

Persons, minds, and morals

Philip Robbins*

Abstract

The concept of a person is, or ought to be, central to the philosophy of social science. Traditionally, investigations into the structure of this concept have been conducted on a largely *a priori* basis. Recently, however, with the advent of research in experimental philosophy, the methodological landscape has begun to change (Knobe et al., 2012). In this paper I report findings from two experimental studies exploring how laypeople think about the relationship between two dimensions of two aspects of personhood: mindedness (cognition, affect) and moral status (agency, patiency).

The concept of a person is, or ought to be, central to the philosophy of social science. Traditionally, investigations into the structure of this concept have been conducted on a largely *a priori* basis. Recently, however, with the advent of research in experimental philosophy, the methodological landscape has begun to change (Knobe et al., 2012). This paper is intended as a contribution to this new trajectory of research, which adopts empirical methods from the social sciences to explore concepts of abiding philosophical interest.

In what follows I report findings from two experimental studies investigating how laypeople think about the relationship between two facets of personhood: mindedness and moral status. Each of these facets of personhood has a two-part structure. Mindedness can be factored into affective or ‘feeling’ capacities, like pain and pleasure, and cognitive or ‘thinking’ capacities, like reasoning and planning. Moral status likewise comes in two forms: agency and patiency. Moral agents commit right or wrong actions and are held accountable as such (e.g.,

*Department of Philosophy, Strickland 426, University of Missouri, Columbia, MO 65211, USA. Email: robbinsp@missouri.edu.

praised or blamed), whereas moral patients are on the receiving end of right or wrong actions and are given consideration on that basis (e.g., sympathized with, offered protection). In contrast to the two dimensions of mindedness, moral agency and moral patiency are not so much capacities as roles played by individuals in the context of social interaction, indeed, roles that tend to be seen as mutually exclusive (Gray & Wegner, 2009, 2011).

Two hypotheses were tested. According to the first hypothesis, people think of moral patiency as more closely linked to sentience (affect) than intelligence (cognition). According to the second hypothesis, people think of moral agency as more closely linked to rationality (cognition) than empathy (affect). The first hypothesis was supported by the data; the second hypothesis was not. Before getting to the studies themselves, though, some stage setting is in order.

1 Minds and morals: A dimensional approach

In philosophy and psychology alike, it is customary to distinguish between thinking, or cognition, and feeling, or affect (Robbins & Jack, 2006; Robbins, 2008). Whether common sense, or folk, thinking about the mind incorporates this distinction, however, is a matter for empirical investigation. Recent work in social psychology and experimental philosophy suggests that it does. Initial evidence for this comes from research on ‘mind perception’ using online survey methodology (Gray et al., 2007; Waytz et al., 2010). In one large-scale study, for instance, participants rated a cast of characters (e.g., a normal adult, a child, a dog, a robot) pairwise on a range of mental capacities. On each trial, participants were shown pictures of two characters, each accompanied by a brief verbal description, and asked which character was more likely to possess a given capacity. After determining a mean rating for each character on each capacity, correlations between the different capacities across characters were computed, and those correlations were then subject to factor analysis. The results of this analysis revealed a distinction between two ‘dimensions of mind’: affective capacities, such as pain and joy (collectively labelled ‘Experience’), and cognitive capacities, such as self-control and planning (collectively labelled ‘Agency’). Participants’ ratings on these two dimensions varied significantly across charac-

ters. Some characters (e.g., a baby, a frog) scored relatively high on feeling and low on thinking, while others exhibited the opposite pattern (a robot, God).

Participants in the study were also asked to make comparative judgments of moral status, both with respect to moral agency (“If both characters had caused a person’s death, which one do you think would be more worthy of punishment?”) and moral patiency (“If you were forced to harm one of these characters, which one would it be more painful for you to harm?”). These two aspects of morality were strongly positively correlated with different dimensions of mind: agency with cognition, and patiency with affect. For example, characters scoring high on cognition and low on affect (like God) were seen primarily as moral agents, whereas affectively rich but cognitively poor characters (like infants) were seen primarily as moral patients (Gray & Wegner, 2009).

