

Journal of Cognition and Neuroethics

Experimental Philosophy, Robert Kane, and the Concept of Free Will

J. Neil Otte

University at Buffalo, The State University of New York

Biography

J. Neil Otte is presently in the Ph.D. program at the University at Buffalo (SUNY) and works in moral psychology, cognitive science, and the history of philosophy. For six years, he was an adjunct lecturer in philosophy at John Jay College of Criminal Justice in New York City. In the last three years, he has been an organizer of the Buffalo Annual Experimental Philosophy Conference, an organizer for a conference on sentiment and reason in early modern philosophy, and has worked as an ontologist for the research foundation, CUBRC.

Acknowledgements

The author wishes to thank Gunnar Björnsson, Gregg Caruso, and Oisín Deery for their supportive comments and conversation during the Free Will Conference at the Insight Institute of Neurosurgery and Neuroscience in Fall of 2014, as well as conference organizer Jami Anderson.

Publication Details

Journal of Cognition and Neuroethics (ISSN: 2166-5087). March, 2015. Volume 3, Issue 1.

Citation

Otte, J. Neil. 2015. "Experimental Philosophy, Robert Kane, and the Concept of Free Will." *Journal of Cognition and Neuroethics* 3 (1): 281–296.

Experimental Philosophy, Robert Kane, and the Concept of Free Will

J. Neil Otte

Abstract

Trends in experimental philosophy have provided new and compelling results that are cause for re-evaluations in contemporary discussions of free will. In this paper, I argue for one such re-evaluation by criticizing Robert Kane's well-known views on free will. I argue that Kane's claims about pre-theoretical intuitions are not supported by empirical findings on two accounts. First, it is unclear that either incompatibilism or compatibilism is more intuitive to nonphilosophers, as different ways of asking about free will and responsibility reveal different answers. Secondly, I discuss how a study by Josh May supporting a cluster concept of free will may provide ethicists with a reason to give up a definitional model, and I discuss a direction future work might take. Both of these objections come from a larger project concerned with understanding the cognitive mechanisms that people employ when they make judgments about agency and responsibility—a project that promises not only to challenge contemporary philosophy, but to inform it.

Keywords

Free will, experimental philosophy, compatibilism, incompatibilism, moral cognition, Robert Kane

Introduction

Studies that employ experimental method to examine non-philosophers' concept of free will have been going on for some time now, and trends in the literature are now forming, giving us reason to re-evaluate some contemporary positions. As an example of how this work lends support to a re-evaluation of the contemporary literature, I take as my focus the influential work of Robert Kane. In this essay, I argue that evidence that nonphilosophers are compatibilists or incompatibilists is presently not forthcoming, that evidence showing that moral judgment affects judgments about happiness, the mental states of others, and knowledge (Beebe & Buckwalter 2010; Knobe ; Phillips, Nyholm, & Liao) might give us pause about accepting Kane's account of self-forming actions, and that work by Josh May supports rejecting a classical account of the concept of free will in favor of a prototype or exemplar account.

I. Kane and the Concept of Free Will

Philosophers often appeal to the ordinary conception of free will. The reason for this is that we're not just interested in any old conception of free will from which we could avoid hard problems and deduce and defend easy platitudes; rather, we're interested in the concept that people actually use, the concept that people employ when they judge an action to have been freely performed, and frequently, the judgment that entails that an action is capable of being praised or blamed. Philosophers who find this train of thought compelling argue that when we judge whether incompatibilism or compatibilism is true or false, it must be judged according to the conception of free will as it is found in the wild.

I use 'concept' here as it is used most often in psychological literature, and not principally as it is often used in other fields like computer science and philosophy. Thus, I take the concept of free will to be the body of knowledge about free will that is characteristically used in the cognitive processes that underwrite our judgments of free will.¹ Popular models of such concepts include definitions, theory-theories, exemplars, and prototypes; and experimental work can help identify whether a particular concept is best captured by one model over others.

