground equivalent moral judgments. That means that in order to comply with the Supervenience Constraint, David not only needs to find a ground for the conviction that Eleanor is charitable (which is provided by the argument (i)-(iii) above), but must also acknowledge that Eleanor is charitable *for exactly the same reasons* Joseph is. And given that he has a suitably grounded conviction that "Joseph is charitable *because* he is willing to help those in need...," he must acknowledge that Eleanor is charitable for just those reasons. Now, since the grounding relation between those reasons and the moral judgment they ground is not intelligible for David, he cannot *just see* that Eleanor is charitable *because* she is willing to help those in need etc. – he will have to reason his way to this conclusion as well.

In order to comply with the Supervenience Constraint David will have to reason along the following lines: (i') Joseph is charitable because [r] (where "[r]" stand for an explanation of the reasons why Joseph is charitable); (ii') Eleanor is charitable as well; (iii') [r] applies to Eleanor as well (and there is no morally significant difference between Eleanor and Joseph); (iv') therefore, Eleanor is charitable because [r]. Notice that (i') and (iv') are expression of a *mere* piece of information, in the sense that although David can enumerate the reasons that constitute [r], and thus provide an answer to the question "why is Eleanor charitable?," he cannot understand the grounding relation between [r] and the judgment that someone is charitable. When he answers that question he is merely repeating a piece of information he has received (just as I can say that water expands when frozen because ice is formed by H_2O molecules arranged geometrically in the form of crystals, even though I do not understand how the later explain the former). This is not, therefore, a piece of reasoning that increases his moral understanding or leads him to make new moral judgments. This is a piece of reasoning about moral *reasons* that is completely disconnected from the production or revision of any moral conviction. And as such, ascribing this mode of reason to *all* competent moral agents (as a condition of moral competence) violates the Doxastic Relevance Constraint.

Of course we sometimes reason from moral convictions to moral reasons, in the sense that we engage in reasoning in order, for instance, to disclose which features of an action that we believe to be wrong explain why it is wrong. This kind of reasoning can have an impact on our moral convictions and the naturalist can provide an explanation of how it contributes to our capacity to track moral facts: disclosing the reasons that ground a moral claim can promote the coherence of our moral convictions, for once we know why a moral judgment is correct we can use that information to further the discussion surrounding controversial moral topics or to revise established moral opinions; and coherence is indispensable if our moral judgments are to regularly track one and the same natural property. But the mode of reasoning illustrated in (i')-(iv') does not promote coherence is this way. In the first place, it does not lead to the conclusion that the fact that [r] applies to someone can explain why this person is charitable; rather, that claim is presupposed by premise (i'). In the second place, the moral agent cannot gain any moral understanding by coming to the conclusion that (iv') – how [r] explains why Eleanor is charitable is as unintelligible for him as it was before, the only difference is that now he can enumerate the reasons why Eleanor is charitable. In the third place, the conclusion (iv') does not lead to any new judgment, for an argument of this form can only specify the reasons why a moral judgment, such as (ii'), we already admit to be correct (perhaps in virtue of a line of reasoning such as (i)-(iii)) is correct.

The point here is not that this form of reasoning is somehow impossible. A competent moral agent may explicitly acknowledge its validity when thinking theoretically about morality. The point is rather that this form of reasoning is absolutely *useless* – it has no practical application. The naturalist cannot include the disposition to this kind of reasoning in the repertoire of modes of reasoning common to every competent moral agent. Given the naturalistic conception of the competent moral agent, the lack of a disposition to engage in reasoning of the form (i')-(iv'), which has no impact over one's moral beliefs, cannot preclude one's moral competence. Someone to whom this kind of reasoning never occurs (or someone who can see no point in engaging in it – since it leads neither to new moral beliefs nor to the disclosure of moral reasons nor to an increase in moral understanding) can be a perfectly competent moral agent, capable of making moral judgments and engaging in meaningful moral discussion. And thus, if the possibility of genuine moral outsourcing were admitted, perfectly competent moral agents could, at least in principle, violate

the Supervenience Constraint. Since compliance with the Supervenience Constraint is a condition for moral competence, the naturalist has no option but to deny the admissibility of genuine moral outsourcing.

