

EXPERIMENTAL PHILOSOPHY AND THE CONCEPT OF MORAL RESPONSIBILITY

Alicia Finch

Abstract: In recent years, so-called experimental philosophers have argued that participants in the moral responsibility debate ought to adopt a new methodology. In particular, they argue, the results of experimental surveys ought to be introduced into the debate. According to the experimental philosophers, these surveys are philosophically relevant because they provide information about the moral responsibility judgments that ordinary people make. Moreover, they argue, an account of moral responsibility is satisfactory only if it is tightly connected to ordinary judgments. The purpose of this paper is to undermine this argument. I will argue that experimental philosophers have not adequately acknowledged the distinction between metaphysics and conceptual analysis; they have not carefully distinguished what-it-is-to-be morally responsible from the concept of moral responsibility. I will draw this distinction, and then argue that metaphysicians qua metaphysicians may both ignore experimental data and offer an account of moral responsibility that satisfies the tight connection desideratum.

1 INTRODUCTION

In recent years, a new philosophical movement, *experimental philosophy*, has rapidly gained momentum. Experimental philosophers “proceed by conducting experimental investigations of the psychological processes underlying people’s intuitions about central philosophical issues” (Knobe and Nichols 2007, 1),¹ and then present their data as relevant to the conduct of philosophical debates. To be clear, they do not imagine that experimental data alone can solve philosophical problems, and they insist that they have no interest in eliminating philosophical argument in favor of data collection. As the experimental philosophers describe themselves, their aim is to add to, and not subtract from, the philosopher’s methodological arsenal. They contend that if philosophers were to consider the results of empirical studies, the level of philosophical discourse could be raised; but they do not suppose that philosophers ought to aban-

¹ In what follows, it will become clear that self-described experimental philosophers also appeal to studies that have been conducted by social scientists.

don any of the argumentative strategies that they currently employ.

When one considers their research program under this description, it seems innovative at best, harmless at worse. After all, one might reason, more information is, *ceteris paribus*, better than less, and experimental philosophers are simply offering their colleagues information that they would not otherwise have. Given that this is so, one might find opposition to experimental philosophy altogether irrational.

And, yet, my purpose here is to oppose experimental philosophy. To be more precise, my purpose is to consider what experimental philosophers have to say about one philosophical debate in particular: the moral responsibility debate.² To be even more precise, my purpose is to challenge just one argument that experimental philosophers make for the claim that participants in the debate ought to consider their data: the *tight connection argument*. The pith of the argument is this: (i) if an account of moral responsibility is satisfactory, it is tightly connected to the moral responsibility judgments that ordinary people make; (ii) in order to give an account of moral responsibility that is (non-accidentally) tightly connected to such judgments, one must consider experimental data; (iii) therefore, in order to give a satisfactory account of moral responsibility, one must consider experimental data about ordinary moral responsibility judgments. Experimental philosophers thus set up a dichotomy: either pay attention to experimental data or offer an account of moral responsibility that is at best accidentally tightly connected to ordinary judgments. With this, they argue that an account of moral responsibility must, in fact, be tightly connected to ordinary moral responsibility judgments. And then they conclude straight-away that philosophers of moral responsibility qua philosophers of moral responsibility ought to pay attention to experimental data.

My purpose here is to argue that experimental philosophers have set up a false dichotomy. A philosopher of moral responsibility may both (i) offer an account of moral responsibility that is non-accidentally tightly connected to ordinary moral responsibility judgments³ and (ii) ignore experimental data.⁴ In presenting the experimental philosophers' argument, I will emphasize that it depends on the thesis that the moral responsibility debate is about nothing other than the *application of the concept* of moral responsibility. In what follows, I will challenge this thesis; in particular, I will argue that the moral responsibility debate includes *metaphysical questions* about moral responsibility, and that these questions are distinct from questions about the application of a concept.

² In what follows, I will consider which questions fall within the debate's parameters. Hereafter, I will refer to participants in the debate as 'philosophers of moral responsibility.'

³ Hereafter, for the sake of rhetorical simplicity, I will leave out the adverb 'non-accidentally' when speaking of the relevant tight connections.

⁴ Perhaps these points go without saying, but, just in case they do not: (i) I am not arguing against experimental philosophy in general. I am only concerned with the relevance of experimental data to one debate in particular; (ii) I do not deny that the experimental philosophers' data are interesting; (iii) I do not deny that it is valuable for social scientists qua social scientists to gather empirical data; (iv) nothing I say here is at odds with or depends upon the success of other challenges to experimental philosophy.