The relationship between concepts and problems in philosophy is a perennial issue, but it has received a particular attention in the last few decades. One reason for this focus is that many philosophers—including Frank Jackson, David Lewis, and David Chalmers—have argued that traction can be gained in solving philosophical problems if we identify the structure of broadly shared concepts. The idea is that if the concepts of rational, well-informed people tend to be more or less the same on a given philosophical issue, then their judgments will reflect principles, platitudes, or well-hewn truths that are also prescriptive. This particular view of conceptual analysis is not the only game in town; revisionists concerning free will think that our intuitions should be resisted or rejected in the face of certain considerations about well-known biases or principles. Manuel Vargas is a representative of this position. But revisionism remains a minority view.

Robert Kane's libertarian account of free will is not revisionist—or, at least, it's not intended to be so. Kane frequently begins explaining his account by reaching back into history to cite important moments in free will debates. When he does this, he is interested in seeing what is important to them; this often means asking, what did

1. In other words, a concept of free will is a subset of the knowledge about free will that we store in long-term memory, namely, the part of our knowledge that is used to categorize a particular action as a free action or as a not-freely-performed action (Belohlavek and Klir 2011; Machery 2009).

philosophers think best characterizes what people wanted free will to do? Why do many people appear to think that free will is necessary for morality and practical reason? And what metaphysical conditions are requisite for the possibility of such a status? Kane then attempts to synthesize this information without loss, asking the question: what now must the world be like in order for this concept to be correctly applied? In addition, Kane has denied that he is even a moderate revisionist who would endorse even “pruning” our everyday concept (Vargas 2005). So I think we should interpret Kane’s account as he intends, and hold that his view is meant to capture the ways ordinary people generally think about free will.²

And in Kane’s view, nonphilosophers are incompatibilists, at least until philosophers come along and convince them otherwise. On this view, the man on the street believes determinism and free will, in the sense necessary for moral responsibility, are incompatible. Kane writes,

In my experience, most ordinary persons start out as natural incompatibilists. They believe there is some kind of conflict between freedom and determinism; and the idea that freedom and responsibility might be compatible with determinism looks to them at first like a ‘quagmire of evasion’ (William James) or ‘a wretched subterfuge’ (Immanuel Kant). Ordinary persons have to be talked out of this natural incompatibilism by the clever arguments of philosophers. (Kane 1999, 217).

Kane is not alone in this opinion about what description best characterizes people’s general concept of free will. Many philosophers have proposed that nonphilosophers are best described as natural incompatibilists. Galen Strawson, for instance, writes that incompatibilism describes “just the kind of freedom that most people ordinarily and unreflectively suppose themselves to possess” (Strawson 1986, 30). Similarly, Derk Pereboom writes, “Beginning students typically recoil at the compatibilist response to the problem of moral responsibility” (Pereboom 2001, xvi).

But in addition to holding that people are natural incompatibilists, Kane supposes the concept of free will to be organized in such a way that it could be given sufficient conditions that may be used across a range of cases. This presupposes that

2. Kane is aware that there are many ways in which people think about free will; my target is at least one common among them that is a) “a significant kind of freedom worth wanting” and b) that is incompatible with determinism (Kane 1998, 14–15).

nonphilosophers—except where they are in error—will betray a cognitive model or concept of free will, which conforms to the structure of a definition. If one is judging correctly according to this concept, and all the necessary components of free will are present in a given situation, we should expect this person to conclude that the action was freely chosen and otherwise that the action was not free.

With this metaphilosophical model at work, Kane presents a two-stage theory that is now quite well-known. Two-stage models of free will posit that we have, at the first stage, a capacity for first generating considerations in a nondeterministic way and then, at a second stage, choosing among considerations according to the determination of our will. Such models purport to explain how agents in a given circumstance can have multiple considerations available to them—thus allowing for randomness in the possible paths an action can take—while also accounting for the fact that their decision is brought about by a sufficiently determined will.

