BETWEEN THE SPECIES

Review
of
*Personhood, Ethics, and Animal Cognition:
Situating Animals in Hare's
Two-Level Utilitarianism*

Gary E. Varner
Oxford University Press 2012
336 pp., cloth

TAL SCRIVEN
California Polytechnic State University
tscriven@calpoly.edu

Volume 16, Issue 1

Jun 2013

© *Between the Species*, 2013
<http://digitalcommons.calpoly.edu/bts/>

This is an important addition to the literature on nonhuman animal welfare and rights with some potential implications for environmental ethics. As the title of his book makes clear, Varner is working within the utilitarian tradition. As such, issues relating to animal sentience, cognition and personhood will take on special significance. He goes beyond the intuitive speculation often found in the philosophical literature and brings a large body of recent empirical research to the table.

The personal setup for the book is of some interest. Utilitarian arguments have always been at the forefront of public debate and policy about animal welfare. The major player in recent applied utilitarian ethics has, of course, been Peter Singer. Singer, in turn, was a graduate student at Oxford under one of the major utilitarian theorists of the twentieth century—R. M. Hare. As time went on, Singer struggled with utilitarian theory especially as it applied to issues of personhood. Meanwhile, Hare struggled with a few issues in applied ethics including the moral status of animals, but he never got beyond a position he tagged as “demi-vegetarianism.” Varner seeks to weave these two thick threads of the utilitarian cloth a little closer together.

The book is divided into three parts. Part I develops a solid reassessment of Hare’s normative theory. Philosophers may find the treatment of Hare’s metaethics a bit on the light side but, for better or worse, interest in the meaning of moral terms has flagged since Hare wrote *The Language of Morals* in 1952. So, I suspect, not too many of the people who will be interested in Varner’s book will lament this lack of metaethical weight. There is a discussion of Hare’s basic views about the logic of moral language but it trails pretty quickly into normative territory.

He proceeds to a nice explanation of Hare's displeasure with the standard method for assessing the adequacy of normative theories. This method was formalized, as much as it probably can be, by Rawls in his *Theory of Justice*. According to Rawls, only a process of "reflective equilibrium" can reasonably assess theories. We bring our pre-theoretical intuitions to bear for and against various theories, like utilitarianism, and, somehow, see which of them best fit those intuitions. The test cases get trotted out, the counterintuitive results are duly logged in the textbooks and, at some point, the theories under analysis, like utilitarianism, are found too counterintuitive to warrant rational assent. This method is on display in almost every ethics class in almost every university. It is used all the time on ordinary moral claims that ordinary people assert all the time. Hare hated it.

Varner notes the "scathing review" Hare wrote of *A Theory of Justice* in 1973. Instead of mere intuitions, Hare thought theories ought to be checked against the very logic of moral discourse. When this is done, Hare argued, only utilitarianism emerges as a reasonable choice. Whatever is right about other theories like Kant's or relativism or any rights theory will be covered in a full analysis under utilitarianism. Moreover, none of the test cases usually thrown in the faces of utilitarians will ultimately stand up. Many of the specific arguments Hare produced can be used even if one is defending utilitarianism through the method of reflective equilibrium. This is good because most utilitarians continue to rely on this method to defend their theory. It is also good because few philosophers will be any better convinced by Varner's rendition of Hare's arguments than they ever were by Hare himself. It is also good that not much of what follows in Varner's book really relies on Hare's position that utilitarianism is the only normative view consistent with the logic of moral discourse.

Less revolutionary, and more generally agreeable, is Hare's "two-level utilitarianism." Variants of the two-level approach are as old as utilitarianism itself. There is a level—Hare called it the "critical level"—at which utilitarianism requires that each act be chosen so as to maximize utility. However, a world in which everyone always *attempted*—act by act—to maximize utility would almost certainly not result in maximal utility. The world is not full of ideally rational and perfectly informed moral agents. The world is full of unexpected consequences that only a perfectly informed and rational "archangel" (to use Hare's term) could foresee. What makes utilitarian sense in the real world of imperfections and surprises is the construction, adoption and advocacy of rules, laws, habits and customs that will, as an "intuitive-level system," most likely maximize utility.