Supporting evidence for the idea that different dimensions of mind are perceived as differing in moral significance comes from a study reported by Knobe and Prinz (2008). In this study, participants read a story about a fictional character interested in the psychological capacities of fish. In one version of the story, the character was described as being curious about how well the fish could remember the location of food sources in their habitat; in a second version, he wanted to know whether fish could feel pain. After reading one version of the story or the other, participants were asked to explain why the character described might want to have this information. The pattern of responses in each condition varied dramatically. In the cognitive (‘memory’) condition, typical answers referred to the character’s interest in predicting, explaining, or controlling the behavior of the fish. In the affective (‘pain’) condition, the primary focus was on the character’s concern about the welfare of the fish. Consistent with the results reported in Gray et al. (2007), these findings suggest that judgments of moral patiency are specifically linked to judgments about whether something can feel, and not to judgments about whether it can think.

The studies described above, though suggestive, take us only so far. Gray et al.’s (2007) study of mind perception provides merely correlational evidence for the dimensional structure of mind perception in relation to moral perception; it does not speak to the possible causal dependence of the latter on the former. The study by Knobe and Prinz (2008), while addressing the effect of

mind perception on moral consideration, does so only indirectly, by varying the perceived psychological state of a hypothetical observer, rather than that of a target. Their study also does not address the question of how people think about the psychological dimensions of moral agency; it only speaks to the issue of moral patiency. The pair of studies described below, whatever other flaws they might contain, do not suffer from these limitations.

2 Experimental studies

To investigate how people think about the relationship between the different dimensions of mindedness and morality, we conducted two studies using an online survey website to collect data. Both studies used a 2×2 design with the independent variable being the type of mindedness present (high or low affect, high or low cognition) and the dependent variable being moral status (patiency in the first study, agency in the second). Participants were presented with a brief scenario describing a hypothetical character, and then asked to provide an initial assessment of the character’s moral status based on the limited information supplied. Afterwards participants were given further information about the character’s psychological profile and asked to update their assessment of the character’s moral status. This was done in order to supply a more sensitive measure of the effect of mindedness information on moral perception.

2.1 The patiency study

For this study, participants read a document voicing concerns about the treatment of lobsters. After reading the document, they were asked whether it was morally wrong to subject lobsters to rough treatment, using a six-point Likert scale (1 = strongly disagree, 6 = strongly agree). Afterwards they were presented with one of four hypothetical scenarios in which scientists had discovered the actual psychological capacities of lobsters, ranging from low intelligence and low sentience to high intelligence and high sentience. In terms of the dichotomous structure of mindedness, then, intelligence and sentience served as proxies for the ‘thinking’ and ‘feeling’ dimensions, respectively. Participants were subsequently asked to make a second patiency rating based upon this new

information, using the same six-point scale. This was done in order to determine the effect of mindedness information on judgments of moral patiency.

The details of the study are as follows. Across all conditions, participants read an actual press release from 2006 announcing that Whole Foods Market, a high-end grocery chain, had decided to stop selling live lobsters due to concerns about the potentially harmful effects of handling and storing the animals prior to sale. After reading this document, participants were asked to make an initial judgment of moral patiency:

Given your own beliefs about lobsters (as opposed to those of the management of Whole Foods, for example), to what extent would you agree or disagree with the following claim?

It is morally wrong to subject farmed lobsters to rough treatment.

Responses to this first question provided a baseline measure of patiency judgment, for later comparison with the judgment made after mindedness information had been added.