Kane cites dissatisfaction with most two-stage theories, which he feels do not go far enough in adequately accounting for the concept of free will. He instead proposes a two-stage model that allows for “dual rational control” or, the ability to do otherwise in precisely the same circumstance (Kane 1985). This involves inserting indeterminacy not only in the moment of the accumulation of alternative considerations, but also in the determination of the decision itself. Although not all decisions will have this indeterminacy and many of our actions will flow directly from a determined character, Kane argues that it is necessary for free will that, at some point in the past, free agents engaged in self-forming actions. In *The Significance of Free Will*, his examples of a businesswoman who decides between self-interest and conscience, and an engineer who decides between his craving for alcohol and his desire to save his marriage and career are intended as examples of such actions. Kane describes these cases as involving “recurrent and connected neural networks” for both sides of the issue, both of which are reflections of the character of the individual. As these networks run their course, a chaotic and amplifying interaction is produced where both networks interfere with each other as they run toward the output of a decision. The result is that “the uncertainty and inner tension that agents feel at such moments are reflected in the indeterminacy of their neural processes” (Kane 1998, 130). Kane believes such a model entails that we may later be responsible for our actions, even where no immediate alternative possibilities exist, provided that our action proceeds from our character, where our character is the product of undetermined self-forming actions.

Kane then argues that the free will debate has been mired in equivocation, since incompatibilists but not compatibilists are committed to ultimate responsibility, a desire

to be the ultimate creators of value and the sources of our own nature (Kane 1996, 58–78). The important question is then whether incompatibilism—particularly the branch committed to a need for ultimate responsibility—best characterizes the substantive issue of free will.

II. Are Nonphilosophers Compatibilist or Incompatibilist?

The short answer is that while the body of evidence has largely suggested that compatibalism best characterizes pre-theoretical intuitions, there is presently no consensus concerning whether incompatibilism or compatibalism best characterizes the ordinary concept of freedom of the will. The reason for this is that different ways of asking about free will provoke different responses. Eddy Nahmias has provided evidence that ordinary intuitions about free will are compatibilist. Nahmias et al. (2006) presented participants with a variety of different scenarios describing deterministic universes and for each scenario, participants were asked whether a person in that scenario acted freely and could be held morally responsible for their action. One scenario described a universe capable of being exactly predicted at any given moment by a supercomputer. In this scenario, a man robs a bank and the participant is asked if the man is morally blameworthy for his action. Most participants (76%), when presented with this information, say yes. In other words, despite the computer’s ability to predict exactly what the man would do, most participants thought that he is still responsible for his action. In follow-up questions, roughly two-thirds of participants claimed that agents in these worlds had free will and slightly more than four-fifths claimed that agents have moral responsibility.

Nahmias takes this to be strong evidence for a compatibilist description of pre-philosophical intuitions. But does characterizing determinism as “a capacity to provide exact prediction of human action” capture what philosophers mean by determinism? If it does, it does so indirectly. Nahmias and colleagues are taking this way of operationalizing free will from Sam Harris’s book, *Free Will*, which employs a scenario like this to motivate the idea that free will is an illusion. But Peter Van Inwagen was closer to the truth when he wrote that determinism, in its most basic formulation, holds that “there is at any instant exactly one physically possible future” (Van Inwagen 1983, 3). Perhaps determinism necessarily entails that a supercomputer could, in principle, have the ability to predict all human action, but the converse doesn’t hold true: in other words, it doesn’t follow that because a supercomputer can predict future human actions that those actions are determined, for perhaps the universe does involve random chance, but not in the relationship between future human action and the predictions of the supercomputer,

which at every instance the computer makes its prediction, fixes a relation between its prediction and some human action. So, one could object that the question Nahmias and colleagues have used to operationalize determinism isn't necessarily getting at determinism, even if it is in the ballpark.

Secondly, libertarians, determinists, and compatibilists among us have friends and spouses who are pretty good at predicting our own behavior. Imagining a superlative capability to predict our actions might just be an imaginative extension of this rather mundane fact for the participants in the study. And since those close to us can often predict what we will say and do without detracting from our capacity to be responsible for those things, why should it follow that an exact prediction should do so?