As Varner admits, the general maneuver at play here was used by Bentham, Mill, Sidgwick and just about every defender of utilitarianism since. It is the base for most of the standard moves utilitarians make to get around the standard objections to the theory. Suppose you could save three people with failing organs by killing an innocent person and harvesting his organs to save the first three. Wouldn't utilitarianism require that? So goes the standard test case. And here goes the standard response: Only if you are prepared to ignore all of the disutility generated by a standardization of this practice along with the routine training of medical professionals to think the practice is morally required. Without the assumption of such ignorance the practice could never be even imagined as felicitous. So it would go with a long list of other counterexamples. Is it OK to frame an innocent person if doing so would stop a crime wave? Is it OK to drown your Aunt Bea in a bathtub in order to secure her inheritance and then do good with it? Is it OK to publicly

torture prisoners for the amusement of the sadistic masses? In each case the utilitarian is likely to respond by pointing out that justifying any of these things would entail setting up a world full of motives, incentives, habits and virtues that would be—from a utilitarian point of view—perverse.

Early on, Varner brings the two-level move to an old division among those concerned with the treatment of animals. Since the publication of Tom Regan's *The Case for Animal Rights* in 1983, there has existed a split between animal welfarists (typified by Singer's position in his 1975 book *Animal Liberation*) and animal rightists like Regan. At the theoretical level this is supposed to be grounded in an irreconcilable split between utilitarians and those, like Regan, who reject utilitarianism in favor of a belief that there are fundamental rights that trump all utilitarian calculations. (Regan is responsible for the Aunt Bea example mentioned in the previous paragraph.) Indeed utilitarians have not always been sagacious in their talk about rights. Bentham famously claimed that talk of rights was nonsense and talk of natural rights was "nonsense on stilts." But then, utilitarians since Mill have gotten used to apologizing for Bentham's occasional misfirings. And, since Mill, utilitarians have usually thought it important to develop a theory of rights. For most, rights are to be grounded in exactly the sorts of customs and laws that exist in the intuitive-level system. So, a utilitarian can say, with a reasonably straight face, that animals have certain rights—and that Aunt Bea has a right to not be killed by her calculating nephew.

I'm not sure any anti-utilitarian is going to be swayed by the arguments Varner presents in Part I. He doesn't seem to expect any such success. But even an anti-utilitarian will recognize that Varner has done a decent job of presenting the view and

defending it as at least one way to approach issues related to the treatment of nonhuman animals by humans.

In Part II, Varner takes up the matter of how morally considerable nonhumans can be sorted into three categories that must be of interest to a utilitarian: persons, near-persons and the merely sentient. One might wonder why a utilitarian should be concerned about these groupings at all, let alone why any utilitarian *must* be concerned about them. There is some history here. Another of Bentham's famous claims is this: quantity of pleasure being equal, "pushpin is as good as poetry." In other words, pleasure is pleasure. It doesn't matter that philosophers have almost universally held that some kinds of pleasures are higher and better than others. In the end, for a utilitarian, the pleasure of a fool is just as good as the pleasure of a bookworm which is, in turn, no better than the pleasure of a pig. The value of a life is to be cashed out entirely in terms of *how much* pleasure it yields, not *what kind* of pleasure it produces.

Mill tried to answer Bentham: "It is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied." The response is troublesome. How can Mill, or any utilitarian, hold a straightforward hedonistic theory of value on the one hand and then deny it by saying that some pleasures are *qualitatively* superior to others? Updating the utilitarian theory of value to preference utilitarianism won't help. Aren't certain higher, distinctively human, preferences more valuable than the lower ones we share with nonhumans? As troubling as Mill's answer is, most utilitarians have thought there is something very right about it. Even animal welfarists like Singer have to admit that, almost always, it would be better to kill a pig rather than a human. One can attempt a Benthamite's exit by arguing that our intuition in favor of the human

is justified on utilitarian grounds because the life of a typical human will produce *more* pleasure than that of a typical pig. But what if the human in question is, as Mill hypothesized, not all that satisfied or satisfying, and the pig is as happy as a pig eating shit?

Maybe there is no coherent way out of this for the utilitarian other than to bite the Benthamite bullet and hold that the only thing that makes the typical human more valuable than the typical farm animal is that the human's life is likely to produce *more* utility. But Varner tries to do something else. What he tries has, at least, the virtue of tying him to a lot of work done by others working in the field (notably, Singer, Jeff McMahan, Michael Tooley, Dale Jamieson and Regan). First, he wants to hold that there is, indeed, some real value attributable to personhood that is not easily and simply reducible to hedonic calculation. Its value emerges only within the context of the intuitive-level system. Second, he finds this value in the same general vicinity it is found by those noted parenthetically above: personhood emerges as a consequence of autonoetic (or self-perceiving) consciousness. This gets him close to, for example, Regan's idea that a being has inherent value if it is the "experiencing subject of a life." Specifically, Varner defines a person as "an individual who deserves special treatment or respect because he, she or it has a biographical sense of self."

