

Folk Intuitions on Free Will*

SHAUN NICHOLS**

ABSTRACT

This paper relies on experimental methods to explore the psychological underpinnings of folk intuitions about free will and responsibility. In different conditions, people give conflicting responses about agency and responsibility. In some contexts, people treat agency as indeterminist; in other contexts, they treat agency as determinist. Furthermore, in some contexts people treat responsibility as incompatible with determinism, and in other contexts people treat responsibility as compatible with determinism. The paper considers possible accounts of the psychological mechanisms that underlie these conflicting responses.

KEYWORDS

Free will, responsibility, folk intuitions, experimental philosophy, conceptual analysis

It is a messy business to chart the psychological underpinnings of the folk intuitions that generate philosophical problems. One tantalizing speculation is that many persistent philosophical problems have their source in conflicting folk intuitions that are generated by different psychological mechanisms. In this paper, I will argue that the cluster of problems surrounding free will and responsibility implicate several different psychological factors. I want to sketch the apparently complex psychological underpinnings of the intuitions implicated in the free will problem. In doing this, I will review recent evidence on folk intuitions, and I'll also present a bit of new evidence. The first order of business, though, is to clarify the terrain of the free will problem and to explain the methodology I'll use for getting at the folk intuitions.

* Thanks to Joshua Knobe, Heidi Maibom, Ron Mallon, Philip Robbins, Tamler Sommers, Paulo Sousa, Manuel Vargas, and Jonathan Weinberg for helpful comments on an earlier draft of this paper.

** Department of Philosophy, University of Arizona. Email: snichols@email.arizona.edu

1. Three projects

Traditionally, the problem of free will has two axes. Both axes pivot on the notion of determinism, according to which every event is an inevitable outcome of the past conditions and the laws of nature. One axis of the free will problem concerns the character of human agency. Some maintain that human decisions are determined by the past. Others maintain that human decisions are at least partly outside the sphere of determinism; for instance, some philosophers maintain that agents enjoy a special kind of causation, *agent causation*, that is not a mere function of deterministic forces. The other axis of the problem concerns the conditions for moral responsibility.¹ Incompatibilists maintain that moral responsibility is impossible for creatures whose decisions are determined. Compatibilists deny this and insist that the truth of determinism would not undermine our moral responsibility.²

Although the issues of agency and responsibility are at the heart of the matter, the philosophical geography is rather complicated. The broad inquiry into free will and responsibility can be divided into three quite different projects: a descriptive project, a substantive project, and a prescriptive project. In this article, I propose to keep these projects entirely distinct. My focus will be mostly on the descriptive project, but it is useful to be clear about the overall nature of the inquiry. For, as we'll see in section 2, many philosophers do not keep the descriptive project distinct from substantive and prescriptive concerns.

The goal of the descriptive project is to determine the character of folk intuitions surrounding agency and responsibility. By uncovering the folk intuitions, one hopes to be able to sketch out the folk theory that underlies these intuitions.³ However, if it turns out that there is no unified

¹ The relevant notion of moral responsibility itself is contentious in the free will debates. For the purposes of this paper, I intend "moral responsibility" to pick out the notion of responsibility that is tied to moral desert, blame, and retributive punishment.

² Compatibilists about responsibility also often maintain that determinism is consistent with *free will*. However, it's possible that compatibilism is true with respect to responsibility but not with respect to free will (see Fischer 1999). Unless otherwise noted, the discussion here will concern compatibilism about responsibility.

³ "Theory" is intended in a weak sense here. Roughly, any internally represented body of information will count as a theory in the intended sense.

folk theory for the domain, then the researcher engaged in the descriptive project will mark this fact. For the aim is to give the best account of the folk intuitions and their psychological underpinnings, even if this means allowing that the folk have blatantly inconsistent intuitions.

The goal of the substantive project is then to determine whether the folk views are correct. Given the folk concepts and the way the world is, does free will exist? Are people morally responsible? Here a number of considerations come into play. Facts about psychology, neuroscience, and physics might all have bearing on whether we are free and responsible. So, for instance, ‘hard determinists’ effectively maintained (i) the folk concepts of *agency* and *moral responsibility* are indeterminist and incompatibilist and (ii) determinism is globally true. From this hard determinists draw the conclusion that the folk concepts of *free will* and *moral responsibility* are in error, i.e., the concepts fail to refer to anything. Hard determinists are thus ‘error theorists’ about free will and moral responsibility. Before this error-theoretic conclusion can be drawn, though, we would also need to settle long-standing issues about concepts and reference. In particular, one needs to know the basic criteria for what determines the reference of a concept. On some approaches (e.g., Lewis 1972, Jackson 1998), if a concept is implicated in a significant set of false beliefs, this can mean that the concept fails to refer. On other approaches (e.g., Lycan 1988, Fodor 1998), on the other hand, even if a concept is implicated in numerous false beliefs, the concept still might refer. Thus, even if the hard determinist is right that our concept of *moral responsibility* is enmeshed in false beliefs, we won’t know whether error theory follows until we know how the reference of a concept is determined.⁴

The prescriptive project is different from both the descriptive and substantive projects. For here the question is whether, given what we know about our concepts and the world, we should revise or preserve our practices that presuppose moral responsibility, like practices of blame, praise, and retributive punishment. Interestingly, simply knowing the answer to the substantive question, “are people really morally responsible?” might not tell us whether we should preserve our practices. For,

⁴ A similar argument is made, with much greater care, in Manuel Vargas’s forthcoming “The Revisionist’s Guide to Moral Responsibility”.

as we saw above, on some accounts of concept-reference, it's not all that easy for a concept to be in error. So it might turn out that our false beliefs about responsibility do not suffice to make it the case that the *responsibility* concept is empty and yet those false beliefs *do* indicate that we should change our practice. After all, rectifying false beliefs about a domain is often sufficient to warrant changing our practices even in the absence of any invocation of error theory. When the Western medical profession came to reject the old view that infants don't feel pain, this should have had (and did have) a significant impact on the treatment of infants. But the medical professionals probably never even contemplated error theory about *infants*, *pain*, or *morality*. Similarly then, if we discovered that people's beliefs about *responsibility* are mistaken in some respects, even if we avoid the error-theoretic conclusion that no one is responsible, it might be that our new beliefs about responsibility give us reasons to alter dramatically our practices towards people. Now consider the other side. As noted earlier, on some accounts of concept-reference, it's fairly easy for a concept to be in error. As a result, on these accounts, it might well turn out that, because of various mistaken beliefs, the folk concept of *the flu* fails to refer.⁵ Nonetheless, it might also be the case that our practices shouldn't significantly change. The same individuals who were said to have the flu should still be kept home from school and work, should still get plenty of rest, and should be given extra consideration around the house. Similarly, then, it might turn out that because of various false beliefs, the folk concept of *responsibility* fails to refer, and yet it's still appropriate to treat people in much the same way.