This hunch, that asking about perfect prediction instead of causation might make a real difference, has already been tested. A paper by Luke Misenheimer (2008) tested whether asking about *perfect prediction* or *complete causation* made a difference in the responses. Subjects were presented with a description of an imaginary world in which people have either, in the causation condition, scientists who have discovered that every action of these beings is caused by things that happened before them, or, in the prediction condition, scientists who have discovered that every action of these beings can be perfectly predicted. Subjects were then asked about an individual being in these worlds who embezzles money, and whether or not they could have done so freely. Misenheimer found that whereas 30% of responses signaled agreement in the causation condition—indicating incompatibilism, 63% of responses signaled agreement in the prediction condition, indicating compatibilism.

This is a large effect, and it clearly supports the hypothesis that language that describes causation explicitly does a much better job at securing incompatibilist responses than does predictive language, which returns compatibilist responses. Why would this be the case? One hypothesis is that causation descriptions are mechanistic and defy what Daniel Dennett has called the intentional stance. Normally, when we deal with agents, we attempt to figure out what mental states they might have, given their environment and their behavior. In particular, we attempt to describe what the agent's desires and objectives are given this behavior. We then attempt to predict what an agent will do on the basis of this analysis. But language involving causation drops to a lower level of abstraction and asks us not to take up the position of evaluating the mental states of an agent, but rather to explain their behavior by reference to something else (e.g., brain states).

In an influential paper, Shaun Nichols and Joshua Knobe (2007) also complicated Nahmias's conclusion. They demonstrated that two ways of presenting the question

of determinism had a considerable effect on whether the response was compatibilist or incompatibilist. They first described two universes: Universe A, in which every decision “is completely caused by what happened before the decision—given the past, each decision has to happen the way that it does,” and Universe B in which “decisions are not completely caused by the past, and each human decision does not have to happen the way it does” (5). In two additional conditions, they described an *abstract* case, in which they solicited an answer to the question: “In Universe A, is it possible for a person to be fully morally responsible for their actions?,” whereas in a *concrete* case they ask whether or not Bill, a man who killed his wife and children in a fire to sleep with his secretary, is morally responsible for the killing.

Nichols and Knobe first asked directly: which universe, A or B, is more like our own? Participants overwhelmingly choose B, the indeterminist universe. They then asked if, in the concrete condition, Bill was morally responsible for the killing of his wife in Universe A, to which 72% of participants said yes—a finding that confirms the previous findings of Nahmias and colleagues. However, in the abstract condition, 86% of participants answered that a person could not be responsible in Universe A.

This finding has held up, even across cultures. Sarkissian et al. (2010) reports finding that in four distinctive cultures there was a consensus among respondents for two theses: a) that we live in an undetermined universe and b) that moral responsibility is not compatible with determinism. For their study, Sarkissian and his colleagues used a total of 231 undergraduate students from the United States, Hong Kong, India, and Columbia, each group divided roughly in half by sex. They were given the descriptions of universes A and B and asked the same questions. Among the four cultural groups, it was found that there was little difference. This is highly surprising. Researchers have found evidence that culture has a large effect on notions of moral responsibility (Miller & Turnbull 1986), what it means to be an individual (Markus & Kitayama 1991), and even what kinds of fallacies we fall for (Nisbett 2003). But Sarkissian et al. (2010) supports the view that whatever psychological mechanisms are guiding the distinctive incompatibilist judgments in the abstract cases and the compatibilist judgments in the concrete cases, they do not vary widely from culture to culture.

This dramatic contradiction between the abstract and the concrete cases has been a major preoccupation in the experimental work on free will and theories abound. Initially, Nichols and Knobe thought that two distinct psychological processes were driving judgments concerning free will in contrary directions. They hypothesized that in the abstract case, people were moved to say individuals in a causally determined universe did not have free will, but as a case became more concrete, a distinct process leads people to

say a person is, in fact, responsible, whether they live in a causally determined universe or not. Their initial interpretation was that our ordinary judgments are incompatibilist, but that an affective response in the concrete condition leads to performance error.