As per the 2×2 between-subjects design, participants were randomly assigned to one of four conditions in the table below (see Table 1):

	<i>High sentience</i>	<i>Low sentience</i>
<i>High intelligence</i>	Maximal	Intermediate (I>S)
<i>Low intelligence</i>	Intermediate (S>I)	Minimal

Table 1. *Patiency factors: Intelligence \times Sentience.*

In the high intelligence/low sentience condition, for example, participants read the following:

Now, imagine that scientists make an interesting new discovery about lobsters. After many years of extensive research, they determine that lobsters are much more intelligent than previously believed. For example, it appears that lobsters have elaborate strategies for foraging, advanced navigational skills, and excellent memories. They are even capable of sophisticated problem solving and planning.

At the same time, however, scientists also discover that lobsters are incapable of feeling. Its not just that lobsters do not experience

pains, pleasures, emotional feelings, or other sensations in anything like the way that people do — it turns out that lobsters do not experience these states at all.

In short, it turns out that lobsters can think, but they cannot feel.

In the low intelligence/high sentience condition, participants read:

Now, imagine that scientists make an interesting new discovery about lobsters. After many years of extensive research, they determine that that lobsters have very limited intelligence. Their memories are extremely poor, and they are incapable of simple planning or problem solving. Because their navigational skills are so weak, they rely on crude foraging strategies to find food.

At the same time, however, scientists also discover that lobsters are much more sensitive creatures than previously believed. For example, it appears that lobsters experience pains, pleasures, emotional feelings, and other sensations in much the same way that people do.

In short, it turns out that lobsters cannot think, but they can feel.

The maximal (high intelligence/high sentience) and minimal (low intelligence/low sentience) conditions followed this same format.

Finally, participants across all conditions read:

Assuming scientists had *actually* made this discovery, to what extent would you then agree or disagree with the following claim?

It is morally wrong to subject farmed lobsters to rough treatment.

This second question, a variant of the first (baseline) measure, was designed to tap respondents' intuitions about the relevance of mindedness to moral patiency.

As noted above, we hypothesized that patiency judgments would depend on whether the target was described as sentient, but not on whether it was described as intelligent. Accordingly, we expected to see an effect of the mindedness manipulation in the feeling (affect) dimension, but not in the thinking (cognition) dimension.

The study collected responses from 112 participants, 63% of them female. The mean age was 27 ($SD = 9.66$). As predicted, a two-factor ANOVA revealed

a significant main effect of the mindedness manipulation on patience, but only in one of the two dimensions. (For similar results from a pilot version of this study, using a different design, see Jack & Robbins, 2006.) Patience judgments varied as a function of whether lobsters were described as sentient, $F(1, 108) = 4.11, p < .05$, but there was no effect for intelligence (see Table 2 and Figure 2), and no interaction effect. Likewise, there was a strong effect of the sentence manipulation on the change in patience judgments relative to baseline, $F(1, 108) = 29.45, p < .001$ (see Table 3 and Figure 3). No analogous effect was observed with the intelligence manipulation; nor was there an interaction effect. These results were confirmed by nonparametric testing, using the Mann-Whitney U test with $\alpha = .05$.

	<i>High sentence</i>	<i>Low sentence</i>
<i>High intelligence</i>	4.11 (1.571)	4.00 (1.440)
<i>Low intelligence</i>	4.25 (1.713)	3.14 (1.604)

Table 2. Mean patience: *Intelligence* \times *Sentience* (SD in parenthesis).

