

Psychopaths and Moral Knowledge

Manuel Vargas

Department of Philosophy

University of San Francisco

mrvargas@usfca.edu

and

Shaun Nichols

Department of Philosophy

University of Arizona

sbn@email.arizona.edu

keywords: psychopaths, moral knowledge, moral responsibility,
exculpation, moral/conventional task

Manuel Vargas is Associate Professor of Philosophy at the University of San Francisco. His research interests include moral psychology, moral responsibility, evil, and topics in Latin American philosophy. He is author (with John Fischer, Robert Kane, and Derk Pereboom) of *Four Views on Free Will* (Blackwell, 2007).

Shaun Nichols is Professor of Philosophy at the University of Arizona. His recent research concerns experimental philosophy, cultural evolution, free will, and cognitive theories of the imagination. He is the author of *Sentimental Rules: On the Natural Foundations of Moral Judgments* (Oxford, 2004) and (with Stephen Stich) *Mindreading* (Oxford, 2003).

Psychopaths and Moral Knowledge

Manuel Vargas and Shaun Nichols

Neil Levy argues that empirical data shows that psychopaths lack the moral knowledge required for moral responsibility. His account is intriguing, and it offers a promising way to think about the significance of psychopaths for work on moral responsibility. In what follows we focus on three lines of concern connected to Levy's account: his interpretation of the data, the scope of exculpation, and the significance of biological explanations for anti-social behavior.

1. The moral/conventional distinction

In this section we'll argue that (1) Levy overclaims what the data show about the moral/conventional distinction; (2) the moral/conventional distinction does not neatly enough track any facts about moral knowledge and moralized judgments to warrant the implications he draws from the data; (3) the nature of morality matters, and cursory consideration of some metaethical issues suggests that we may be better off supposing psychopaths can be responsible, at least sometimes; and (4) even if all the preceding points (1-3) are false, it is still not clear why a grasp of purely conventional norms is not sufficient to ground responsibility, in at least some cases.

First, Levy draws on data from Blair to conclude that psychopathsⁱ “fail to grasp the [moral/conventional] distinction; for them, all transgressions are rule-dependent” (MS, 8). They “lack the ability to distinguish moral from conventional transgressions” (MS, 8). But Blair's data don't show anything this strong. In Blair's classic study (Blair 1995), what he finds is that a group of psychopathic prisoners does not show a significant difference on any of the standard measures used in the moral/conventional task. However, it's crucial to note that there were only 10 subjects in the group, and so it's possible (indeed, quite likely given the overall responses) that with a bigger sample, a distinction would emerge between the moral and conventional items for psychopaths. What Blair does find (see especially Blair 1997) is a diminished sensitivity or capacity with respect to making the distinction, comparing across children who have or lack psychopathic tendencies.

However, even if psychopaths have a diminished appreciation of moral considerations, it is a further issue whether their diminished capacity leaves them with a partial or variable capacity that suffices for responsibility in some contexts, even if not in others. We return to this issue in the next section.

It is worth bearing in mind that experiments on psychopathologies usually produce data that is less ordered than we might hope for. For example, it is not as though all autistic children fail the false belief task. Nor do psychopaths miss every case of the moral/conventional task. Rather, psychopathologies tend to show relatively diminished response. This, of course, does not undermine the importance of Blair's results for discerning the psychological mechanisms implicated in moral judgment. The fact that a defective emotion system is correlated with a defective moral system is a very interesting fact for theorizing in moral psychology. However, it is misleading to say that autistic children *lack* theory of mind or that psychopaths *cannot draw* the moral/conventional distinction.

Second, although the moral/conventional task might reveal important psychological capacities, it is important to notice the domain of the “moral” in this task is much narrower than the domain of morality as we think of it in philosophy. There are moral notions that do not map on to the distinction as it is conceived of and tested for by psychological researchers: aretaic notions and particular welfarist notions, for example. There are also cases where the overlap between the moral and the conventional is vague. What of duties to the self; what of the possibility that some moral reasons are prudential, or reflect some idealized observer's recommendations of prudence?

We take it that all of this points to a third line of concern: one's metaethics matters. Suppose one accepts that some genuinely moral reasons are prudential (a line that would be familiar to ethicists in ancient Greece). And, suppose we thought that psychopaths could govern their conduct in light of prudential norms to a degree that is comparable to at least parts of the non-psychopathic population that we ordinarily hold responsible. It seems plausible to say that psychopaths would have at least a limited kind of moral knowledge in light of their ability to apprehend reasons that are both moral and prudential. But, this conclusion is at odds with Levy's account, so we need to hear from Levy why this conclusion is not the right one to accept.