5. The Problem for Moral Naturalism

It should be clear enough that the inadmissibility of genuine moral outsourcing the naturalist is forced to acknowledge is not a moral inadmissibility. It is not the case that the fact that a moral sage holds that *p* (where "*p*" is a moral claim) can ground the conviction that *p* but that we should not settle with a conviction grounded in this way. What the naturalist is forced to admit is that the fact that a moral sage holds that *p cannot* ground the conviction that *p*, in such a way that if one can offer no ground for the conviction that *p* other than the opinion of the moral sage, then one has a groundless moral conviction.

How could the naturalist explain this? If moral judgments are just fact-stating judgments, why should it not be possible to ground the moral conviction that *p* on the fact that someone who is better at tracking the relevant facts holds that *p*? Fact-stating judgments can usually be outsourced, and, if moral judgments are just this kind of judgment, moral outsourcing should be possible as well.

The naturalist could claim that not all fact-stating judgments can be outsourced: when it comes to highly controversial topics, appeals to expertise or testimony are problematic. One cannot, for instance, outsource the judgment that the Riemann Hypothesis is true simply because this is a judgment about a controversial issue. The naturalist could, then, point out that morality is a domain in which disagreement is pervasive and claim that this explains the inadmissibility of moral outsourcing.[10]

This strategy will not work. Notice first that, while it is indeed inadmissible to outsource some mathematical judgments (or, for that matter, some judgments about the impact of an economic policy) simply because those judgments are controversial, that does not entail that we cannot outsource mathematical (or economic) judgments in general. One can outsource the judgment that the formula "$4/3 \, \pi \, r^3$" allows one to calculate the volume

---

[10] Thanks to an anonymous referee for bringing this point to my attention.

of a sphere (indeed, it is very likely that most high-school students can offer no ground for such a conviction besides their teachers testimony) just as one could outsource the judgment that "putting a ceiling on rents reduces the quality of available housing" (which is agreed upon by the majority of economists). The fact that some judgments about a subject matter are controversial is compatible with the possibility of outsourcing judgments about that subject matter. Even if morality is plagued with disagreement, there could still be a core of less controversial judgments – and the naturalist would still owe us an explanation of why those judgments cannot be outsourced.

Perhaps the naturalist could argue that there is no intermediate class of moral judgments in between those judgments whose truth is evident to every competent moral agent and those which are controversial to such an extent that they cannot be outsourced. However, I do not see how, given the naturalistic image of moral judgment, one could back this claim up. Given that moral judgments ascribe natural properties and that some people may be better at tracking the relevant properties, it is quite clear that some people may be better moral judges than others. If these people came to agree upon a moral judgment whose truth is not evident to an otherwise competent moral agent and if this agent cannot properly assess the considerations put forward in favor of the judgment, why couldn't he outsource this particular judgment? The naturalist should provide an answer to this question, but it is not clear which answer he could provide.[11] This situation would be analogous to that in which someone with a reasonable understanding of mathematics, capable of understanding a mathematical theorem but incapable of assessing its proof, grounds the judgment that the theorem is correct on the fact that mathematical experts agree that it is correct; or to that in which someone capable of understanding a sociological hypothesis but incapable of assessing the complex sociological evidence that back it up grounds the

---

[11] Copp, for instance, far from explaining why outsourcing should not be possible in this case, is willing to admit the possibility of moral expertise and the possibility of relying on the testimony of moral experts. See Copp (2007, 82-83).

judgment that the hypothesis is correct on the fact that sociologists that were trained to assess this kind of data agree in declaring the hypothesis correct. This is the kind of situation Moral Naturalism makes room for and, indeed, the kind of situation one would expect given the naturalist's approximation of moral theory to the status of a special science.[12]

But suppose we were willing to admit that every possible moral judgment one could find occasion to outsource is controversial – could the naturalist rely on this claim to explain the inadmissibility of moral outsourcing? No, for the success of Moral Naturalism depends on the claim that the community of users of moral terms *can* reach a significant degree of agreement in moral judgment.