Having drawn this distinction, I will argue that a metaphysician qua metaphysician is able to provide an account of moral responsibility that is tightly connected to ordinary judgments, and she is able to do so without considering any experimental data at all. I will conclude, then, that the experimental philosophers' tight connection argument does not succeed.

2 SIGNIFICANCE

But even if this is true, one might wonder how significant this thesis is. After all, it does not imply that philosophers of moral responsibility are any the worse for considering the data. And it is surely not worthwhile to provide an excuse for (lazy) mainstream philosophers to continue in a state of ignorance, nor is it worthwhile to provide a justification for irrational conservatism about philosophical methodology. Moreover, even if experimental philosophers must abandon the tight connection argument, they are able to offer other arguments on behalf of their position. It seems worthwhile, then, to consider the dialectical significance of the project at hand.

Toward this end, it is of the first importance to acknowledge how harshly experimental philosophers criticize mainstream philosophers of moral responsibility. Let us be clear: when experimental philosophers argue that philosophers of moral responsibility ought to consider experimental data, they are implicitly arguing that philosophers of moral responsibility are doing something *wrong* if they fail to do so. Granted, they do not suggest that all philosophers ought to go out and conduct experiments, nor do they suggest that every philosophical debate would improve if participants were to attend to experimental data. Indeed, they go out of their way to insist that their project is not at odds with mainstream philosophy. And, yet, the prominent experimental philosophers Joshua Knobe and Shaun Nichols have this to say in "An Experimental Philosophy Manifesto":

If problems like free will, moral responsibility, and personal identity flow from commonsense, then to understand these problems, it would be myopic to look only at the responses of philosophers.^[5] Rather, to understand the intuitions that are at the core of philosophical problems, one would surely want to look at different groups to see whether interesting patterns of similarity and difference emerge. (2007, 7)

With this remark, the authors make it difficult to deny that they aim at undermining the mainstream approach to the moral responsibility debate. No one wants to be myopic, after all. But, if Knobe and Nichols are correct, myopia and experimental philosophy are the only dialectical options for a philosopher of moral responsibility. Granted, there are presentations of the argument that do not rely on expressions as loaded as 'myopic.' However, even the kinder, gentler variations

⁵ Knobe and Nichols are here referring to responses to questions about hypothetical scenarios in which an agent is alleged to be morally responsible.

of the argument present the same line of thought that Knobe and Nichols do.

With the tight connection argument, experimental philosophers attempt to divide philosophers of moral responsibility into two camps: those who pay attention to experimental data, and those who disregard ordinary judgments about who is morally responsible for what. If the debate were construed in this way, philosophers who care about ordinary judgments would come to see non-experiment work as disconnected from (and perhaps even at odds with) their own. And, given that there are many philosophers of moral responsibility who do, in fact, care about ordinary judgments,⁶ non-experimental work would likely be marginalized. This would be a terrible shame if, as I will argue, the dichotomy is false. Perfectly good philosophy might be ignored for no good reason, and philosophers might go to great lengths to familiarize themselves with data that, as a matter of fact, they need not consider.⁷

Perhaps it seems alarmist to raise such suggestions. Nonetheless, it is not unimaginable that experimental philosophy will eventually become the dominant methodology within the moral responsibility debate. After all, it is not uncommon for a trend to take on a life of its own, and experimental philosophers do have *arguments* for their position: it is not as if they are trying to exploit anyone's irrational impulses. And so it seems that the arguments of the experimental philosophers deserve attention, and, again, the task of this paper is consider one argument in particular. It is to this argument that I now turn.

3 THE CONTOURS OF THE ARGUMENT

In order to appreciate the full force of the tight connection argument, one ought to consider the experimental philosophers' project in more general terms. Toward this end, one ought to consider how the experimental philosophers themselves characterize their project. Of particular importance is this: although they construe philosophy as the study of the *application of concepts*, they are suspicious of the methodology of conceptual analysis.

In their manifesto, Knobe and Nichols (2007) offer these introductory remarks:

Experimental philosophers are certainly not the first to think that important philosophical lessons can be learned by looking carefully at ordinary people's intuitions about cases. This methodological approach has a long history within the research program sometimes known as 'conceptual analysis.' It may be helpful, then, to begin by discussing the ways in which experimental philosophy departs from this earlier program. (4)

And, as Knobe and Nichols suggest, the goal of a conceptual analysis is (more or less) to provide a set of conditions that are individually necessary and jointly

⁶ I am such a philosopher.