He puts some flesh on this idea by drawing on Marya Schechtman's *The Constitution of Selves*. Under this view full personhood can only emerge in individuals who can locate themselves within a story of their lives. The special, and moral, significance of the sort of stories persons use to locate themselves in the world of other beings and other selves explains why there is something special about Socrates that is lacking in

the pig. Socrates lives in a story. The pig can only have a story if a person like Socrates tells it. Even then, the story of the pig is not something the pig is really aware of. It cannot enrich the existence of the pig with a complex narrative that ties together the past, the present and the future. It cannot bring meaning to the whole of the pig's life and to the daily burdens of the beast. As such, we cannot have respect for the life of the pig in the same way we can for the person. The person's life has a trajectory that rises from awareness of the past and projects itself as plans for the future. Interference with these plans damages the person and wanton disregard for these plans shows disrespect of the person.

Varner doesn't mention this but this autonoetic tack ties him in not just with people currently working in the realm of animal ethics but with the utilitarian tradition going back to Mill. Anyone who has been confused by Mill's claim that some pleasures are qualitatively better than others has probably been able to guess that Mill thinks the confusion is somehow settled by an appeal to his views about individuality. As these views emerge in the enigmatically romantic third chapter of *On Liberty* Mill stresses the importance of struggling to develop one's own tastes, opinions and character. It is a small step to think that he is talking, at least in part, about the development of one's own story. There are other connections Varner does mention to Aristotle, Nietzsche and MacIntyre. In sum, there is nothing particularly odd about this approach to personhood. It has a pretty rich and respectable story of its own among philosophers.

One consequence that Varner draws from the autonoetic approach is that linguistic understanding appears to be essentially related to personhood. After all, there can be no story without the words needed to tell it. As he puts it, "competence in a

natural language is a necessary (but obviously not sufficient) condition on personhood.” A whiff of Descartes is perhaps in the air as Varner concludes that no nonhuman can be a person because none of them appear to be able to master in any compelling way enough of a natural language to communicate or understand a story. But, of course, Varner is not going to argue, as Descartes did, that nonhumans are, therefore, non-conscious and, therefore, just machines. Also, he is not going to rest his case on a simple assertion that no nonhumans are capable of mastering a natural language. Varner makes this claim on the basis of research.

This same body of research indicates that nonhumans approach personhood to varying degrees. Some are close enough to be called near-persons. Others fall into the set of the merely sentient. Varner does a good job of presenting the empirical research now available to a community of thinkers (i.e., ethicists) who may know little of it. His descriptions of the relevant studies are clear, as are his arguments about why this research really matters for the issues at hand. Early in Part II Varner presents some of the research that bears on the issue of mere sentience. The results are nicely summarized in three tables and pretty closely match the intuitions of philosophers who have been writing in the area. Vertebrates pretty clearly seem to be sentient under any objective and empirical standard. The only slight surprise is that there appears to be less question about the sentience of fish, snakes and birds than some might have imagined. Insects and earthworms don’t do so well. Octopi and squid occupy a grey area.

When it comes to determining near-personhood, things are a little less intuitive. The idea that anything like autonoetic consciousness can be empirically measured across species is

bound to strike some philosophers as *prima facie* nutty. But there is research that purports to do just that and it is at least worth a look from anyone serious about the matter. Varner is aware of the philosophical bog he is wading into and spends a lot of time giving good arguments that the studies he reports really are measuring what they say they are, viz., the likelihood of self-consciousness. He describes three kinds of experimental examinations on the issue: tests of episodic memory, tests of self-recognition and tests for the presence of a “theory of mind.” The third is less spooky than it sounds. The presence of a “theory of mind” is indicated when a scheming, planning and manipulative individual shows awareness that other individuals seem to be scheming, planning and manipulating things as well. In the end, the group of near-persons seems to contain some of the usual suspects: primates, dolphins and maybe elephants. But there is also a surprise: corvids (i.e., ravens, jays and the like) and maybe parrots.