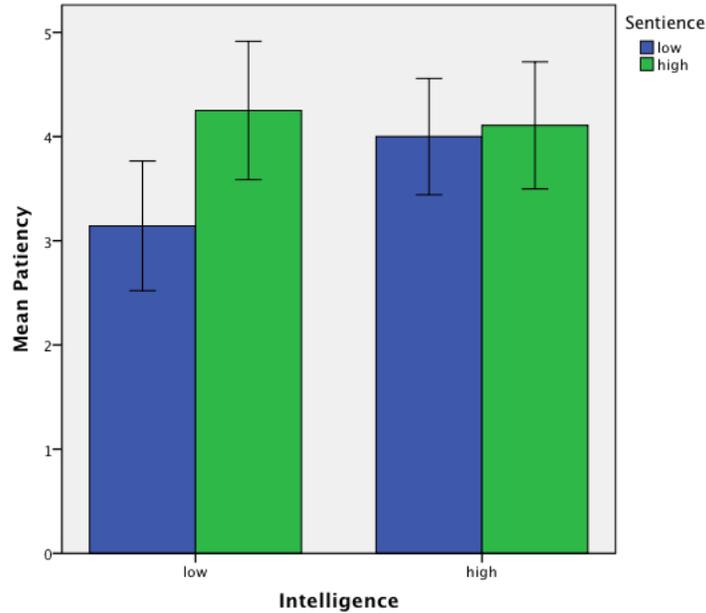


Figure 2. Mean patience.

	<i>High sentence</i>	<i>Low sentence</i>
<i>High intelligence</i>	.43 (1.230)	-.36 (.989)
<i>Low intelligence</i>	.50 (1.036)	-1.11 (1.370)

Table 3. *Mean change in patency: Intelligence \times Sentience.*

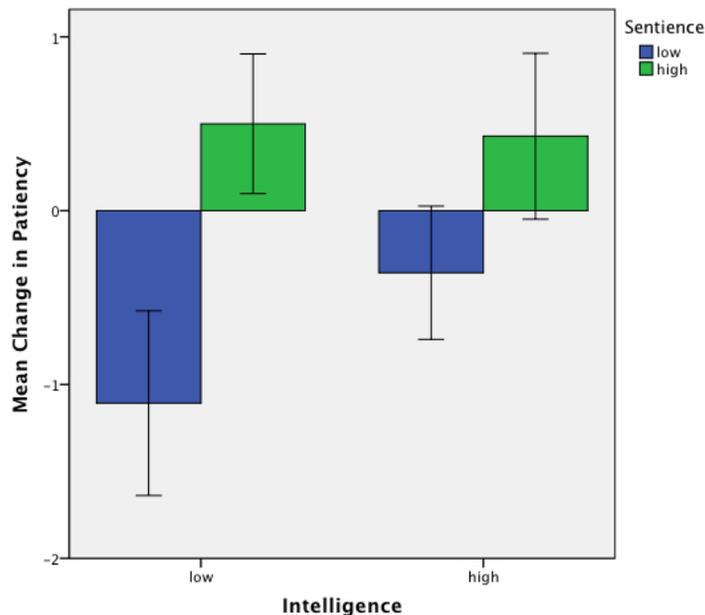


Figure 3. *Mean change in patency.*

A word of caution about the interpretation of these data: It is tempting to look at the patency and change in patency scores across the four conditions and conclude that the attribution of patency is sensitive to information about the target’s intelligence, at least to some extent. In the two low sentience conditions, for example, mean patency ranges from 3.14 (low intelligence) to 4.00 (high intelligence), a contrast that is comparable in magnitude to the contrast between the two low intelligence conditions, where it ranges between 3.14 (low sentience) and 4.25 (high sentience). Given that there was no interaction effect, such a conclusion is unwarranted. It’s worth noting, however, that though no interaction effect was detected at the chosen level of significance ($\alpha = 0.05$), there was an interaction effect at $\alpha = 0.1$ ($p = .098$) — suggesting that such an effect might be detected at $\alpha = .05$ in a larger sample. The same point applies to the effect of the manipulation on change in patency, where the interaction between factors nearly reached significance ($p = .065$).

As far as the commonsense conception of moral patency goes, however, it seems that the philosopher Jeremy Bentham got it right when he wrote, con-

cerning the moral status of animals: “The question is not, Can they *reason*? nor, Can they *talk*? but, Can they *suffer*?” (Bentham, 1789). In other words, what matters most for moral consideration is the capacity to feel, not the capacity to think — or so the folk tend to assume. Whether this view of moral patiency is correct, of course, is a separate issue (but see Bernstein, 1998, for defense of a similar position).