If this interpretation is correct, then our judgments appear systematically produced in no small part by feeling states. This would give us reason to doubt the two-stage theorist's proposal that our everyday concept of free will can be assessed by a strict consideration of facts regarding agency (e.g., what the agent knew, whether their actions flowed from their character, whether they had alternative possibilities open to them, etc.). Moreover, it would then provide reason to be suspect of Kane's particular description of self-forming actions, whose descriptions involve cases of internal moral conflict and the creation of meaningful contributions to the character of the agent. Such descriptions stir our empathy for an agent, and may contribute to the sympathy or disgust we feel at their decision, and if this hypothesis is correct, these moral sentiments may lead us to attribute an exculpatory or inculpatory status. In other words—to the skeptic's rejoicing—it would appear on this interpretation that Kane's descriptions of self-forming actions are playing to an error theory of free will judgments rather than revealing a deeply latent principle in everyday psychology.

Fortunately for Kane, the evidence has not borne out this conclusion and Knobe has since retracted this proposal. The initial conclusion that affect was the driving force behind compatibilist intuitions in the concrete cases has been challenged and largely set aside, despite its initial plausibility. A recent meta-analysis of twenty-nine studies shows that high affect cases do indeed produce compatibilist responses, but that this response is actually quite small ($d = .18$) (Feltz & Cova 2012).

Nahmias and colleagues have also questioned Knobe's original study for its language. They have found evidence that when the wording that describes determinism is phrased in mechanical or reductionist terms, this can lead to a *bypassing error*, whereby participants assume that the person in the scenario can not be acting according to their own motivations. Nahmias and Murray (2010) presented participants with the following statement: "In Universe A, what a person believes has no effect on what he or she ends up being caused to do." Surprisingly, they found a majority of participants tend to agree. In other words, participants likely infer that causal determination prohibits individuals from acting according to their beliefs and desires. Nahmias has argued that this may be one reason why the free will debate is so bedeviling, namely, incompatibilists will frequently give descriptions of determinism that present it alongside a mechanical or materially reductionist picture, and this description occasions the bypassing error and gives the appearance of a strong incompatibilist response. Likewise, the incompatibilism

that Knobe and Nichols' original study may have found could be the result of bypassing due to the wording of the vignette, rather than a reflection of a more general cognitive process reflecting typical judgments about determinism.³

More recently, Knobe and Nahmias have offered two opposing theories to explain the available evidence to date (Machery & O'Neill 2014, 69). Whether either is correct, time will tell. But the foregoing discussion here should be sufficient to highlight the pitfalls of assigning nonphilosopher's to one camp in this debate or another on the basis of anecdotal evidence. And while there are now repeated findings in the experimental literature, we should pause to note that what is of interest to philosophers doing this experimental work is not merely what label to attach to survey responses. Rather, they are using experimental methods in partnership with traditional philosophy to both test and generate hypotheses about the underlying structures of the concepts involved. The debate concerning free will is genuinely complex and an easy answer is likely to be elusive. It is likely that there will be some canonical descriptions of determinism that will not bear out certain responses, and other canonical descriptions that will.

III. Alternative Concept Models

I turn now to a second part of Kane's claim: that the ordinary concept of free will is best characterized by a series of necessary conditions. Such a concept model is often called a definitional model, or the classic model of concepts. A frequent example philosophers invoke is the concept of a bachelor, which includes two components: male and single. When these two components are present, then there is a categorization under bachelor and when one or more is absent, then there is no categorization under the heading bachelor. Similarly, Kane's account of free will holds that barring error, we should expect people to assert under some conditions that an action is free and that an agent is morally responsible, and to withhold this judgment when those conditions are not present. How much of the specifics of Kane's account are required to be in these conditions is, I think, up for debate. However, it appears Kane must hold that at a minimum, these conditions should include that an agent has alternative possibilities present at the time of action or previously, and that an agent has a capacity to act according to what they deem to be their own best reasons. When these two conditions are present—perhaps with some addition—then we should expect to see a majority of respondents declaring that an

3. However, evidence of bypassing has been recently challenged by Gunnar Björnsson and Derk Pereboom (2014), who provide evidence that the wording of Nahmias's studies can be understood in ways that involve *passing through* rather than *bypassing* the agent's decision, desires, and beliefs.

agent acted freely and can be morally responsible for an act, and when one or more of these conditions is absent, this declaration should not be found.