This becomes clear when we consider the variety of Moral Naturalism that relies on a causal semantics for moral terms. According to this approach, the natural property a given moral predicate refers to is the moral property that causally regulates its use. If this variety of Moral Naturalism can be vindicated, then it must be the case that the use of a given moral term by competent moral agents in a community is regulated by a particular natural property (otherwise, the predicate would have no determinate meaning). And, if the use of a moral predicate is causally regulated by a natural property, the agents that use this predicate will tend, over time, to agree in their judgments. That does not mean that the naturalist will deny the evident fact that there is significant moral disagreement, but that he commits himself to the claim that, if Moral Naturalism is attainable, then our modes of moral reflection will (if properly carried out) lead to a significant convergence in moral opinion.[13]

---

[12] This approximation is quite clear in Brink (1984) and Boyd (1988). Copp is explicit in claiming that his theory implies that we can settle moral questions by appeal to complex sociological evidence (Copp, 2007, p.81).

[13] Boyd states that "a term t refers to a kind (property, relation, etc.) k just in case there exist causal mechanisms whose tendency is to bring it about, over time, that what is predicated of the term t will be approximately true of k" (Boyd, 1988, p.195). Brink (1984, p.116-117) and Sayre-McCord (1997, p.323) clearly indicate that the vindication of Moral Naturalism depends on possibility of reflective equilibrium leading to a convergence in moral opinion.

The naturalist could then appeal to the current disagreement in order to explain why moral outsourcing is problematic *now,* but he presupposes that it is possible to resolve a great deal of disagreement – and once a greater convergence is achieved the naturalist will not be able to explain the inadmissibility of moral outsourcing.

The same is true of Jackson's analytic approach, which presupposes that folk morality will converge, under critical reflection, on a "mature folk morality". What about Copp? He does not commit himself to a causal semantics for moral terms and there is no evident sense in which the success of his naturalistic semantics for moral terms is tied to the possibility of a convergence in moral opinion. It should be noted, however, that he is not in any way inclined to declare moral outsourcing problematic. He actually claims that it "should not be surprising" that there are moral experts and that some people defer to them (Copp, 2007, p.82) and that the fact that his theory implies that there can be moral expertise is not an objection to it (Copp, 2007, p.83). We can understand why he does not believe that the current pervasive moral disagreement renders moral outsourcing inadmissible if we recall his strategy to account for the possibility of moral knowledge. According to Copp, we are capable of acquiring moral knowledge because our moral beliefs are the product of reliable belief-generating mechanisms. If these mechanisms are to generate knowledge, they must link our moral beliefs to what are actual moral facts (Copp, 2007, p.83-84). And if that is the case, then moral facts causally regulate our moral beliefs, via these belief-generating mechanisms. If moral knowledge is at all possible, our moral beliefs should tend to, over time, approximate moral facts – and, thus, we should expect convergence in moral belief.[14]

Therefore, even if it were possible to account for the inadmissibility of moral outsourcing by appeal to the idea that moral disagreement is pervasive, naturalists could not adopt that strategy because they are committed to the claim that convergence in moral belief is possible and forthcoming.

---

[14] Copp actually has a story to tell about the mechanisms that tend to lead to convergence in moral opinion over time. See Copp (2007, p.84-86).

How else could the naturalist accommodate the inadmissibility of moral outsourcing? Given that, according to the naturalist, moral judgment is just a species of fact-stating judgments, any distinctive traits of moral judgments should be explained not in terms of a special semantic function, but in terms of the kind of fact these judgments track. After all, the distinction between moral judgments and other kinds of descriptive judgments derives exclusively from differences between the features of reality they aim at describing.

The ban on moral outsourcing cannot, however, be accounted for in terms of the usual candidates for natural moral properties. Moral Naturalists tend to favor consequencialist conceptions of moral properties. Boyd and Brink conceived of goodness as a functional property characterized in terms of its tendency to foster a number of human goods, such as physical, psychological and social needs (Boyd, 1988, p.203) or the maintenance and flourishing of human organisms (Brink, 1984, p.121-2). But if that was correct, being able to calculate the impact of, say, an action on these needs would be enough for someone to be able to make this kind of judgment. Obviously, some people may be much better than others in performing such calculations. If someone has proved to be much better than I am in making these calculations (if we have ascertained that his predictions about the impact of particular actions on general well-being are always more accurate than mine), why couldn't I ground my conviction that a given action will promote the maintenance of human life on that fact that he holds that opinion? I could; just as I can ground my convictions about tomorrow's weather on the weather forecast. There is no reason not to outsource the consequencialist calculation. And if moral judgments were reducible to the claim that an action tends to have such an effect, the ban on moral outsourcing would be unintelligible.