2.2 The agency study

Our second study was designed to mirror the first, albeit with a focus on moral agency rather than moral patiency, and with a different dimensionalization of mindedness. Rather than factoring mindedness into intelligence and sentience, we zeroed in on a pair of psychological capacities that seem potentially important to moral responsibility: rationality and empathy.

The rationale for this choice of factors was three-fold. First, there is the idea that the capacity for sophisticated practical reasoning is a prerequisite for intentional agency in general, and hence moral agency in particular. This is suggested by the fact that persons lacking this capacity, such as floridly psychotic or severely retarded individuals, are not held to the same moral standards as the rest of us. Thus, practical rationality seems essential to moral agency. Second, it seems intuitively obvious that the affective or feeling capacities in general do not suffice for moral agency. If they did, then moral agency would be seen as much more pervasive than it actually is; for example, it would be commonly attributed to non-human animals (e.g., cats, dogs, and other mammals). Hence, a generic affective or feeling dimension of mind is unlikely to bear any interesting relation to the agentic dimension of morality. Third, there is the (roughly Humean, or ‘sentimentalist’) idea that empathy enters directly into moral judgment and decision-making. This is suggested, at least, by the stereotype of the ‘morally insane’ (i.e., criminal psychopaths) as Machiavellian social calculators, devoid of compassion or fellow-feeling (Hare, 1993; Blair et al., 2005). Such individuals might be regarded as less responsible for their actions on the grounds that, due to their lack of empathy (as opposed to lack of rationality), they do not genuinely understand the difference between right and wrong — or at least that they are not motivated to conform to moral norms in the way empathic individuals are.

If so, moral agency might well depend on a kind of affective mindedness, namely, emotional sensitivity to the morally relevant features of social situations.

By factoring mindedness in this way, we hoped to give the agentic analogue of Bentham's view (roughly, the idea that what matters most to moral agency is feeling, rather than thinking) some initial plausibility. Nonetheless, our working hypothesis was that moral agency is, on the commonsense conception, the mirror image of moral patiency. Hence, we expected that the mindedness manipulation would be effective in the cognitive (rationality) dimension only, with affect (empathy) paling into insignificance.

In this study, participants began by reading a brief description of a fictional character named John, who has a history of anti-social behavior. They were then asked to indicate whether John was morally responsible for his bad behavior, also using a six-point Likert scale (1 = not responsible at all, 6 = fully responsible). Participants were then presented with one of four hypothetical scenarios in which John's psychological profile was revealed in more detail, ranging from low rationality plus low empathy to high rationality plus high empathy, and asked to update their agency assessment.

The details of the study are as follows. All participants read the following passage:

John has a long history of anti-social behavior. As a child, he frequently bullied and terrorized his classmates at school. He often got into physical fights — which he would invariably win, due to his superior strength and his fierce determination to dominate others. As an adult, working in the executive division of a large multinational corporation, John maintained his reputation for ruthlessness. Quick to anger, he routinely lashed out at those working under him when he judged their performance to be inadequate. Family and friends received similarly rough treatment at his hands.

To what extent do you think John is morally responsible for his anti-social behavior?

As in the first study, answers to this question provided a baseline measure of agency judgment, for later comparison with a judgment made after updating.

Participants were randomly assigned to one of four conditions in the table below (see Table 4):

	<i>High empathy</i>	<i>Low empathy</i>
<i>High rationality</i>	Maximal	Intermediate (R>E)
<i>Low rationality</i>	Intermediate (E>R)	Minimal

Table 4. *Agency factors: Rationality \times Empathy.*

In the high rationality/low empathy condition, for example, they read:

One day, John decides to leave his job in the private sector and start a new career in the military. As part of the application process, he is given a battery of psychological tests. Most of the tests are designed to assess the capacity for reasoning, problem-solving, and decision-making. But not all of the tests are like this. Some of the tests are designed to assess the capacity to empathize with other people — to understand how others feel, and to respond sensitively to their emotional needs.