⁷ Unless, of course, one of the experimental philosophers' other arguments succeeds. But that is not the issue at hand.

sufficient for the concept's application. Indeed:

Typically, a conceptual analysis attempts to identify precisely the meaning of a concept by breaking the concept into its essential components, components which themselves typically involve further concepts. (4)

A skeletal depiction of the methodology of conceptual analysis is something like this: (i) Someone proposes a set of necessary and sufficient conditions for the application of a concept *C*; (ii) someone challenges this analysis by describing a hypothetical scenario, i.e., a thought experiment, that is supposed to elicit some intuition about whether the concept *C* applies; (iii) it is made clear that the intuition would be false if the analysis in question were adequate; (iv) the analysis is deemed “counterintuitive,” and those who will not give it up are described as “biting the bullet;” (v) those who are uncomfortable with the bullet biting begin the process again.

It is particularly relevant, in this context, that, “Conceptual analysis is supposed to be a distinctively *a priori* activity that many take to be the essence of philosophy” (Margolis and Laurence 2011). This method is *a priori* in that philosophers themselves evaluate, from the armchair, which responses to hypothetical scenarios count as “intuitive.” At the heart of experimental philosophy is a challenge to this *a priori* approach. Instead of relying on nothing other than the intuitions of philosophers, they insist, one should also consider the intuitions of ordinary people.⁸ Or, at least, one should consider such intuitions about concepts that ordinary people apply.⁹ And, of course, the concept of moral responsibility is a paradigmatic example of a concept of this kind.

It is at this point in the dialectic that the tight connection argument becomes salient:

For many standard philosophical problems—for example, problems concerning free will, personal identity, knowledge, and morality—if it weren't for commonsense intuitions, there wouldn't be a felt philosophical problem. The problem of moral responsibility, for instance . . . arises because people think of themselves as morally responsible, and this seems at odds with other important and plausible world views. (8–9)

Of course, as I have already said, Knobe and Nichols construe a non-experimental approach to the moral responsibility debate as “myopic.”

In another version of the tight connection argument, Eddy Nahmias offers that “[one] reason philosophers should be interested in ordinary intuitions is that the target concept of interest to most philosophers is the type(s) of freedom or control relevant to holding oneself and others morally responsible for their

⁸ Strictly speaking, this is only one way that experimental philosophers might gather “hard data.” They also employ the results of fMRI scans. However, in the context at hand, this method of data collection is not relevant.

⁹ As Knobe and Nichols acknowledge, ordinary intuitions about cognitive architecture seem to lack philosophical significance.

actions” (2011, 557). Moreover, in a recent blog post, Nahmias articulates nearly the same point:

Some of the reasons for caring about what the folk think [include¹⁰ that] many philosophers believe (and I agree) that the best theory of free will is one that connects it tightly to our beliefs and practices regarding moral responsibility (MR) If so, then a good theory of free will needs to know what our beliefs and practices are regarding attributions of MR and the sort of control and choice capacities we think are required for such attributions. Lots of experimental work (not just x-phi) seems relevant to approach these questions. (June 20, 2011)

With remarks such as these, experimental philosophers affirm their commitment to the claim that an account of moral responsibility must be tightly connected to ordinary judgments.

And in order to gather information about these judgments, they contend, one ought to conduct *experimental surveys*. With respect to the moral responsibility debate, experimental subjects should be presented with descriptions of various hypothetical scenarios, and asked a series of questions about whether agents in those scenarios are morally responsible for their actions.¹¹ Answers to these questions would constitute experimental data about ordinary intuitions. While conceptual analysts are content to consider nothing but philosophers’ intuitions about hypothetical scenarios, experimental philosophers are in a position to consider ordinary intuitions as well. By gathering these data, philosophers gain a richer understanding of the application of the concept of moral responsibility.

At this point, one might suppose that the only important difference between experimental philosophers and conceptual analysts is the means by which they collect their data: lab versus armchair. But, as Knobe and Nichols make clear, this is an oversimplification of the experimental philosophers’ project: “Over time, experimental philosophers have developed a way of thinking about these issues that departs in truly substantial respects from the approaches familiar from conceptual analysis” (4). Instead of offering analyses:

[The experimental philosophers’] aim is usually to provide an account of the factors that influence applications of a concept, and in particular, the *internal psychological processes* that underlie such applications. Progress here is measured not in terms of the precision with which one can characterize the actual patterns of people’s intuitions but in terms of the degree to which one can achieve explanatory depth. Typically, one starts out with a fairly superficial characterization of certain patterns in people’s intuitions. (4)

¹⁰ I want to be clear that the quotation that follows is elliptical, and that Nahmias offers many reasons for thinking that the experimental data are relevant to the moral responsibility debate.