Nor does Copp's claim that the moral predicate "correct" ascribes to the action evaluated the relational property of "being required by the code of rules, whatever it is, the currency of which in the society in question would best contribute to the society's ability to meet its needs" (Copp, 2007, p.238) fares any better in this respect. If that claim were correct, then being able to determine which set of rules would best contribute to a society's ability to meet its needs and which actions are required or prohibited by that set of rules would be enough for someone to be able to make moral judgments about actions. But surely we can outsource judgments about the impact of certain social norms on the possibility of peaceful cooperative interaction among the members of a given society – to mention one social need Copp has in mind (Copp, 2007,

p.237). We do sometimes outsource questions of this kind to social scientists or economist.

The problem for Moral Naturalism is, then, simple but serious: the naturalist must deny the admissibility of genuine moral outsourcing, but the ban on moral outsourcing makes no sense under the assumption that Moral Naturalism is the correct characterization of moral judgments. If no answer to this problem can be provided, we have a good reason to reject Moral Naturalism.

Refeences

BRINK, D. "Moral realism and the sceptical arguments from disagreement and queerness". *Australasian Journal of Philosophy*, 62, pp. 111-125, 1984.

_____ "Realism, Naturalism, and Moral Semantics". *Social Philosophy and Policy*, 18, pp. 154-176, 2001.

BOYD, R. "How to Be a Moral Realist". In G. Sayre-McCord (ed) (1988), pp.181-228

COPP, D. "Milk, Honey, and the Good Life on Moral Twin Earth". *Synthese*, 124, pp. 113-137, 2000.

_____ *Morality in a Natural World*. New York: Cambridge University Press, 2007.

_____ "Normative Naturalism and Normative Nihilism". In S. Kirchin (ed) (2017), pp. 28-53.

DRIVER, J. "Autonomy and the asymmetry problem for moral expertise". *Philosophical Studies*, 128, pp. 619-644, 2006.

HILLS, A. "Moral Testimony and Moral Epistemology" *Ethics*, 120, pp. 94-127, 2009.

HOPKINS, R. "What is Wrong With Moral Testimony?". *Philosophy and Phenomenological Research*, 74, pp. 611-634, 2007.

HOWELL, R. "Google Morals, Virtue, and the Asymmetry of Deference". *Noûs*, 48, pp. 389-415, 2014.

JACKSON, F. *From Metaphysics to Ethics*. New York: Oxford University Press, 1998.

JONES, K. "Second-Hand Moral Knowledge". *The Journal of Philosophy*, 96, pp. 55-78, 1999.

_____ & SCHROETER, F. "Moral Expertise". *Analyse & Kritik*, 34, pp. 217-230, 2012.

Kirchin, S. *Reading Parfit On What Matters*. New York: Routledge, 2017.

MCGRATH, S. "Skepticism about Moral Expertise as a Puzzle for Moral Realism". *Journal of Philosophy*, 108, pp. 111-137, 2011.

NICKEL, P. "Moral Testimony and its Authority". *Ethical Theory and Moral Practice*, 4, pp. 253-266, 2001.

SAYRE-MCCORD, G. "Different Kinds of Kind Terms: A Reply to Sosa and Kim". *Philosophical Issues,* 8, pp. 313-323, 1997.

_____ *Essays on Moral Realism*. Ithaca: Cornell University Press, 1988

SLIWA, P. "In Defense of Moral Testimony". *Philosophical Studies*, 158, pp. 175-195, 2012.

_____ "Moral Understanding as Knowing Right from Wrong". *Ethics*, 127, pp. 521-552, 2017.

STURGEON, N. "Moore on Ethical Naturalism". *Ethics*, 113, pp. 528-556, 2003.