There are several routes available to Levy. He could reject this picture of moral knowledge, endorsing one that disallows prudential concerns to ever count as moral. However, we would then want to hear more about why this is a plausible view. Alternately, he might accept that prudential concerns can sometimes count as moral, but then argue that there is a difference in how psychopaths and non-psychopaths regard their moral knowledge, and that it is in *this* difference that the non-responsibility of psychopaths is to be secured. How might this latter argument go? Well, here is one way: Levy could argue that psychopaths do not (cannot?) think of moral norms (even the prudential ones) as marking out or constituting reasons as *moral* reasons, even if they are in fact sometimes perceiving and responding to prudential concerns that are also, in fact, bits of moral knowledge.

Although this route would allow Levy to draw a distinction between the responsibility of psychopathic and non-psychopathic populations, the consequences of going this route strike us as unforgiving. Consider discussion of so-called “Huck Finn” cases, where an agent does the right thing (e.g., freeing Jim from slavery) but believes that doing so is immoral and anti-social act. The ordinary reaction, it seems to us, is to (morally) praise such Huck-style agents. However, if Levy were to argue that a moral judgment can only be genuinely moral if the agent thinks of the involved knowledge or judgments as specifically moral, then it looks like Huck fails to meet this standard. So again, we would face an intuition face-off similar to the one concerning fairness: we could either give up on the idea that Huck is praiseworthy or we could give up on the idea that blameworthy agents need to think of their reasons for action in some particular circumstance as moral (Cf. Nichols and Vargas).ⁱⁱ We are inclined to favor the latter possibility, but this forces us back to the conclusion that we should suspect that psychopaths are indeed morally responsible in a range of cases. At any rate, we would like to hear more about from Levy about his picture of morality and whether prudential considerations can count, and if so, what he wants to say about prudential knowledge grounding limited responsibility in the case of psychopaths.

Fourth and finally, Levy maintains that psychopaths do understand conventional violations. We would like to know why the ability to track conventions, along with some pedestrian facts about

reactive attitudes and the utility of social norms, is not sufficient to justify holding psychopaths responsible. Levy writes that “For psychopaths, all offenses are merely conventional, and therefore – from their point of view – none of them are all that serious” (ms 9). Here's our point, though: we can and do take convention violations very seriously. Think about conventions surrounding sexual fidelity, personal privacy, and property ownership. It also emphasizes a further question Levy should consider: given that we do think it is appropriate to blame and punish transgressors of conventional rules, why isn't this enough to blame and punish psychopaths? Levy might reply that we cannot direct *moral* blame at transgressions of conventional rules. But this is not obvious, at least when we consider sexual, property, and privacy conventions. Still, suppose that Levy has an explanation for our praising and blaming attitudes in these conventional contexts. (There are several ways he might do this- perhaps the reactions are only pseudo-moral, or perhaps, moral judgments are inappropriately being applied moral cases.) We would then want to know more about how moral blame and regular blame differ, and why we can't cast moral blame on agents that knowingly flaunt conventions that we value.

To be sure, the issues we have raised turn on some complicated issues in metaethics and moral psychology. Fully replying to them is perhaps beyond the reasonable scope of a single paper. Nevertheless, they are clearly relevant to the issue as Levy has framed it. Even if psychopaths fail to apprehend those moral considerations captured by the moral/conventional distinction, it does not tell us anything about the psychopath's sensitivity to a range of at least apparently moral notions not captured by the distinction. Inasmuch as the psychopath's responsibility turns on the ability to possess moral knowledge (or, alternately, to apprehend moral reasons), partial or variable sensitivity to the range of moral considerations not captured by the moral/conventional distinction might properly ground responsibility attributions in those other domains. Partial knowledge of the full range of moral reasons might be sufficient to ground complete blame, at least in some circumstances. Indeed, this may be the normal case for all of us.

2. The scope of exculpation

Levy holds that the empirical data about psychopaths shows that they cannot control their behavior in light of moral knowledge, and that this is sufficient to show that they are not morally responsible.

In the present section we will grant the assumption that moral knowledge is required for moral responsibility. What we have argued is that the evidence is insufficient to show that psychopaths lack moral knowledge altogether (as opposed to having a decreased grasp of it). In this section, we will suggest that there is a different approach one might take to some of these issues, one that might constitute a friendly amendment to some of Levy's account.