John gets high scores on the rationality tests but low scores on the empathy tests. His psychological profile indicates that he has a strong capacity for reasoning, but a limited capacity to empathize with others.

In the other three conditions, participants read variants of this passage in which John’s psychological testing yielded suitably different results. Participants in all conditions were then asked:

Given this new information about his psychological profile, to what extent do you now think John is morally responsible for his anti-social behavior?

As with the previous study, this second question was meant to reveal the effect of mindedness information on judgments of moral status.

In this case, it was hypothesized that judgments of moral agency would depend primarily on whether the target was rational, and only secondarily (if at all) on whether the target was empathic. As in the patience study, then, we

expected to see an effect of the mindedness manipulation in one dimension but not the other.

We collected responses from 84 participants, 56% of whom were female. The mean age was 31 (SD = 11.61). There was no main effect of the mindedness manipulation on agency, in either dimension (rationality or empathy), and no interaction effect. There was, however, a strong effect of the mindedness manipulation on the change of agency ratings relative to baseline, in both dimensions: rationality, $F(1, 80) = 13.08$, $p = .001$; empathy, $F(1, 80) = 11.07$, $p = .001$ (see Table 5 and Figure 5). As before, these results were confirmed by nonparametric analysis, using the Mann-Whitney U test with $\alpha = .05$.

	<i>High empathy</i>	<i>Low empathy</i>
<i>High rationality</i>	.47 (.750)	-.33 (.577)
<i>Low rationality</i>	-.38 (.865)	-.67 (.796)

Table 5. *Agency: Rationality × Empathy.*

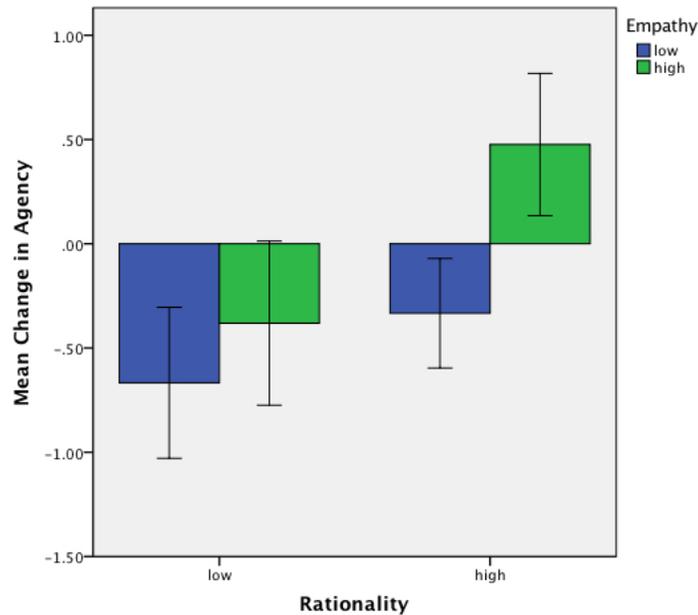


Figure 5. *Mean change in agency.*

The overall findings were at odds with our hypothesis that the cognitive dimension of mindedness (rationality) contributes to the attribution of moral agency in a way that its affective counterpart (empathy) does not. It appears

that both factors play a role in the commonsense conception of moral agency. Accordingly, the study provides indirect support for a Humean view of moral judgment, according to which moral-cognitive competence depends as much on the capacity to feel as it does on the capacity to think.