¹¹ For the sake of simplicity, the focus of this paper will be responsibility for actions, as opposed to the consequences of actions, omissions, or the consequences of omissions. In the context at hand, this assumption is harmless.

With this, they offer an example of the kind of superficial characterization they have in mind: “People are more inclined to regard an agent as morally responsible when the case is described in vivid and concrete detail than they are when the case is described more abstractly” (5). But, they suggest, the goal of the experimental philosopher is to provide a “deeper explanation.” For instance: “People are more inclined to regard an agent as morally responsible when they have a strong affective reaction to his or her transgression” (5). Of course, if they could find an even deeper explanation, that would be all the better. And so they might conduct experiments to test a hypothesis like this: “People’s intuitions about moral responsibility are shaped by the interaction of two different systems—one that employs an abstract theory, another that relies more on immediate affective reactions” (5). Because their aim is to understand how the concept of moral responsibility is applied, and not just what the concept of moral responsibility is, their goals are broader than those of conceptual analysts. Unless one acknowledges this point, one cannot provide a fair assessment of experimental philosophy in general or the tight connection argument in particular.

4 THE DATA AND ITS SIGNIFICANCE

The next step in understanding the experimental philosophers’ project is a consideration of the data they have gathered, and the conclusions that have been drawn from them. Although detailed discussions of their findings are available, it is enough, for our purposes, to consider just a few of their findings.¹² Knobe and Doris (2010) report that:

- When subjects are placed in a “concrete condition” and told that determinism is true, they are more likely to ascribe moral responsibility than they are when they are placed in an “abstract condition” and told that determinism is true. In one study, subjects in an “abstract condition” were simply told that determinism is true, and then they were asked whether agents can be morally responsible for their actions. Subjects in the “concrete condition,” by contrast, were first told that determinism is true, and then told the story of a particular man, Bill, who burned his family alive so that he could pursue a sexual relationship with his secretary. Subjects were then asked whether Bill was morally responsible for his actions. While relatively few subjects in the abstract condition thought that moral responsibility is possible if determinism is true, an overwhelming majority of subjects blamed Bill. (See Knobe and Doris 2011, 329–30.)
- Subjects are more likely to ascribe blame for foreseen but unintended negative side effects than they are to ascribe praise for fore-

¹² See Knobe and Doris 2010, and the works that they cite, for a thorough discussion of the experimental data.

seen but unintended positive side effects. In one of the experiments, some of the subjects were placed in a “harm condition:” they were presented with a hypothetical scenario in which an executive is told that an action she intends to perform will harm the environment; she responds that she does not care, as her only concern is profit maximization. Subjects in the “help condition” were presented with a structurally similar scenario: an executive is told that the action she intends to perform will help the environment, and she responds that she does not care, as her only concern is profit maximization. Subjects in the “harm condition” were more likely to ascribe blame than subjects in the (structurally similar) “help condition” were to ascribe praise. (331–2)

- “[Subjects] gave agents considerably less blame for morally bad behaviors when those behaviors were the result of overwhelming emotion than when they were the result of calm deliberation. But for morally good behaviors, there was no corresponding effect” (334).
- “Subjects were considerably more willing to assign blame for morally bad intentions that were not acted upon than they were to assign praise for morally good intentions” (335).
- “Four decades of research on this topic points unequivocally to the conclusion that . . . people are willing to say that an agent is responsible for a severe harm (e.g., the death of a child) even when that agent’s behavior was only very slightly negligent whereas they refused to say that an agent is responsible for mild harms (e.g. a fender bender) unless the agent was very negligent indeed” (337).
- Data indicate that there is a significant correlation between marital satisfaction and blame ascription: the more satisfied subjects are in the marriages, the less likely they are to ascribe blame to their spouses when they are presented with hypothetical scenarios in which their spouses engage in bad behavior. Moreover, other studies “seem to indicate that indicate that people’s judgments about particular cases can vary dramatically depending on their relationship to the agent [being judged]” (345).

With these data in tow, experimental philosophers go on to formulate hypotheses about how ordinary people apply the concept of moral responsibility.

As one would expect, experimental philosophers focus on the inconsistency of ordinary applications of the concept. And it is at this point that they are in a position to set up a dichotomy between experimental philosophy, on the one hand, and disconnection from ordinary judgments, on the other. Granted, they may not explicitly state that such a dichotomy exists; however, it is clear that they rely on an argument something like this: (i) if an analysis of the concept of

moral responsibility were possible, it would be possible to give necessary and sufficient conditions for the application of the concept; (ii) if it were possible to give necessary and sufficient conditions for the application of the concept, then people would, by and large, apply the concept consistently; (iii) the data indicate that ordinary people do not, by and large, apply the concept consistently; (iv) therefore, it is not possible to give an analysis of the concept of moral responsibility that is tightly connected to ordinary judgments.