We'll begin with a brief sketch of a view of responsibility that is somewhat different than Levy's but perhaps it can be rendered in a way compatible with Levy's own view. On the view we are suggesting the matter of anyone's responsibility depends on, at least, the relevant justified norms, the context, and our (justified) expectations about how easy or difficult it is to comply with the norm. On this view, there is no reason to suppose there is a fixed, justified standard by which the capacity for moral knowledge and reasons sensitivity is to be measured, and required, across all contexts. Instead, there are likely varied standards for appropriately determining praise and blame. The conditions that appropriately govern blaming in a society where agents are ordinarily acculturated with anti-racist attitudes may be different than the justified norms in societies with racist attitudes; these societies may yet share a further, but distinct norm with different standards of praise and blame about some other subject matter, say, serial monogamy.

Part of the reason for this is that capacities (for moral knowledge, for apprehending moral considerations, for governing conduct in light of these things) are best understood as capacities *in a context*. The ability to swim is not an ability to swim in all possible contexts, but rather, an ability to swim in certain contexts and not others. That we can swim in ordinary ocean water is no guarantee that we can swim in a vat of acid, or in a highly viscous liquid. Moreover, how we properly think about what counts as a capacity is partly a function of our aims. Sometimes we are only interested in capacities in a very general way, as when we ask if someone knows how to read. Other times, the capacity is conceived of in a very particular and narrow way. Famously, a good portion of the free will debate turns on the virtues of conceiving of the responsibility-relevant notion of the capacity

“to do otherwise” in a way that holds fixed the full facts about the arrangement of the universe, including its past and the laws of nature.

If one accepts something like this picture, it points to a different way of thinking about things than is suggested by Levy's account. For example, we may discover that psychopaths have variable capacities to possess moral knowledge, and variable capacities to respond to and control their behavior in light of it. In some contexts, psychopaths may do as well as most of us with approximately equal effort. In other contexts it may be harder for them to comply with the demands of morality, but perhaps the difficulty will not exceed what is required by the justified norms of responsibility. In other contexts we may discover that the difficulty of compliance, or that the appropriately governing notion of compliance, cannot be satisfied. In those circumstances psychopaths will not count as responsible agents.

To clarify: we are not arguing that we cannot discover that psychopaths altogether lack the capacities required for moral responsibility in all possible contexts. We are agnostic about what the data will show. Our points are merely these: (1) we see no reason to start from the supposition that such sweeping exculpation will necessarily be the order of the day, even if we are convinced that psychopaths cannot draw the moral/conventional distinction and (2) nothing about the existing data speaks against the sort of picture we have offered; indeed, if the data only suggests a diminution of capacity (as we suggested above), then the sort of account we have gestured at better captures the nuances of the existing data.

3. Tumors and the creeping threat of biological explanations

Before concluding, we wish to consider the significance of a class of thought experiments for thinking about moral responsibility.

Levy makes use of an example he borrows from Reznick, where we are asked to imagine a boy who, over time, becomes increasingly aggressive and anti-social as a result of a tumor (ms 15). The example plays several roles. First, it is a part of Levy's critique of attributionist approaches to responsibility; it is used to generate an argument for the importance of an agent's history for moral

responsibility. Second, the example might be taken to suggest a somewhat more implicit argument that we believe, and ought to believe, that behavior rooted in mental disorders is exculpatory. In this section, we will suggest that what the tumor example shows is not particularly clear, and that it points to a general problem of hasty exculpation.

(As an aside, but one relevant here and elsewhere, it seems to us that much of the dispute Levy has with “attributionists” would evaporate if we distinguish between two issues: (1) responsible agency, i.e., whether an agent is appropriately subject to the norms of responsibility and (2) the edicts of the norms of responsibility, i.e., what the justified norms of responsibility say about praising and blaming particular actions, character traits, and so on. The latter might (somewhat misleadingly) be called “blameworthiness.”ⁱⁱⁱ Presumably, all or nearly all attributionists will be committed to there being *some* account of responsible agency. It is implausible to suppose that they hold we can attribute responsibility willy-nilly to anything we like. So, when they talk about failure to see the force of morally relevant reasons as not exculpatory, presumably they are assuming that we are talking about creatures that are morally responsible agents. What is at stake is their *blameworthiness* and not their responsible agency. In contrast, Levy's concern seems to be for the latter. Still, we do not want to deny that one might have worries about particular attributionist conceptions of responsible agency, or worries that many of them have failed to sufficiently articulate some account of responsible agency.)

Consider the tumor case. Levy writes, “Now, it is clear that when his tumor is discovered at autopsy, we would cease blaming Billy for his vicious behavior” (15). Perhaps, but we might inquire about which “we” he has in mind. The folk, or theorists of different stripes? Suppose he means the former. In that case, we would like Levy to specify why we the latter, the theorists, have any special reason to care about the opinions of “we” the former. Moreover, there is some reason to be skeptical that we can reliably predict folk attitudes about ascribing responsibility. Work by Nahmias et. al. (forthcoming), and Nichols and Knobe (forthcoming) among others, strongly suggests that folk beliefs about responsibility attributions can depart in interesting ways from the expectations of theorists, and seem sensitive to a variety of factors that are not always obvious. So, it

seems to us, we should be cautious when ascribing beliefs to the folk.