2. The Descriptive Project: A Methodological Preamble

The goal of the descriptive project is to determine the character of folk intuitions surrounding agency and responsibility. This project has important similarities to more traditional philosophical endeavors in conceptual analysis. Indeed, the descriptive project that I'll promote has significant debts to methods used in conceptual analysis. But the project also diverges

⁵ Thanks to Steve Downes for this example.

in significant ways from the traditional armchair approaches. To see how, let's begin with the traditional approaches.

2.1. *Traditional conceptual analysis*

On the dominant approach to analyzing philosophically important folk concepts, the basic method is to consult one's intuitions about possible cases, and this is supposed to provide the foundation for giving an analysis of our concepts. Frank Jackson (1998; 2001a,b) has given perhaps the most developed account of this method. He writes:

How should we identify our ordinary conception? The only possible answer, I think, is by appeal to what seems to us most obvious and central about free action, determinism, belief, or whatever, as revealed by our intuitions about possible cases. Intuitions about how various cases, including various merely possible cases, are correctly described in terms of free action, determinism, and belief, are precisely what reveal our ordinary conceptions of free action, determinism, and belief, or, as it is often put nowadays, our folk theory of them. For what guides me in describing an action as free is revealed by my intuitions about whether various possible cases are or are not cases of free action. Thus my intuitions about possible cases reveal my theory of free action . . . To the extent that our intuitions coincide, they reveal our shared theory. To the extent that our intuitions coincide with those of the folk, they reveal the folk theory (Jackson 1998, 31-2).

Thus, to get an analysis of a philosophically interesting concept, I consider my own intuitions about various actual and possible cases. This will gradually lay bare our ordinary conceptions – the folk theory. Jackson notes that there will be some question about whether my own intuitions are representative, but he suggests that this isn't a major problem with the method since "often we know that our own case is typical and so can generalize from it to others" (1998, 37).

There are a number of methodological problems that have been levied against this sort of armchair conceptual analysis. I'll list three (see also Stich & Weinberg 2001). First, practitioners of this method will be blind to the distinctive contributions of culture. By the very armchair nature of the enterprise, one will be unable to recognize which features of one's concepts are culturally local. However, there is reason to suspect that there is significant inter-cultural variation in philosophically

central intuitions (Weinberg et al. 2001; Machery et al. 2004). Second, the assumption that one's own intuitions are typical runs up against the fact that even within cultures there are apparently salient individual differences in intuitions about philosophically important concepts (see the results of Machery et al. 2004 and Knobe 2003a). Third, a philosopher's intuitions about cases probably sometimes reflect the effects of indoctrination.⁶ Thus, the armchair approach seems vulnerable to vitiating effects of training. The core concern that underlies all of these worries is that one can't tell from the armchair the extent to which one's own intuitions reveal a folk theory.

Although I fully share the above misgivings about the traditional project, for present purposes, an equally important shortcoming of prevailing approaches to conceptual analysis is that they fail to sustain the descriptive project as purely descriptive. Rather, philosophers taking an armchair approach often fuse revision into the project of conceptual analysis, providing "reforming analyses." Jackson is quite explicit in building substantive considerations into the project of conceptual analysis, saying "conceptual analysis has a prescriptive dimension" (2001a, 618). Similarly Allan Gibbard writes:

An analysis can be offered not as a bald statement of fact about what people mean, but as a proposal. Where a term is problematical, a new and clearer sense may serve its purposes – or some of them. No unique analysis need be correct; rather, we can expect some analyses to work better than others (Gibbard 1990, 32).

Indeed, Gibbard goes on to write "Any philosophical analysis strains its concept" (Gibbard 1990, 32). Fortuitously, Jackson uses the concept of free will as an example in reforming analysis. He maintains that the ordinary notion of free will is incompatibilist and perhaps "embodies some kind of confusion" (2001a, 618), but he claims that compatibilists nonetheless provide a successful analysis of the concept of free will. So, even though compatibilist accounts conflict with some folk intuitions, a

⁶ In pilot work, Gary Bartlett (unpublished data) found some suggestive results along these lines. Undergraduates enrolled in an epistemology course completed a set of epistemic thought experiments once at the beginning of the course and again at the end of the course. On key questions, student intuitions shifted to become more like the views of the professor.

compatibilist can give a proper analysis of the notion of free will (e.g., 2001b, 661). Thus this form of conceptual analysis delivers proposals that are effectively of the form, “*T* is the folk theory, suitably revised to eliminate various mistakes.” Both Gibbard and Jackson suggest that conceptual analysis can accordingly be thought of as what Quine dubs ‘paraphrasing’. When we paraphrase, the objective is not “synonymy, but just approximate fulfillment of likely purposes of the original sentences” (Quine 1960, 224; quoted in Jackson 1998, 45; see also Gibbard 1990, 32). Thus, on this approach to conceptual analysis, we provide analyses that will partly reform the folk concept.⁷

Perhaps the project of generating reforming analyses is important for some concerns, but it is a dangerous path to take if the goal is to understand folk concepts. When we try to determine the nature of philosophically important concepts used by the folk, we are basically doing a kind of cognitive anthropology. Elsewhere in cognitive anthropology, we wouldn’t brook reforming analyses. Imagine that an anthropologist returned from an isolated isle and provided the following executive summary of his research: “The indigenous people have a concept that is identical to our concept of anger, once their concept is revised to eliminate various errors.” The problem with this isn’t so much that it would patronize the indigenous people. The problem is it would take all the fun out of anthropology and sap it of much of its interest. We want to know when a culture’s emotion concepts differ from ours. A similar lesson applies to the project of characterizing philosophically important folk concepts. If substantive and prescriptive considerations are fused into the project of conceptual analysis, we will be quick to eliminate inconsistencies in folk intuitions about a domain. But if folk intuitions concerning philosophically important issues are wrong or inconsistent, that’s a richly interesting anthropological fact that we should acknowledge, contemplate, and investigate, not sweep away in a blur of revision.⁸

⁷ In section 1, I charted three projects in the broad inquiry: a descriptive project, a substantive project, and a prescriptive project. Gibbard and Jackson maintain that conceptual analysis isn’t purely descriptive, but they do not explicitly address whether both substantive and prescriptive considerations might come into play in generating a reforming analysis.