In an original study, Josh May has found evidence for a different concept model characterizing free will. May (2013) provides experimental evidence motivating a cluster theory of free will, according to which components of the concept work as features of free will rather than as necessary conditions. These findings support a challenge to traditional philosophical accounts of free will that frequently assume that people work from tacit definitions of free will—definitions that philosophical accounts of free will are often depicted as illuminating rather than revising.

In his study, May tested two aspects of free will that have been critical in the literature: ensurance and liberty. Ensurance is that feature of free will that is marked by the agent's control over their actions. One can think of ensurance as the capacity to act based on her own mental states. Liberty is marked by having alternative possibilities, where an agent has the power to make a genuine choice between more than one option. May hypothesized that non-philosophers would be highly moved to respond that a person acted freely if she had both liberty and ensurance. In cases where both factors were missing, there would be little assertion that she had free will, but in cases where only one factor was present, May predicted that judgments would be mixed.

Using a factorial design, May employed four vignettes and randomly assigned non-philosopher participants. Each vignette featured two paragraphs, where each described a universe that is re-created over and over again from some initial conditions. In the first paragraph, either liberty was present or it was not. When liberty was present, the scenario holds that the laws of nature "needn't" cause the exact same events to happen again, but when liberty was absence, the laws of nature "must" cause the exact same events to happen again. In a second paragraph, a subject named Jill is described as either, in the ensurance condition, "deliberating and deciding" to steal a necklace or, in the no ensurance condition, being "brainwashed to have a powerful urge" to steal a necklace. Each participant in the study was then asked whether "Jill stole the necklace freely"⁴ (May 2014, 10).

As expected, May found that the mutual presence of ensurance and liberty encourages judgments that Jill acted freely,⁵ their mutual absence finds a nearly complete

4. Responses were recorded using a 7-pt Likert scale, with "disagree completely" and "agree completely" at 1 and 7 respectively and a middle value of "in-between."

5. $\bar{x} = 6.6$; $\sigma = .89$; total in agreement: 96%.

absence of judgment that Jill acted freely,⁶ but their individual presence finds mixed results and a lack of consensus.⁷ His findings support the hypothesis that the concept of free will has a cluster structure, whose features motivate judgment independently of one another.

Definition models of the ordinary concept of free will can't explain why the presence of ensurance or liberty, in the absence of the other, would result in sizeable minorities of participants responding that an agent had acted freely. The response that this large minority of participants (49% in one case) were simply in error in the case where one feature is present but not the other is unsatisfactory, for the definition model holder can't say this large minority is in error while simultaneously wishing to argue that they are merely clarifying a position rooted in ordinary cognition about free will. The usual response from a two-stage theorist is that a case like the brainwashed Jill case is simply not a case of free will because Jill's actions are not rightfully hers due to a disruption in the flow of actions from her character, her motivations, her reasons, or her desires and beliefs. What this response doesn't account for, given this data, is the fact that this disruption leaves participants without consensus about the free will of an agent, rather than a consensus for the lack of free will of an agent.