Moreover, they might add, it is not as if inconsistencies in the application of the concept can be explained away by the thesis that, when it comes to applying the concept, there are a few outliers who are making mistakes. The data indicate that inconsistent judgments are the norm. And so it is that philosophers of moral responsibility must make a choice: (i) insist on pursuing the project of conceptual analysis or (ii) join experimental philosophers in taking ordinary judgments into account. There are, they suggest, no other options.

5 DRAWING A DISTINCTION

In construing the debate in this way, experimental philosophers assume that, in the context of the moral responsibility debate, there is no important distinction between the *concept* of moral responsibility and what it is to *be* morally responsible; that is, they make the assumption that there is no important distinction between metaphysics and the study of concepts.¹³ In what follows, I will argue that this is a mistake.

Of course, as anyone familiar with the history of philosophy knows, it is far from obvious how to draw the distinction in question. In order to address this issue adequately, one ought to go back at least as far as Kant, and one ought not to stop before reaching the middle of the twentieth century. According to the standard narrative, it was at this point in the history of philosophy that metaphysics became (relatively) respectable again, after having been derided by analytic philosophers for decades. Of course it is impossible to provide such a discussion here.

I will begin, then, by simply gesturing at Kant's distinction between *transcendental* and *critical metaphysics*:

Whereas transcendent metaphysics seeks to characterize a reality that transcends sense experience, critical metaphysics has as its task the delineation of the most general features of our thought and knowledge. It seeks to identify the most general *concepts* at work in our representation of the world, the relationships that obtain among those *concepts*, and the presuppositions of their objective employment (*italics added*). (Loux 2006, 7)¹⁴

¹³ Or, at least, they had better be making this assumption: if not, they are cutting metaphysicians out of their discussion for no good reason. And it would be most uncharitable to suppose that they are doing so, given how many participants in the debate describe themselves as metaphysicians.

¹⁴ Neither Loux nor I intend this as a piece of Kant scholarship. Loux is painting in broad strokes, and so am I.

Kant's distinction corresponds, roughly, to the distinction between metaphysics *simpliciter* and the study of concepts. Metaphysicians study "reality in itself" while students of concepts study the constituents of thoughts.

But what exactly is "reality in itself"? Once again, it would be best to conduct a careful examination of the history of philosophy, and, once again, this task is beyond the scope of this paper. This time, I will begin by gesturing at Aristotle. In the fourth book of his *Metaphysics*, he calls it the study of *thing that is qua thing that is*, and, in other passages, he describes it as the study of first principles. But by the time that the rationalists come along, metaphysics is construed (roughly) as the study of the nature and existence of, for instance, God, the soul, matter, causation, and freedom of the will. None of these suggestions is entirely perspicuous, so it is perhaps best just to use these rather cryptic expressions to describe the subject matter of metaphysics: what *is*, what is *real*, what *exists*, what *is the case*, what *there is*; and it is perhaps worthwhile to note the use of italics, and to emphasize that if one were to utter any of these expressions aloud, one ought to use a metaphysical tone of voice.

At this point, it might be useful to articulate the intended distinction in extremely crude terms: metaphysics is the study of what is "in the world," while a concept is "in the mind." To study concepts is to study how the mind carves up the world; to study metaphysics is (supposed to be) to study how the world itself is carved up. Metaphysics is (supposed to be) the study of the nature and structure of reality itself.

In any case, this distinction is relevant to the present discussion because a participant in the moral responsibility debate might aim to study what-it-is-to-be morally responsible, but the experimental philosophers seem to focus on nothing but the application of the concept of the same. This is evident in the way that they set up the dichotomy between conceptual analysis, on the one hand, and the empirical study of the use of concepts, on the other. They argue against conceptual analysis, and conclude that experimental philosophy is the better way. But if, as I will argue, the metaphysician can offer a third way, this inference is far too quick.

6 THE METAPHYSICS OF MORAL RESPONSIBILITY

In order to explain how metaphysicians of moral responsibility are relevantly distinct from those who study the concept, it is best to begin with a discussion of what they may agree about. First, they must agree about the claim that:

- (MR) An agent *S* is *morally responsible* for performing an act *A* =_{df.}
 (i) *S* is (*morally*) *blameworthy* for performing *A* or (ii) *S* is (*morally*) *praiseworthy* for performing *A*.