⁸ There are, of course, other disadvantages with giving a reforming analysis without first doing the purely descriptive project. For arriving at the best view about how to

2.2. *Empirical conceptual analysis*

Recently, an alternative approach to analyzing philosophically important concepts has emerged. A handful of philosophers have begun using empirical methods to do conceptual analysis. The basic idea has been to adopt techniques from the social sciences to explore folk intuitions surrounding philosophically interesting concepts (e.g., Knobe 2003a,b, Nichols 2004a, Weinberg et al. 2001). This approach avoids many of the shortcomings of armchair analysis. It also provides a manifestly superior framework for uncovering the psychological underpinnings of philosophical intuitions.

On the empirical approach, one doesn't simply assume that one's intuitions are representative of the folk; rather, interviews and anonymous surveys are used to get a more objective measure of the folk intuitions. Indeed, one can easily appropriate aspects of the traditional method of cases into this empirical approach. This also naturally lends itself to exploring whether there is cultural variation or individual differences in intuitions surrounding philosophically important concepts. If one's own intuitions are culturally local, then this can be uncovered by applying the empirical approach. Further, the possibility of training effects on intuitions doesn't vitiate the empirical project; on the contrary, it poses an interesting research question – how does philosophical training influence one's intuitions? Furthermore, like the social sciences generally, the empirical approach is quite deliberately a purely descriptive enterprise. We are trying to discern exactly what the folk think about the cases. Attempts at revision are inappropriate for this part of the inquiry.

Empirical conceptual analysis thus seems more methodologically reputable than traditional conceptual analysis. But for our purposes, there's another reason that the traditional method is inferior. For traditional conceptual analysis doesn't allow us to plumb the psychological underpinnings of the folk intuitions that underlie philosophical problems. Some of the most interesting questions about the "psychology of philosophy"

reform folk concepts presumably depends on having a clear-eyed view of what the folk concepts *are*. Furthermore, since people inevitably disagree about *how* to reform the folk concepts, those prescriptive disputes can be made much sharper if the descriptive project is done independently.

simply won't be available to traditional conceptual analysts. From the armchair we can't discern which psychological mechanisms subserve the intuitions. Nor can we determine whether conflicting intuitions have their origins in different psychological mechanisms. Empirical methods do allow us to investigate these matters. For instance, we can vary experimental conditions to determine which factors will influence people's intuitions. We can also use such methods to glean which psychological mechanisms are implicated in generating the intuitions. The empirical approach allows us to investigate matters that are entirely closed off to the armchair theorist.

Thus, I think that the empirical approach is promising indeed. Nonetheless, the best way to defend its promise is to deliver the goods. What this involves is developing detailed psychological models that generate predictions about specific kinds of cases, and then doing experiments to test those predictions. In what follows, I will present some early and awkward attempts to meet this challenge.

3. Agency

The fundamental issue about agency concerns whether or not an agent's choices are determined. A closely related issue concerns whether or not free will is consistent with determinism. But since there are many different notions of free will, the debate here often descends into squabbles about which kind of free will is under consideration. Rather than debate about the proper concept of free will, it will be more fruitful to ask more directly about folk views on the nature of agency. Do the folk treat choice as deterministic or indeterministic? The answer, I submit, is that they do both (see also Honderich 1988). When engaged in the practical process of predicting and explaining behavior, people treat choice as deterministic, but in other contexts, people seem to regard choice as indeterminist. Let's start with the indeterminist intuition.

3.1. *Choice as not determined*

In pilot studies, adults were presented with physical events (e.g. a pot of water coming to boil) and choice events (e.g. a man chooses vanilla

ice cream), and they were asked how much they agreed with a deterministic claim of the form, “if everything in the world was the same up until the event occurred, the event had to occur.” Adults registered stronger disagreement with the deterministic claim about the choice event than the physical event. Based on these pilot results, I ran a similar task on 4-5 year old children. They were presented with scenarios of physical events, spontaneous choices, and moral choices. In one of the moral choice scenarios, Mary chooses to steal a candy bar. After correctly answering some comprehension questions, children were asked, “Okay, now imagine that all of that was exactly the same and that what Mary wanted was exactly the same. If everything in the world was the same right up until she chose to steal, did Mary have to choose to steal?” In one of the physical event cases, a pot of water is put on a stove and boils. After comprehension questions, the children were asked “Okay, now imagine that all of that was exactly the same. If everything in the world was the same right up until the water boiled, did the water have to boil?” In this study, children were more likely to say that the physical events *had to happen* than that the moral choice events *had to happen* (Nichols 2004b). This provides some support for the view that people will treat choice as indeterminist under certain conditions.

Further support for the claim that people regard choice as indeterminist comes from recent experiments that Joshua Knobe and I have conducted (Nichols & Knobe forthcoming).⁹ We presented subjects with a questionnaire that depicted both a determinist universe (A) and an indeterminist universe (B). The descriptions went as follows:

Imagine a universe (Universe A) in which everything that happens is completely caused by whatever happened before it. This is true from the very beginning of the universe, so what happened in the beginning of the universe caused what happened next, and so on right up until the present. For example one day John decided to have French Fries at lunch. Like everything else, this decision was completely caused by what happened before it. So, if everything in this universe was exactly the same up until John made his decision, then it *had to happen* that John would decide to have French Fries.

⁹ Further details and discussion of this experiment will be given in section 4.

Now imagine a universe (Universe B) in which *almost* everything that happens is completely caused by whatever happened before it. The one exception is human decision making. For example, one day Mary decided to have French Fries at lunch. Since a person's decision in this universe is not completely caused by what happened before it, even if everything in the universe was exactly the same up until Mary made her decision, it *did not have to happen* that Mary would decide to have French Fries. She could have decided to have something different.

The key difference, then, is that in Universe A every decision is completely caused by what happened before the decision – given the past, each decision *has to happen* the way that it does. By contrast, in Universe B, decisions are not completely caused by the past, and each human decision *does not have to happen* the way that it does.

After this description, subjects were asked, “Which of these universes do you think is most like ours?” The vast majority of subjects answered that the *indeterminist* universe (Universe B) is most like ours. Note that the only feature of the universe that is indeterminist is choice. So, the responses indicate that people are committed precisely to the idea that choice is indeterminist. But it would be premature to conclude that people are consistently indeterminist about choice. For, as we will see, different kinds of questions seem to provoke determinist responses.

3.2. *Choice as determined*

A number of philosophers have maintained not only that determinism is true, but also that it's the common view. Hume is characteristically deft on the point:

It is universally acknowledged that there is a great uniformity among the actions of men, in all nations and ages, and that human nature remains still the same, in its principles and operations. The same motives always produce the same actions: The same events follow from the same causes . . .

it appears, not only that the conjunction between motives and voluntary actions is as regular and uniform as that between the cause and effect in any part of nature; but also that this regular conjunction has been universally acknowledged among mankind, and has never been the subject of dispute, either in philosophy or common life (Hume, 1955/1743, pp. 83, 88).