3 Ramifications

The experimental work described above suggests that commonsense thinking about the relation between mindedness and morality is more complex than one might initially suppose. In particular, it suggests that the two dimensions of mind perception (thinking and feeling) do not map onto the two dimensions of moral perception (agency and patiency) in a straightforward way. Though Bentham's view of the singular importance of sentience for moral patiency seems to capture the folk conception fairly well, the idea that high-level cognitive capacities play a comparably dominant role in the folk conception of moral agency does not appear to be correct. Instead, it seems that folk judgments of moral agency are also sensitive to perception of an individual's capacity to feel (not just think) her way through the social world, using empathy. Where that capacity is perceived as defective, the perception of moral responsibility is diminished as well.

Where does this leave us with respect to traditional debates about the metaphysics of moral agency and patiency? A first, preliminary remark is this. If we think of the metaphysics of moral agency and patiency as an essentially descriptive (as opposed to prescriptive) enterprise, and if we think of the task of descriptive metaphysics as articulating the structure of our thought about the world (Strawson, 1959), then it becomes hard to see how metaphysical debates in this area could fail to be illuminated by experimental findings like those presented above. A second thing to note is the prominent role played by thought experiments, or intuition pumps, in the literature. For example, Bernstein (1998) makes liberal use of such devices in defending an 'experientialist' view of moral patiency, according to which sentience suffices for moral consideration. A foundational assumption of Bernstein's project, like so many projects in philosophical ethics, is that the results of thought experiments constitute an

important part of the evidence base for theory construction. The same goes for philosophical investigations into the nature of moral agency (Fischer, 1999). Third, philosophical appeals to intuition are typically made on the assumption that such intuitions are pre-theoretic, or at least relatively theory-neutral. And those are precisely the sort of intuitions that experimental philosophy is in a good position to identify, insofar as it targets the judgments of individuals without philosophical training.

More generally, experimental philosophy offers powerful tools for mining the conceptual foundations of the social sciences. It does so in a somewhat surprising, even paradoxical, manner: by applying the empirical methods of the social sciences themselves. In this sense it may appear a bit like a puppy chasing its own tail. But not all circularity is vicious, and this case seems more virtuous than most.

References

- Bentham, J. (1789). *Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Bernstein, M. (1998). *On Moral Considerability: An Essay on Who Morally Matters*. Oxford: Oxford University Press.
- Blair, J., Mitchell, D., and Blair, K. (2005). *The Psychopath: Emotion and the Brain*. Malden, Mass.: Wiley-Blackwell.
- Fischer, J. M. (1999). Recent work on moral responsibility. *Ethics*, 110: 93–139.
- Gray, H., Gray, K., and Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315: 619.
- Gray, K., and Wegner, D. M. (2009). Moral typecasting: Divergent perceptions of moral agents and moral patients. *Journal of Personality and Social Psychology*, 96: 505–520.
- Gray, K., and Wegner, D. M. (2011). To escape blame, don't be a hero — Be a victim. *Journal of Experimental Social Psychology*, 47: 516–519.

Hare, R. (1993). *Without Conscience: The Disturbing World of the Psychopaths Among Us*. New York: Guilford Press.

Jack, A. I. and Robbins, P. (2006). [Intuitions about complex cognition.] Unpublished data.

Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., and Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, 63: 81–99.

Robbins, P. (2008). Consciousness and the social mind. *Cognitive Systems Research*, 9: 15–23.

Robbins, P., and Jack, A. I. (2006). The phenomenal stance. *Philosophical Studies*, 127: 59–85.

Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.

Waytz, A., Gray, K., Epley, N., and Wegner, D. M. (2010). Causes and consequences of mind perception. *Trends in Cognitive Sciences*, 14: 383–388.

Acknowledgments

Versions of this material were presented in talks at Case Western Reserve University, Truman State University, and the University of Texas at Arlington, as well as at my home institution, the University of Missouri. I am grateful to the audiences on those occasions for critical feedback. The research reported in this paper was carried out in the MU Experimental Philosophy Lab with the able assistance of Eddie Kirsch. Thanks also to Tony Jack and Josh Knobe for advice and moral support.