Moreover, they might both accept, for instance, Derk Pereboom's characterization of moral responsibility:

- (MR*) For an agent to be morally responsibility for an action is for it to belong to her in such a way that she would deserve blame if she

understood that it was morally wrong, and she would deserve credit or perhaps praise if she understood that it was morally exemplary. (Pereboom 2007, 86)

After all, it seems that one might reasonably construe (MR*) as *either* the concept *or* the what-it-is-to-be of moral responsibility.

The issues become a bit more complicated, though, when one considers a statement like this:

(MR**) An agent *S* is morally responsible for performing an act *A* only if (i) *A* is either morally good or morally bad; (ii) *S* performs *A* intentionally; (iii) it was up to *S* whether *S* performs *A*; and (iv) *S* had or has a reasonable belief about whether *A* is morally good or morally bad.¹⁵

The complication is that it is not clear whether this statement describes necessary conditions for *being* morally responsible or necessary conditions for the application of the *concept* of moral responsibility. But, for the sake of the present discussion, it seems best not to dwell on this. No one in the moral responsibility debate seems concerned about this question, and so it seems to be assumed that (MR**) expresses necessary conditions for both an agent's being morally responsible and the application of the concept of moral responsibility.

From here on out, though, the distinction between concept and what-it-is-to-be cannot be ignored. In order to see that this is so, it is perhaps best to begin by considering the "up to" condition. As anyone familiar with the moral responsibility debate knows (and, indeed, as the preceding quotations from the experimental philosophers suggest), questions about this condition have long dominated the moral responsibility debate. But the student of concepts and the metaphysician will approach this question entirely differently.

The metaphysician qua metaphysician is concerned about questions like these: Is it possible that determinism is true and it *is* up to an agent which action she performs? Is it possible that indeterminism is true and it *is* up to an agent which action she performs? Moreover, a metaphysician might ask: *is* it up to an agent which action she performs if she lacks "alternative possibilities"? Must an agent *be* the source of her action if it *is* up to her which action she performs? If it is up to an agent whether she performs an action only if she *is* able to do otherwise, what is the modal force of *able*? These are questions about the what-it-is-to-be of its being up to an agent whether she performs. How must the world *be* if it is up to an agent which action she performs? How does the agent

¹⁵ This statement is meant only to gesture at necessary conditions for an agent's being morally responsible for an act; it is not meant to be precise. There are, after all, a number of questions about how an agent's responsibility for an act *now* might be *derived from* her earlier acts, decisions, omissions, failures of attention, or the like. Philosophers of moral responsibility often talk about "tracing" an agent's responsibility for an act (back) to earlier events. Moreover, this statement is not meant to imply that one can perform an act freely without performing it intentionally. In the context at hand, there is no harm in being pleonastic. And I cannot emphasize enough that this statement is neutral with respect to the truth of either a "source theory" or a "leeway theory" of "up to."

stand, out in the world (so to speak), with respect to an action?

Let us suppose that metaphysicians keep asking these questions, and that they do not stop until they have a full account of the necessary and sufficient conditions for its *being* up to an agent which action she performs. And let us suppose that, not only do they investigate the “up to” condition, they keep going and attempt to give metaphysical accounts of the other necessary conditions for *being* morally responsible. That is, let us suppose that they attempt to give a full metaphysical account that provides both the necessary and sufficient conditions for being morally responsible.

And now let us ask: are they attempting to provide necessary and sufficient conditions for the application of the concept? This much seems true: at least some metaphysicians are attempting to provide necessary and sufficient conditions for the *correct* application of the concept, and they assume that someone correctly applies a concept if and only if it is consistent with how the world *really is*. To judge truly that someone is responsible is to apply the concept of moral responsibility correctly.

Of course, given that these metaphysicians think that the world really is a certain way, and that people make false judgments about how the world is, they are committed to the view that people apply the concept incorrectly. Or, they might say, people apply the concept inconsistently. Some people apply it one way; other people apply it in an opposing way; and, hence, their applications of the concept are inconsistent. But if they proceed on the assumption that people apply the concept of moral responsibility inconsistently, these metaphysicians cannot reasonably suppose that there is a correct analysis of the concept. If the concept is applied inconsistently, there is no set of necessary and sufficient conditions that captures the way that the concept is applied.

With this, let us return to the experimental philosophers’ picture of the moral responsibility debate: there are no necessary and sufficient conditions for the application of the concept of moral responsibility, given that ordinary people make inconsistent judgments. In some conditions, ordinary people judge one way; in other conditions, they judge other ways. This, they contend, is at odds with the claim that one can provide a conceptual analysis, that is, a set of necessary and sufficient conditions for the application of the concept of moral responsibility.