Note here that Hume is not simply saying that people are determinists about physical events; he's saying that people are determinists about motives and voluntary actions. Thus, Hume's claim runs exactly against the trend of the previous section. However, I suspect that there's something deeply right about Hume's claim that people acknowledge psychological determinism.

There has been an enormous body of work investigating how children (and adults) predict and explain behavior. It's a striking fact about this work on "mindreading" that the notion of indeterminist choice is not invoked in any of the prevailing models of how children predict and explain behavior (see, e.g., Gopnik & Meltzoff 1997, Leslie 1995, Nichols & Stich 2003). That is, none of these models maintains that people routinely invoke the notion of indeterminist choice when they predict and explain behavior. It's possible that this has been an outrageous oversight of researchers on mindreading. But another possibility is that the prevailing neglect of indeterminist choice is apt. The notion of indeterminist choice might really be absent from mindreading processes. The prevailing approach to mindreading does make an interesting prediction. According to determinism, psychological duplicates (i.e., those who have all the same psychological states) will always make the same decision when given the same options. Indeterminists maintain that psychological duplicates might make different decisions because of their free will. Thus, if indeterminist choice plays no role in mindreading (as suggested by prevailing accounts), then people should predict that psychological duplicates will always make the same decisions when presented with the same options. It's not trivial to test this, but I have some preliminary data to report.

Experiment 1

In this experiment, subjects were asked to predict whether psychological duplicates will make the same decisions. What we want to know is whether subjects think that there is a 100% probability that psychological duplicates will make the same decision given the same options. Since people reason better with frequencies than with probabilities (e.g. Gigerenzer & Hoffrage 1995), subjects were asked about myriad psychological dupli-

cates. Based on prevailing accounts of mindreading, the prediction was that subjects would respond as determinists, predicting that the psychological duplicates would make the same decisions.

Method

Participants: 30 undergraduate in introductory philosophy classes at the College of Charleston participated in this study.

Procedure: Participants were given questionnaires in a classroom. The questionnaire begins as follows:¹⁰

Background: For all of these questions, you are to imagine that the earth is only one of 10,000 planets that are very similar. All of these planets are governed by the very same laws of nature (e.g., laws of physics, biology, psychology) as on earth. They are also part of very similar solar systems. And many of the individuals on earth have very similar parallel individuals on these other planets. But none of the planets is exactly the same as earth or as any of the other 10,000 planets.

Question 1. On each of the 10,000 planets, there is a person named Jerry Grames. There are some physical differences between all these different Jerrys, but the physical differences are not readily detectible, and everything they have experienced throughout their lives has looked and sounded exactly the same. Indeed, at the psychological level, all of these different Jerrys have been exactly the same up until now. That is, they have all had the same beliefs, desires, thoughts, perceptions, and intentions. Each of them has wanted to learn a new skill, and each of them has just been considering the idea of learning to walk a tightrope. At this moment, the Jerry on earth decides to learn to walk a tightrope. How many of the other Jerrys do you think decided the same thing?

Circle your answer:

All of them Many but not all of them Few or none of them

Explain your answer in detail (use several sentences) along the following lines:

- If you said “all”, explain in detail why you predict that all would make the same decision.

¹⁰ There were subsequent questions on the survey as well, but the subsequent questions don't bear on the question of interest here.

- If you said “many but not all” or “few or none”, explain in detail why you think that not all of them would make the same decision. What would lead some of them not to make the same decision?

Scoring

An answer of “all of them” was coded as determinist, any other answer was coded as indeterminist. All explanations were screened to determine whether the participant abided by the conditions of the thought experiment.

Results

Of the 30 participants, 13 gave explanations that indicated that they failed to abide by the conditions of the thought experiment. Initial analyses thus focused on the remaining 17. Of these 17, 14 gave the determinist response, and this differs significantly from what would be expected by chance alone.¹¹

Given the high rate of compliance failure, it’s worth considering in a bit more detail the 13 subjects who failed to abide by the conditions of the experiment. Two misinterpreted the language – they took “until now” to carry the implicature that something changed. But for the remaining 11 subjects, there was an intriguing pattern – 8 of these subjects offered explanations that adverted to differences between the cases. For example, one wrote, “it is unlikely that they have all had the same experiences and encounters with others . . . that will alter their thought processes.” Of course, this counts as a violation of the conditions of the thought experiment. But what’s more interesting is that this explanation fits in precisely with deterministic explanations. That is, the way they explain why some of the Jerrys will decide differently is by invoking additional causal factors that might have affected the decision. However, since there were so many people who failed to abide by the conditions of the thought experiment, it was important to investigate the matter in a more interactive framework that allowed for correcting the subject.

¹¹ χ^2 goodness-of-fit ($N = 17$) = 7.118, $p < .01$.

Experiment 2

This experiment again explored whether subjects would give deterministic responses in predicting the decisions of psychological duplicates. In this case, though, the experiment was conducted over email so that subjects could be corrected if they failed to comply with the conditions of the thought experiment.

Method

Participants: 8 undergraduate students recruited through an introductory philosophy class at the University of Utah participated in this study.

Procedure: Participants were asked to email the investigator to arrange to complete a survey and follow-up questions over email. Participants were sent Part I below as part of an email message. If their explanation for their response violated the conditions of the thought experiment, they were then sent Part II below.

Part I

Imagine that there is another planet that is really similar to earth. And on this other planet, there are people who are a lot like people on earth. In some cases, the people are almost exactly the same. But no one on either planet knows about the other planet. Now imagine that there is a boy on earth and a boy a lot like him on the other planet. If you could look at them both, you wouldn't be able to tell any differences between them. Let's just call them #1 and #2. There are some small differences between #1 and #2. For instance, #1 has a few more hairs on his head than #2. No one, including the boys themselves, would even notice the difference. Also, some of their insides are a little different. Like, #1's stomach has a different shape than #2's stomach. But no one could tell this either.

Now, even though there are these sorts of differences between #1 & #2, **EVERYTHING THEY THINK, WANT, REMEMBER, AND FEEL RIGHT NOW IS EXACTLY THE SAME.** So, if #1 wants gum, #2 wants gum. If #2 is feeling cold, #1 is feeling cold. If #1 likes basketball, #2 likes basketball. And so on. Okay, #1 & #2 are each trying to figure out what to do today, and each of them has been thinking about climbing a tree. At this moment, #1 decides to climb a tree.

1. Will #2 definitely make the same decision as #1 (namely, to climb a tree) or is it possible that he'll decide something else? Put an X by your answer below:

#2 will definitely decide the same thing as #1.