Now suppose the correct interpretation of Levy's "we-who-exculpate" is theorists. Here the truth of the claim is even harder to assess. As described, the case simply does not tell us enough to make a judgment, at least not on a wide range of theories. On theories that tie responsible agency to rational agency, or to knowledge of moral reasons, we would clearly need to know more. For example, we would need to know whether the tumor damaged reasons-responsiveness or apprehension of moral knowledge, diminished these capacities, or eliminated them altogether. Bear in mind that it is implausible that the tumor causes behavior all by itself; what causal powers it has are part of a causal field that includes the brain and the nervous system, as well as the context of action in which the agent operates. If all the tumor does is exacerbate the disposition to anti-social behavior, but it does not eliminate the capacity (whatever that comes to) to override those dispositions, then we could explain away the growth of anti-social behavior without necessarily having any justification for exculpating that behavior.

The tumor might be another one of any number of things that might make compliance with the demands of morality more difficult, but not necessarily exculpating. For example, we might discover that listening to loud music encourages anti-social behavior. Such a discovery does not seem sufficient grounds for exculpation. It is difficult to see why a tumor whose morally salient causal effects are not different than loud rock music should count as exculpatory. One might be tempted to point to the tumor's biological nature or the fact that it is located under the agent's skin. We do not deny that these considerations can give the example some extra layer of creepiness and intuitive pull towards exculpation. It is difficult to see what principled reason there is, though, for considering biology or location to matter when it has no difference in the causal effects.

We take it that all of this points to a lesson that is perhaps underappreciated in the literature at large. The lesson is this: the more we learn about the biological basis of behavior, the more we may want to exculpate disorders with identifiable biological roots. Now Levy doesn't make this an error of assumption- part of the contribution of his paper is that he actually argues for it. Nevertheless, the urge to exculpate psychopaths and others with disorders that are at least partly

biologically rooted may outstrip both the empirical facts and our best normative theories. Proceed with caution.

References

- Arpaly, Nomy. 2004. *Unprincipled Virtue*. Oxford: Oxford.
- Blair, R.J.R. 1995. A Cognitive Development Approach to Morality: Investigating the Psychopath. *Cognition* 57: 1-29.
- Blair, R.J.R. 1997. Moral Reasoning In the Child With Psychopathic Tendencies. *Personality and Individual Differences* 22: 731-739.
- Fischer, J., R. Kane, D. Pereboom, and M. Vargas. 2007. *Four Views on Free Will*. Oxford: Blackwell.
- Nahmias, E. and S. Morris, T. Nadelhoffer, and J. Turner. (Forthcoming). Is Incompatibilism Intuitive? *Philosophy and Phenomenological Research*.
- Nichols, S. and M. Vargas. 2007. How to Be Fair to Psychopaths. *Philosophy, Psychiatry, and Psychology* (this volume): pp.-pp.
- Nichols, S. and J. Knobe. (Forthcoming). Moral Responsibility and Determinism: The Cognitive Science of Intuitions. *Nous*.
- Strawson, P.F. 1962. Freedom and Resentment, *Proceedings of the British Academy*. Xlvii: 1-20.
- Wallace, R. J. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard.
- Watson, G. 1987. Responsibility and the Limits of Evil. In F. Schoeman, ed. *Responsibility, Character, and the Emotions*. New York: Cambridge: 256-287.

-
- i Levy identifies psychopathy with antisocial personality disorder. But the diagnostic criteria are significantly different, and as a result, the diagnosed populations are not co-extensive, at least in the work of many contemporary psychopathy researchers. The first is, at best, a subset of the latter, although even this is unlikely to be true for there may be some anti-socials that are not psychopaths and some psychopaths that are not anti-socials.
 - ii Perhaps what this shows is that we should care about moral knowledge *de re* and not *de dicto*. If so, then psychopaths may well have *de re* moral knowledge, even if they never have *de dicto* moral knowledge. See Arpaly (2004) for this distinction as applied to moral reasons. If what matters is moral knowledge *de re*, then we will need to do some work in moral epistemology before we can settle whether psychopaths are morally responsible.
 - iii The roots of this distinction can be found in Strawson (1963), and it has been accepted in one form or another in a good portion of the literature inspired by Strawson, e.g., Watson (1987) and Wallace (1994). See also Fischer, Kane, Pereboom, and Vargas (2007), especially chapters 6 and 8.