And so it is that there are metaphysicians who wholeheartedly agree with experimental philosophers that there is no analysis of the concept of moral responsibility, and who agree with the experimental philosophers that the fact that people make inconsistent judgments is evidence for this thesis. Indeed, they might say: we have been assuming all along that people apply the concept inconsistently and, hence, that a conceptual analysis of moral responsibility is impossible. We did not need experimental data to know that people apply the concept of moral responsibility inconsistently. There is, after all, plenty of anecdotal evidence that ordinary people make conflicting judgments all the time. But this just means that people are very often wrong about who *is* morally responsible for what. There are conditions that an agent must meet for *being*

responsible, and plenty of people make moral responsibility judgments without knowing whether the agent meets those conditions. This is to be expected, given that it takes a great deal of metaphysical work to figure out what these conditions are. But there is a “fact” about how the world is, and it is in virtue of this “fact” that a judgment about moral responsibility is true or false. The world is either one way or another, and so it is impossible that conflicting judgments are true. Nothing we say is undermined by data that shows that people make conflicting judgments about moral responsibility. Again, we have been assuming all along that people judge inconsistently.

7 METAPHYSICS AND ORDINARY JUDGMENTS

But, one might wonder, what does this have to do with the tight connection argument? After all, it is not clear that an account of moral responsibility can satisfy the tight-connection desideratum if it is based on the assumption that ordinary people are very often wrong in their moral responsibility judgments. And the metaphysician I have been considering is committed to precisely this assumption.

It is at this point, I contend, that one must consider the different ways in which an account of moral responsibility might be connected to so-called ordinary judgments: the account might be connected to (i) the content of the judgments that so-called ordinary people make when they are in a lab or (ii) the way that people (both ordinary and not) go about making judgments over the course of their ordinary, everyday lives.

With respect to the latter, it is crucial to acknowledge how common it is for people to have exchanges like these:

- 1a. I feel so guilty! It’s all my fault that it happened.
- 1b. It’s not your fault. You didn’t really have a choice.
- 2a. You can’t hold it against her. She didn’t even know that she was doing anything wrong.
- 2b. She did too know it was wrong and she did it anyway. She has no excuse.
- 3a. You can’t blame me! I didn’t do it on purpose!
- 3b. What do you mean you didn’t do it on purpose? It’s not like you just *accidentally* looked online and copied something from Wikipedia into your paper.
- 4a. You have made me proud!
- 4b. You’re just saying that because you’re my mother. I don’t deserve praise for doing what anyone would have done in the circumstances. Really, what choice did I have?
- 5a. I’m really proud of myself! I wanted to take the easy way out, but I knew it was wrong, so I didn’t.
- 5b. I don’t know why she’s so proud of herself. Really, it was the least she could do.

- 6a. I really don't deserve all this praise. I wasn't trying to do anything special. I was just doing my job.
- 6b. You went above and beyond, and you should give yourself credit for that.

Indeed, it seems that most of the moral responsibility judgments that people issue are issued in the context of such exchanges. Moreover, the most obvious way to understand these exchanges is this: people are making conflicting judgments about who is morally responsible, and each person thinks that what she is saying is true, and that what the other person is saying is false. In other words, it is reasonable to think that when people have conversations like this, they implicitly assume that the parties with whom they disagree are *wrong*. And if this is what ordinary people implicitly assume, then there is an important sense in which the metaphysicians I have been considering are able to give an account of moral responsibility that satisfies the tight connection desideratum. After all, both the metaphysicians and ordinary people agree that people are very often wrong in their judgments about who is morally responsible for what. Moreover, it seems as if ordinary people agree with the metaphysicians in thinking that a judgment is right or wrong based on how the world *really is*.

But, one might ask, do ordinary people actually assume that their opponents' judgments are false, and their own judgments are true? And do they assume that judgments are right or wrong based on "how the world really is?" These are reasonable questions, given how many people profess to be moral relativists, and how many people express distaste at the idea of "judging" others to be at fault. In fact, might not the experimental philosopher conduct studies that suggest that ordinary people do not take themselves to be making true/false judgments in the cases I consider? And if the experimental philosopher were to do so, and if the data suggested that ordinary people do not take themselves to be doing this, would not the metaphysician be vulnerable to the objection after all?

The first response is that they would actually have to go out and do this in order to make the tight connection argument work; as it is, the non-experimental metaphysician has a response.