It's possible that #2 will decide something else.

2. Please explain your answer in detail (use several sentences) along the following lines:

- If you said "definitely decide the same thing", explain in detail why you predict that #2 would definitely decide the same as #1.
- If you said "it's possible that #2 will decide something else" explain in detail why you think that #2 might make a different decision than #1. How might it happen that #2 would decide something different?

Part II

Thanks for the reply. At the risk of trying your patience, I need to point out that there is one problem with your explanation. #1 and #2 are psychological duplicates, so their wants, beliefs, desires, and feelings were all exactly the same. So if that's right, if all their psychological states are the same up until the moment of their decision, is it possible that #2 will decide differently than #1?

YES / NO

Scoring

For part I, the answer "#2 will definitely decide the same thing as #1" was coded as determinist; the answer "It's possible that #2 will decide something else" was coded as indeterminist. For part II, an answer of 'No' was coded as determinist and 'Yes' as indeterminist.

Results

For Part I, 6 responded as determinists, and 2 responded as indeterminists. Both subjects who responded as indeterminists gave explanations indicating that they failed to abide by the conditions of the thought experiment. In both cases, they adverted to differences between the case of #1 and #2 that might allow for different behavior. Both of these participants were then sent Part II. On this question, one of the participants now gave the determinist reply. The other participant sustained his indeterminist answer, and now gave a different justification that appealed

explicitly to free will.¹² Thus, when we consider subjects who either initially abide by the conditions of the thought experiment or who are corrected, we find that 7 out of 8 gave a determinist response when asked to predict the decision of a psychological duplicate. As in experiment 1, this differs significantly from what would be expected by chance alone.¹³

These two experiments came out exactly as the prevailing approach would predict. The appeal to indeterminist choice was extremely rare when subjects were asked to predict the decision of psychological duplicates. This provides some rudimentary support for Hume's claim that people really accept psychological determinism. But of course this is only part of the story. Although people respond as psychological determinists in some conditions, they respond as psychological indeterminists in other conditions. Indeed this tension might be exactly what generates the intuitive problem of free will. On the one hand, it seems plausible to us that our decisions flow directly from our psychology. On the other hand, it seems plausible to us that our decisions are not simply a consequence of past forces.

4. Moral Responsibility

When we turn to responsibility, the issue is not what agency is like, but rather, what agency would have to be like for us to be morally responsible. Again, determinism is central. Compatibilists maintain that determinism is entirely consistent with moral responsibility. Incompatibilists maintain that none of us is morally responsible if our decisions are determined. Folk intuitions have been at the fore of recent discussions of responsibility, and there is considerable dispute about what the folk intuitions are.

Incompatibilists maintain that their view is the natural, intuitive view (Kane 1999, Pereboom 2001, Strawson 1986). However, the view that incompatibilism is intuitive clashes with the extant empirical work on

¹² As it happened, although this student had not taken any philosophy classes before, he had independently studied the free will debate.

¹³ χ^2 goodness-of-fit (1, $N = 8$) = 4.500, $p < .05$, two-tailed.

folk intuitions about responsibility and determinism. In early work, Wayne Viney and colleagues found that people whom they identified as determinists were just as punitive as indeterminists and just as likely to be retributivists (Viney et al. 1982, 1988). This result suggests that determinists sustain typical practices of responsibility attribution. Unfortunately, the measure for identifying determinists was flawed. For the central determinist claim (“I believe in determinism”) was only one item out of seven that contributed to the overall score on the “free will-determinist scale” (Viney et al. 1982). Thus, people who reject the central determinist claim might still end up on the “determinist” end of the spectrum.¹⁴ A different approach has been taken by Eddy Nahmias and colleagues (Nahmias et al. 2005). They present subjects with a description of a deterministic universe and ask whether a miscreant in that universe is morally blameworthy. One of their scenarios goes like this:

Imagine that in the next century we discover all the laws of nature, and we build a supercomputer which can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any future time. It can look at everything about the way the world is and predict everything about how it will be with 100% accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25th, 2150 A.D., twenty years before Jeremy Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 PM on January 26th, 2195. As always, the supercomputer’s prediction is correct; Jeremy robs Fidelity Bank at 6:00 PM on January 26th, 2195.

After reading the scenario, 83% of subjects replied that Jeremy was morally blameworthy for robbing the bank. This looks like striking of evidence against the claim that people are intuitively incompatibilist.

¹⁴ Thanks to David Jackson for drawing my attention to the problems with these measures. In a later study, Viney and colleagues actually attempt to correct the problem (Haynes et al. 2003). Unfortunately, the new study errs in the opposite direction. The criteria used for identifying determinists in this study is too conservative and likely would exclude compatibilists from the set of determinists. Thus it fails to address whether compatibilism is the intuitive view of determinists.

Furthermore, Nahmias and colleagues got similar results using two other scenarios.

Thus, incompatibilists maintain that the folk have incompatibilist intuitions, but the recent results from Nahmias and colleagues indicate that the folk have compatibilist intuitions. Who is right here? In a recent paper, Joshua Knobe and I maintain that *both* parties are right. The folk can be led to report compatibilist intuitions and they can also be led to report incompatibilist intuitions. Our suspicion was that in the absence of emotional triggers, people would respond as incompatibilists. But we also thought that by making the scenario emotionally salient, people would be more likely to respond as compatibilists.

Simplifying somewhat, our experiment had two conditions, an affect-neutral condition and an affect-laden condition. We presented participants with descriptions of two universes, a determinist universe (Universe A) and an indeterminist universe (Universe B) (see section 3.1 for the full descriptions). Participants in the affect neutral condition were asked, “In Universe A, is it possible for a person to be fully morally responsible for their actions?” Participants in an affect-laden condition were given the following question: “In Universe A, Bill stabs his wife and children so that he can be with his secretary. Is it possible that Bill is fully morally responsible for killing his family?” We found that most people in the affect-neutral condition responded as incompatibilists. That is, they tended to deny that it is possible for a person in Universe A to be fully morally responsible. We also found that subjects in the affect-laden condition were significantly more likely to give compatibilist responses than subjects in the affect-neutral condition (Nichols & Knobe forthcoming).

Thus, as with folk intuitions about agency, the situation is complex and interesting. There is no simple answer to the question, “is compatibilism intuitive?” Under some conditions, people’s intuitions tend to compatibilism; under other conditions they tend to incompatibilism. The important subsequent task for the descriptive project is to try to glean the psychological mechanisms that are implicated in generating the intuitions.