These results motivate a cluster theory, but it isn't yet clear from this study which concept model best describes judgments concerning free will. What May has found evidence of are typicality effects in our judgments concerning free will, according to which purported cases of action are not either free or not free, but are rather more or less clearly free according to how much they share certain features common to our concept. Much of the research on typicality and concepts began in earnest in the 1970s with the work of Eleanor Rosch, who found evidence that categories are graded. Both penguins and robins, for instance, are birds, but a robin is more like a bird than a penguin. This gradation also allows psychologists to explain how *systematically* inconsistent people can be in their reasoning. For instance, participants in a study may assent that a dentist's chair is a chair and that chairs are instances of furniture while denying that a dentist chair is a piece of furniture. This intransitivity would lead one to think that people were simply in error about furniture if it was thought that people really had concepts that took the form of definitions, but if we think of the concepts for each of these items as clusters of features, then we can see that they have overlapping characteristics in some areas that

6. $\bar{x}=3.17$; $\sigma=2.25$; total in agreement: 30%.

7. No ensurance, but liberty: $\bar{x}=3.98$; $\sigma=2.25$; total in agreement: 39% and no liberty, but ensurance: $\bar{x}=4.67$; $\sigma=2.32$; total in agreement: 49%.

do not overlap in others, so this inconsistent triad is actually consistent at the level of conceptualization.

Two important models for thinking about cluster theories include exemplars and prototypes. Exemplar models hold that when we categorize an instance under a concept, we look in long-term memory at a collection of the best possible instances of that concept. We then measure the similarity of a given instance to the set of these examples and if a threshold for similarity is achieved, then that instance is categorized under the concept. Prototypes, on the other hand, are best thought of as collections of property data rather than as mental representations of exemplars. James Hampton's formal prototype model is a useful for thinking about prototypes (Hampton 1993). May does not cite Hampton's model but his thinking about the issue is captured by this formalism quite well. Hampton's model includes a similarity measure and a decision rule. The similarity measure, $S(x, C)$, of an instance x to a category C is defined in terms of values $w(x, i)$, which includes the weight of the value (e.g., ensurance) possessed by x for attribute i of the prototype (e.g., free will).

$$S(x, C) = \sum_i w(x, i)$$

This measurement of similarity can then be used in a decision rule, which states that if the similarity measure (S) of an instance (x) to a category (C) is greater than some threshold (t) then that instance is a member of the category.

$$S(x, C) > t \Rightarrow x \in C$$

In coming into contact with objects in the world, we extract from them statistical information concerning how regular a class of properties are associated with objects of a kind. Birds, for instance, can often fly and do so in groups. Birds typically dive into water and build nests; they have beaks, feathers, and forward-facing eyes and are often small. This property description then gets an ordering based on weights. Feathers and beaks have a heavier weight, whereas flying in groups and relative size have lighter weights. This allows us to account for how humming birds and ostriches can both be birds, while being less typical birds than common songbirds.

This way of approaching the question of free will makes it clear that May's study is really only the beginning of a promising empirical approach. Future studies should look not only at ensurance and liberty as attributes whose weights are measured, but also at the many components philosophers have long thought—sometimes controversially—to be essential to exercising free will. In addition, May does not discuss the thesis of Edouard

Machery's book, *Doing without Concepts* (2009). Machery proposes a heterogeneity thesis concerning concepts, which holds that most categories (e.g., free actions) are represented by several concepts that belong to kinds that have little in common. Each of these kinds, which may include prototypes, exemplars, and theories, can be involved in a variety of cognitive processes, including learning, recalling, revising, induction, and judgments. If this thesis is true, it implies that the concept of free will may well consist of distinct bodies of knowledge rather than as a single body of knowledge. This also means it is possible that future experimental philosophers will discover that Kane's significant conception of free will does reflect some bodies of knowledge regarding free will, while failing to adequately capture others.