But, more to the point, it seems worthwhile to acknowledge this: when people make such judgments, they certainly *seem* to be assuming that there is a fact of the matter about whether someone is morally responsible. No matter what they would say if they were asked about the truth value of their claims, to observe them in action is to get the impression that they proceed on the assumption that they are right and their opponents are wrong. They do, after all, offer *reasons* on behalf of their positions.

Of course, one might wonder why we should privilege the way that ordinary people behave in their ordinary lives over the responses that they give in the lab. But this brings us right back to where the experimental philosopher began: because our goal is to come up with an account of moral responsibility that is tightly connected to the *ordinary* judgments that people make in their day-to-day lives.

The idea, then, is this: if one looks at the way that people issue judgments in their ordinary lives, one sees that they proceed on the assumption that the people who disagree with them make false judgments. And, of course, one can gather this from anecdotal evidence, and so one does not need to go into the lab, or look at experimental data, to know that this is so. But if an account of moral responsibility is “tightly connected” to this aspect of ordinary moral responsibility judgments, it satisfies the tight connection desideratum. Metaphysicians qua metaphysicians can offer such an account. Given that this is so, the experimental philosophers’ tight connection argument does not succeed.

8 CONCLUSION

With this, it ought to be clear that the experimental philosophers are far too hasty in making the tight connection argument. Because they do not consider the distinction between metaphysics and the study of concepts, they do not recognize all the dialectical options available to a participant in the moral responsibility debate. Moreover, they do not acknowledge that there is more than one way that an account of moral responsibility might satisfy the tight connection desideratum. They seem to assume that satisfying this desideratum requires coming up with an account of moral responsibility that is connected to the content of the judgments that ordinary people make. But, as I have argued, there is an entirely different approach: one may offer an account that is consistent with the ordinary view that some moral responsibility judgments are right, and some are wrong. Metaphysicians, of course, are in a position to offer an account of moral responsibility that makes contact with *this* aspect of ordinary judgments. Moreover, it is not at all clear that experimental philosophers are in a position to do so: they must provide some reason for thinking that they are. In the light of these considerations, then, it seems as if the tight connection argument does not succeed. Once one considers the debate from a different perspective (that of the metaphysician), one can see that the dialectical situation is not nearly so simple as the experimental philosopher takes it to be.¹⁶

Alicia Finch
 Department of Philosophy
 Northern Illinois University
 DeKalb, IL 60115 USA
 E-mail: afinch@niu.edu

References

- Knobe, Joshua 2003. “Intentional Action and Side Effects in Ordinary Language.” *Analysis* 63 (3): 190–193.
 Knobe, Joshua & Shaun Nichols. 2007. “An Experimental Philosophy Manifesto.” In *Experimental*

¹⁶ I am grateful to Susan Brower-Toland, Mary Beth Finch, Kevin Timpe, and Margaret Watkins for very helpful comments on earlier drafts of this paper.

- Philosophy*, edited by Joshua Knobe & Shaun Nichols, 1–23. New York: Oxford UP.
- Knobe, Joshua & John M. Doris. 2010. “Responsibility.” In *The Moral Psychology Handbook*, edited by John M. Morris, 321–54. New York: Oxford University Press.
- Loux, Michael. 2006. “Introduction.” *Metaphysics: A Contemporary Introduction*. NY: Routledge, 1–16.
- Margolis, Eric and Stephen Laurence. 2011. “Concepts,” in *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*, edited by Edward N. Zalta. <http://plato.stanford.edu/archives/fall2011/entries/concepts/>.
- Nahmias, Eddy. 2011. “Intuitions about Free Will, Determinism, and Bypassing.” In *The Oxford Handbook of Free Will* (2nd ed.), edited by Robert Kane, 555–76. New York: Oxford UP.
- _____. June 20, 2011. “Why Folk Intuitions Matter to the Free Will Debate and Why We Should Study Them Empirically.” *Flickers of Freedom* Blog. <http://agencyandresponsibility.typepad.com/flickers-of-freedom>.
- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, & Jason Turner. 2006. “Is Incompatibilism Intuitive?” *Philosophy and Phenomenological Research* 73 (1): 28–53.
- _____. 2005. “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility.” *Philosophical Psychology*. 18 (5): 561–84.
- Nichols, Shaun & Joshua Knobe. 2007. “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions.” *Noûs* 41 (4): 663–85.
- Pereboom, Derek. 2007. “Hard Incompatibilism.” In *Four Views on Free Will*, edited by John Martin Fischer, Robert Kane, Derek Pereboom, and Manuel Vargas, 85–124. Malden, MA: Blackwell.