5. Mechanisms

In the previous two sections, I've maintained that people have intuitions on all sides of the free will issue. Under different conditions, people respond as determinists, indeterminists, compatibilists, and incompatibilists. From one perspective, this seems like a flatly discouraging result – people look to be utterly confused on the issues. However, if we consider the psychological underpinnings of these intuitions, matters become less chaotic.

Psychological determinist intuitions

In section 3.2, I charted the striking fact that in the mindreading literature, none of the prevailing accounts of how people predict decisions acknowledge a role for the concept of indeterminist choice in decision prediction. Further, I reported some preliminary data indicating that people make deterministic predictions about the decisions of psychological duplicates. The source of these psychological deterministic intuitions is presumably the mindreading system. That system generates decision predictions without deploying the notion of indeterminist choice. There is even a plausible story about why mindreading systems would not be built to invoke indeterminist factors. For in the domain of folk prediction and explanation, it's likely that the notion of indeterminism plays no useful role. Of course folk psychological predictions often go wrong, but indeterminism is no more predictively useful here than it is in folk biology, folk physics, or folk geology. This point is only reinforced if we think of the origins of the mindreading system. The capacity for mindreading is plausibly a product of our biological evolutionary heritage (e.g., Leslie, 1995; Nichols & Stich, 2003). Adding an indeterminist factor into the mindreading system would contribute no additional predictive power, so there would be no call for Mother Nature to graft this into the system.

Compatibilist intuitions

When we turn to compatibilist intuitions about responsibility, another set of mechanisms seem to be implicated. Knobe and I suggested that

affective mechanisms contribute to compatibilist responses in our experiment. Indeed, the specific aim of the manipulation in that experiment was to vary the amount of affective response. However, even if it's true that affect contributes to compatibilist intuitions, there are quite different accounts of how this works. Compatibilist intuitions might result from an affect-generated performance error or compatibilist intuitions might come from an affective competence (Nichols & Knobe forthcoming).

Perhaps the most obvious affect-based explanation is the performance-error model on which affect distorts our normal reasoning about responsibility and leads us to give responses that don't reflect the folk theory. This interpretation fits with a large body of evidence in social psychology that indicates that affect distorts our reasoning processes. Indeed, there is even research indicating that affect distorts our reasoning about responsibility. In an elegant study, Jennifer Lerner and colleagues compared responsibility attributions in affect-laden and affect-neutral conditions. In the affect-laden condition, subjects were shown an upsetting video that aroused negative emotions; in the affect-neutral condition, subjects were shown an emotionally bland video. Lerner and colleagues found that the subjects in the affect-laden condition held agents in unrelated scenarios more responsible (Lerner et al. 1998). Clearly the best explanation of their findings is that the emotions distort the subjects' reasoning about responsibility. Thus, it's natural to suppose that something similar happens when compatibilist intuitions are elicited.

Although the performance error model is the most obvious candidate, there is an important rival. In recent work in moral psychology, several philosophers and psychologists have suggested that normal fluency with moral judgment depends on an affective competence (Blair 1995, Nichols 2004c, Prinz forthcoming). Psychopaths, who lack certain affective tendencies, also show abnormal performance on standard tests of moral judgment. Thus, moral competency plausibly involves an affective competency. This raises the possibility that the normal fluency with responsibility judgments also depends on an affective capacity. On this approach, the affect-neutral condition in our experiment might be thought to bypass the normal affective competence associated with responsibility attributions. The genuine competency for responsibility attribution shows itself in the affect-laden condition.

Both of these affect-based accounts are currently viable. If either is right, then compatibilist intuitions have an importantly different psychological source from determinist intuitions and from incompatibilist intuitions.¹⁵

Indeterminist and incompatibilist intuitions

The source of our intuitions of psychological indeterminism and incompatibilism are harder to identify. One significant point can be drawn from the experiment that Knobe and I did: intuitions about indeterminism and incompatibilism seem to flow from cold cognitive processes. For most subjects in the affect-neutral condition expressed both the incompatibilist view and the view that our universe resembles the universe in which decisions aren't determined.

The fact that indeterminist and incompatibilist intuitions emerge in cold cognition does not take us very far. We certainly do not know which cognitive mechanisms are operative when people arrive at their intuitions about indeterminism and incompatibilism. Indeed, we don't even know whether the implicated mechanisms are domain specific or domain general. It is to be hoped that future work will consider those issues more directly. The one area that has received a modicum of discussion, though, is how we acquire the belief in indeterminist agency. So we can give some attention to the question of acquisition.

The traditional explanation for how we come to believe in indeterminist agency is that it comes from introspection (e.g. Reid 1969 [1788], Holbach 1970 [1770]). When we introspect it does not seem to us that our choices are determined. Introspection fails to reveal any deterministic underpinnings of my decision making. Although this is true, it can hardly be a complete explanation of how we acquire the belief in inde-

¹⁵ One might alternatively try to give an account of the compatibilist intuitions that does not rely on affect. Knobe and I suggested the possibility that compatibilist responses might be driven by a domain specific module dedicated to the attribution of responsibility (Nichols & Knobe forthcoming). To be sure, there are affect-free ways to describe the results. For incompatibilist intuitions are elicited by a task that presents the material in an entirely abstract way, and the cases that attract compatibilist responses involve concrete cases with an individual actor. Perhaps the concrete/abstract dimension is what matters, and not the affect-neutral/affect-laden dimension. I've set aside broader discussion of these options given the space limitations.

terminist choice. The fact that we don't perceive deterministic decision-making processes doesn't yet explain why we would believe that our decisions are *indeterministically* generated. For we often think processes are deterministic even when we don't perceive deterministic causal transactions.¹⁶ Even some of our behaviors have this quality. When my eye twitches, I have no idea what causes it, but this doesn't remotely lead me to think that eye twitches are indeterministically generated. Thus the fact that we don't perceive a deterministic process of decision making must be supplemented to explain the intuition that our decisions are not determined. The natural supplement would be to maintain that we have a standing belief that we *do* have introspective access to all the causal processes underlying our own decision making. If we have a standing belief in such introspective transparency and we fail to introspect deterministic processes, we might infer that there are none.

This combination of presumed introspective-access coupled with introspective failure might provide an explanation for how we come to believe in our own indeterminist agency. But there are several shortcomings. First, we would need evidence that people do in fact have a standing belief that we have introspective access to all the causal processes underlying our own decision making. Second, the account also has to assume that people carry out the required inference to arrive at the view that agency is indeterminist. Third, this account would only apply to oneself. It would need to be supplemented to explain why we think that *other* people's decisions are indeterministic. Finally, this explanation of the indeterminist intuition will not extend easily to explain the acquisition of the incompatibilist belief that determinism precludes responsibility. Why did we end up with a notion of responsibility that is directly tied to the apparent indeterminist feature of our agency? None of these shortcomings is decisive, of course. There are various stories one might tell about how we extend indeterminist choice to others and how we tie responsibility to this notion. But I want to put a rival account on the table.