Conclusion

I have argued that presently available studies on the shared concept of free will do not provide clear empirical support that people are generally incompatibilist; rather, different cognitive processes—reflected in different ways of setting up individual questions—produce different results. Whether or not some of these different ways of asking questions are relevant is a question for further reflection and argument. I have also argued that there may be reason—however under-supported—to doubt the intuitive force of Kane's notion of self-forming actions, due to the affect such descriptions produce, and the role such affect might have in over-riding our usual responses. Finally, I have argued that May provides compelling evidence that the structure of nonphilosopher's concept of free will is distinct from the structure described by Kane. A similar argument has been made now by many others, including Stephen Stich, Alvin Goldman, and Mark Johnson, who have each argued that the demise of the definitional model should give us reason to reject normative ethical theories that presuppose nonphilosophers to represent normative categories using necessary and sufficient conditions. Further evidence may continue to demonstrate a similar argument for philosophical accounts of action that presuppose a definitional model.

References

- Beebe, James, and Wesley Buckwalter. 2010. "The epistemic side-effect effect." *Mind & Language* 25 (4): 474–498.
- Belohlavek, Radim, and George J. Klir. 2011. *Concepts and Fuzzy Logic*. Cambridge: The MIT Press.
- Björnsson, Gunnar, and Derk Pereboom. 2014. "Free Will Skepticism and Bypassing." *Moral Psychology*, Volume 4, edited by Walter Sinnott-Armstrong, 27–36. Cambridge: The MIT Press.
- Feltz, Adam, and Florian Cova. 2012. "When and how affective reactions impact judgments about free will and determinism: A Meta-analysis." Unpublished manuscript.
- Kane, Robert. 1985. *Free Will and Values*. New York: SUNY Press.
- Kane, Robert. 1998. *The Significance of Free Will*. New York: Oxford University Press.
- Kane, Robert. 1999. "Responsibility, luck, and chance: reflections on free will and indeterminism." *Journal of Philosophy* 96 (5): 217–240.
- Hampton, James A. 1979. "Polymorphous concepts in semantic memory." *Journal of Verbal Learning and Verbal Behavior* 18 (4): 441–461.
- Hampton, James A. 1996. "Testing the prototype theory of concepts." *Journal of Memory and Language* 34: 686–708.
- Machery, Edouard. 2009. *Doing without Concepts*. New York: Oxford University Press.
- Machery, Edouard, and Elizabeth O'Neill. 2014. *Current Controversies in Experimental Philosophy*. London: Routledge.
- May, Joshua. 2014. "On the very concept of free will." *Synthese* 191 (12): 2849–2866.
- Markus, Hazel, and Shinobu Kitayama. 1991. "Culture and the self: implications for cognition, emotion, and motivation." *Psychological Review* 98 (2): 224–253.
- Miller, Dale T., and William Turnbull. 1986. "Expectancies and interpersonal processes." *Annual Review of Psychology* 37 (7): 233–256.
- Misenheimer, Luke. 2008. "Predictability, causation, and free will." Accessed 6/10/14: http://philosophy.berkeley.edu/file/551/misenheimer-free_will.pdf.
- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner. 2006. "Is incompatibilism intuitive?" *Philosophy and Phenomenological Research*. LXXIII (1): 28–53.

- Nahmias, Eddy, Justin Coates, and Trevor Kvaran. 2007. "Free will, moral responsibility and mechanism: Experiments on folk intuitions." *Midwest Studies in Philosophy* XXXI (1): 214–242.
- Nisbett, Richard. 2003. *The Geography of Thought*. New York: Free Press.
- Park, John Jung. 2011. "Prototypes, exemplars, and theoretical & applied ethics." *Neuroethics* 6 (2): 237–247.
- Pereboom, Derk. 2001. *Living without Free Will*. Cambridge: Cambridge University Press.
- Rosch, Eleanor. 1975. "Cognitive representations of semantic categories." *Journal of Experimental Psychology: General* 104 (3): 192–233.
- Sarkissian, Hagop, et al. 2010. "Is belief in free will a cultural universal?" *Mind & Language* 25 (3): 346–358.
- Strawson, Galen. 2010. *Freedom and Belief*. New York: Oxford University Press.
- Van Inwagen, Peter. 1983. *An Essay on Free Will*. Oxford: Clarendon Press.
- Vargas, Manuel. 2005. "The revisionist's guide to responsibility." *Philosophical Studies* 125 (3): 399–429.