In Nichols (2004b) I suggested, *very tentatively*, an alternative account on which the acquisition of the belief in indeterminist choice might derive

¹⁶ Indeed, on some views, we *never* perceive the causal powers that we presume to underwrite deterministic processes.

from a prior belief in obligation. According to a famous Kantian argument, we can prove that we have indeterminist choice from the maxim “*ought implies can*” and the fact that we ought to follow the moral law. The idea is that we can’t be obligated to do the impossible, and if determinism is true, it is impossible for us ever to do other than we are determined to do. Thus, if we say that a person *ought* to have behaved differently, this implies that the person *could have done otherwise* (in an indeterminist sense). The suggestion in Nichols (2004b) was that, despite the dubiousness of the Kantian argument as a *proof* of indeterminist choice, it might provide an account of how we come to *believe in* indeterminist choice. There is plenty of evidence that even young children think that people *ought* to behave in certain ways (e.g. Nichols 2004c, Nucci 2001). Indeed, the child applies notions of obligation in a variety of contexts including contexts of moral transgressions (you shouldn’t kick people), conventional transgressions (you shouldn’t eat steak with your hands), and even simple cases of advice (you should put on sunscreen). If children apply some notion of obligation that carries the Kantian implication *could have done otherwise* (in an indeterminist sense), then the child has the essential ingredients for coming to believe that decisions are not determined. Unfortunately, while there is abundant research showing that children apply obligation concepts, there is no evidence yet confirming the idea that children embrace the Kantian maxim.¹⁷

The obligation-based account of how we come to believe in indeterminist choice is far from being confirmed. But it does have a couple of advantages over the introspection-based account. First, whereas the introspection-based account only directly explains how I come to think that *my* decisions are indeterminist, the obligation-based account explains why we attribute indeterminist choice to people across the board. We need no auxiliary hypotheses to explain how I extend the notion of indeterminist choice outside of myself. Second, the obligation-based approach might also provide a basis for explaining how we arrive at the incompatibilist intuition. For it’s plausibly a folk view that people are only morally blameworthy for actions that they *ought not* to have performed;

¹⁷ It is likely that children embrace *some* kind of ought-implies-can view. If you ask whether it was wrong for the paraplegic not to swim to save a drowning victim, children will presumably say that it’s not wrong because he *couldn’t* swim. But it will be harder to show that children think that obligations carry the implication of *indeterminist-can*.

given this view, the Kantian maxim implies that we are only blameworthy for an action if we *could have done otherwise* (see, e.g., Widerker 1991). So the incompatibilist intuition that determinism undercuts responsibility would be a natural inference to draw from the maxim. Thus, on the obligation-based account, we get an attractive economy – the account would help explain how we acquire both the indeterminist and the incompatibilist intuitions. Nonetheless, I should stress, that this is all extremely speculative, especially given that we have no evidence on whether children embrace the maxim that *ought* implies *can*.

Thus, the mechanisms underlying the intuitions concerning free will are extremely diverse. Intuitions of psychological determinism plausibly emerge from mindreading mechanisms. Affective mechanisms seem to facilitate compatibilist intuitions. For indeterminist and incompatibilist intuitions, things are much less clear. But it certainly seems to be something different again. One possibility is that we acquire our beliefs in indeterminism and incompatibilism from a prior notion of obligation. Another possibility is that we acquire the belief in indeterminist choice from introspection and something else again is responsible for our incompatibilist intuitions. In any case, it seems clear that a rich and varied set of psychological mechanisms is at play in generating the intuitions that drive the problem of free will.

6. Beyond description

Before bringing this essay to a close, I want to consider briefly some issues beyond the descriptive project. If the broad contours of the descriptive account are right, then one thing is clear – the folk intuitions do not present a coherent theory of agency or responsibility. On the contrary, the folk seem to have inconsistent intuitions about agency and responsibility. Thus, at least *prima facie*, when we pursue the substantive project, we have to concede that parts of the folk view are simply wrong. *Something* has to go. Determining what should be rejected is too large a question to consider here, but if it's true that the folk views on responsibility are not coherent, we are also faced with a pressing prescriptive question. Should we revise our practices of punishing and blaming people? The issues here are extremely difficult and subtle, and I want to make a very limited point.

In doing prescriptive ethics, the dominant philosophical approach for decades has been the method of reflective equilibrium. When we have intuitions that don't fit together, we try to work through various cases until we arrive at a prescriptive view that achieves some sort of psychological (and hopefully rational) stability for us. In the case of responsibility and determinism, what we've found is precisely that people have inconsistent intuitions on the matter. So when we turn our attention to resolving the inconsistencies, reflective equilibrium is an obvious method to deploy.

The method of reflective equilibrium is typically practiced by trained philosophers. But Knobe and I thought it would be interesting to see how the philosophically naïve would reconcile the conflicting intuitions. Thus, we ran some pilot studies in which philosophically naïve subjects were given our original materials with a few changes. We included both the affect-neutral question ("In Universe A, is it possible for a person to be fully morally responsible for their actions?") and the affect-laden question ("Is it possible that Bill is fully morally responsible for killing his family?"). After the subjects responded, the text explained that people often answer these questions in inconsistent ways and that we were interested in how people would resolve the inconsistency. The results of the pilot study were entirely mixed. About half of the subjects opted for the compatibilist response as their considered view, and half opted for the incompatibilist response.

Thus, lay responses on this matter seem only to emphasize the difficulty of settling the prescriptive question. A further intriguing item can be wrung out of the experiment that Knobe and I did, along with other pilot data. Very few people in our studies maintained that our universe is deterministic. But over the course of several studies and several more pilot studies, twelve people who were given the affect-neutral question maintained that our universe is most like the deterministic universe. Of those twelve, ten either responded as compatibilists on the initial affect-neutral question or responded as compatibilists under conditions of reflective equilibrium.¹⁸ This is hardly conclusive evidence (indeed, it

¹⁸ Eight of these subjects were in a reflective equilibrium study. The other four were in an affect-neutral condition of the original experiments. All four of those subjects

is hardly evidence at all), but it helps to reinforce the suggestive findings from Viney and colleagues (1982, 1988) that determinists tend to settle into compatibilist views about responsibility.

Of course, none of the foregoing shows that we *should* be compatibilists, all things considered. For there are many considerations that should inform our prescription here, not just how the folk resolve the question. Nonetheless, if the folk do not tend to resolve this issue by embracing incompatibilism as their considered view, then we need to be cautious about taking the folk view to establish incompatibilism. For the folk view is fractured, and when the folk attempt to patch things together, incompatibilism is hardly the universal resolution.

7. Conclusion

This paper has promoted a rather ambitious agenda. Drawing on a burst of recent research, I've tried to chart the psychology underlying the central intuitions in the free will debate. I've maintained that the conflicting intuitions that drive the free will problem are generated by importantly distinct psychological mechanisms. The specific proposals that I've made here are no doubt mistaken in important ways. But I am much more sanguine about the general approach. Experimental techniques and evidence will, I'm confident, help us to discover the psychological underpinnings of the folk intuitions that generate philosophical problems. This cannot be done in a global fashion. Rather, it must be pursued on a case by case basis. For it's likely that different philosophical problems implicate different psychological factors. Indeed, in the single case of folk intuitions concerning free will, the extant research seems to reveal a strikingly multifarious set of underlying mechanisms.

responded as compatibilists to the affect-neutral question. In the reflective equilibrium experiments, the eight subjects were evenly split on the initial affect-neutral question. Those who responded as compatibilists to the affect-neutral question sustained their compatibilism under reflective equilibrium. However, two who responded as incompatibilists to the affect-neutral question switched to a compatibilist response under reflective equilibrium.

REFERENCES

- BLAIR, R.
1995 "A Cognitive Developmental Approach to Morality: Investigating the Psychopath." *Cognition*, 57.
- FISCHER, J.
1999 "Recent work on Moral Responsibility." *Ethics* 110, 93-139.
- FODOR, J.
1998 *Concepts*. Oxford: Oxford University Press.
- GIGERENZER, G. & HOFFRAGE, U.
1995 "How to Improve Bayesian Reasoning without Instruction," *Psychological Review*, 102, 684-704.
- GIBBARD, A.
1990 *Wise Choices, Apt Feelings*. Cambridge, Mass: Harvard University Press.
- GOPNIK, A. & MELTZOFF A.
1997 *Words, Thoughts and Theories*. Cambridge, MA: MIT Press.
- HAYNES, S., ROJAS, D., & VINEY, W.
2003 "Free Will, Determinism, and Punishment," *Psychological Reports*, 93, 1013-1021.
- HOLBACH, P.
1970 [1770] *The System of Nature: Or, Laws of the Moral and Physical World*. Translated by H. D. Robinson. New York, B. Franklin.
- HONDERICH, T.
1988 *A Theory of Determinism*. New York: Oxford University Press.
- HUME, D.
1955 [1743] *An Enquiry concerning Human Understanding*. L. Selby-Bigge (ed.). Oxford: Clarendon Press.
- JACKSON, F.
1998 *From Metaphysics to Ethics: a Defence of Conceptual Analysis*. Oxford: Oxford University Press.
2001a "Précis of From Metaphysics to Ethics." *Philosophy and Phenomenological Research* 62(3): 617-624.
2001b "Responses." *Philosophy and Phenomenological Research* 62(3): 653-664.
- KANE, R.
1999 "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism." *Journal of Philosophy* 96: 217-240.
- KNOBE, JOSHUA
2003a "Intentional Action and Side-Effects in Ordinary Language." *Analysis* 63: 190-193.
2003b "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16: 309-324.
- LERNER, J., GOLDBERG, J., & TETLOCK, P.
1998 "Sober Second Thought: The Effects of Accountability, Anger, and Authoritarianism on Attributions of Responsibility," *Personality and Social Psychology Bulletin*, 24, 563-574.

- LESLIE, A.
1995 "A Theory of Agency." In D. Sperber, D. Premack and A. Premack (eds.) *Causal Cognition*. Oxford University Press.
- LEWIS, D.
1972 "Psychophysical and Theoretical Identifications," *Australasian Journal of Philosophy*, 50, 249-58.
- LYCAN, W.
1988 *Judgement and Justification*. Cambridge: Cambridge University Press.
- MACHERY, E., MALLON, R., NICHOLS, S., AND STICH, S.
2004 "Semantics, Cross-Cultural Style." *Cognition*, 92, B1-B12.
- NAGEL, T.
1979 "Subjective and Objective." In his *Mortal Questions*. Cambridge: Cambridge University Press.
- NAHMIA, E., MORRIS, S., NADELHOFFER, T., AND TURNER, J.
2005 Surveying freedom: folk intuitions about free will and moral responsibility. *Philosophical Psychology* 18: 561-584.
- NICHOLS, S.
2004a "Folk Concepts and Intuitions: From Philosophy to Cognitive Science." *Trends in Cognitive Sciences*.
2004b "The Folk Psychology of Free Will: Fits and Starts." *Mind & Language*, 19, 473-502.
2004c *Sentimental Rules: On the Natural Foundations of Moral Judgment*. New York: Oxford University Press.
- NICHOLS, S. AND KNOBE, J.
forthcoming "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions."
- NICHOLS, S. AND STICH, S.
2003 *Mindreading*. Oxford: Oxford University Press.
- NUCCI, L.
2001 *Education in the Moral Domain*. Cambridge: Cambridge University Press.
- PEREBOOM, D.
2001 *Living without Free Will*. Cambridge: Cambridge University Press.
- PRINZ, J.
forthcoming *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- QUINE, W.
1960 *Word & Object*. MIT Press.
- REID, T.
1969 [1788] *Essays on the Active Powers of the Human Mind*. Cambridge, Massachusetts: MIT Press.
- STRAWSON, G.
1986 *Freedom and Belief*. Oxford: Oxford University Press.
- STICH, S. AND J. WEINBERG
2001 "Jackson's Empirical Assumptions," *Philosophy & Phenomenological Research*, 62.
- VARGAS, M.
forthcoming "The Revisionist's Guide to Moral Responsibility," *Philosophical Studies*.

VINEY, W., WALDMAN, D., & BARCHILON, J.

1982 "Attitudes toward Punishment in Relation to Beliefs in Free Will and Determinism" *Human Relations*, 35, 939-49.

VINEY, W., PARKER-MARTIN, P., DOTTE, S. D. H.

1988 "Beliefs in Free Will and Determinism and Lack of Relation to Punishment Rationale and Magnitude." *Journal of General Psychology* 115, 15-23.

WEINBERG, J., NICHOLS, S. AND STICH, S.

2001 "Normativity and Epistemic Intuitions," *Philosophical Topics*, 29, 429-460.

WIDERKER, D.

1991 "Frankfurt on 'Ought Implies Can' and Alternate Possibilities," *Analysis*, 51, 222-224.

Copyright of *Journal of Cognition & Culture* is the property of Brill Academic Publishers and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.