Part III ties the first two Parts of the book together by formulating intuitive-level system rules in light of the analyses of personhood, near-personhood and mere sentience. He begins by arguing that there are good reasons for any society to have an intuitive-level rule against the unnecessary killing of sentient individuals. Such a rule, of course, throws a lot of weight on the word “necessary” and, ultimately, such a rule will not give us much of a practical tool. It will, however, remind us that all sentient beings ought to count from a moral point of view.

The question of whether or not we are justified in killing animals for food brings Varner to an issue that Singer struggled with for about twenty years and never successfully resolved. On the face of it, utilitarianism would seem to justify the killing of animals for food. As long as cattle, for example, are treated

humanely, painlessly slaughtered and then replaced with more cattle who will be treated similarly what is morally wrong with the practice from a utilitarian point of view? Aren't all animals "replaceable" in this sense? In fact, aren't all persons similarly replaceable? Varner does a nice job of pulling apart the philosophical maneuvers Singer attempted between the writing of *Animal Liberation* in 1975 and that of *Practical Ethics* in 1993. In the end he argues that Singer needs to concede to the objections piled up by Parfit, Hart and Hare. A consistent utilitarian must treat all individuals, persons or not, as ultimately replaceable. In the end, at the critical level, long-term total utility must decide all issues. If this entails weighing the utility associated with actual beings against that of merely potential (and thus, currently non-existent) "beings" then so be it. Two corvids in the bush in the future may, indeed, be better than one in the hand at the present.

An annoying consequence of this concession arises immediately. If the interests of future beings must be weighed in against those of current beings then do current beings have an obligation to sacrifice their own best interests in order to bring those future beings into existence and, thusly (we suppose), increase the sum total of utility over the long term? Here the Harefied Singer faces Parfit's "repugnant conclusion." Do we have an obligation to drive human population up to, say, 40 billion even though the average per capita utility would be cut in half relative to the world we live in now with its population of around 7 billion? The obvious (and, for a lot of people, obviously wrong) answer is "yes."

There are lots of ways to avoid the repugnant conclusion but Varner uses this as an opportunity to display the merits of two-level utilitarian analysis. While the 40-billion-person world

may be required at the critical level, there are lots of good reasons, at the intuitive level, to leave reproductive decisions to the individuals who would have to do the reproducing and tend (either as parents or taxpayers) to the offspring that result. This solution is not trouble-free. Neither are the rejoinders that might expose the trouble. Nor are the counter-rejoinders that would follow.

Similar reasoning is on display in the concluding sections of the book when Varner moves to the issue of sustainable agriculture. From the two-level point of view, questions about how many people ought to be raising how many animals is a matter of how we can best engineer our intuitive-level laws, regulations and expectations in a way so as to approach the critical-level ideal. Dashing immediately to the conundrums presented by the ideal is fun but not entirely useful. The real work of ethics in the real world is a matter of sweating through the details in an orderly and, ultimately, utilitarian way.

At the beginning and at the end of the book, Varner realizes that philosophers who hate utilitarianism are not going to be happy with this approach or the sequel promised by Varner. Nor are philosophers inclined to larger views like Taylor's biocentrism or the ecocentrism of Leopold and Callicott. They are all likely to wonder why people like Varner are so obsessed with the details of small things like pain or personhood. Vertebrates, after all, make up only about three percent of the animals on Earth. Throw in the plants, fungi, bacteria and the rest and the percent of living things that meet the minimal qualifications for sentience or personhood shrinks even more. Why are a few components in the package of survival mechanisms that are vital to so few creatures of such overwhelmingly moral significance in the broadest scheme of life on this planet? The

promise of a sequel centered on the blatantly anthropocentric concept of “sustainability” is not encouraging.

That outburst out of the way, I think there is something terribly right about the Harean point of view that Varner so effectively applies in this book. I’m not sure that it is *logic* that forces serious thought in a utilitarian direction but something does. Early on (p. 86) Varner mentions the work of the psychologist Jonathan Haidt and Haidt’s argument that some deeply embedded intuitions drive us this way. This isn’t very Harean but it may be very right.

Whatever compels us toward utilitarianism, the theory is extraordinarily durable. Counterexamples are easily dreamt up but hard to sustain. Even when the theory seems to say the wrong thing it locates a serious problem in need of further analysis. Varner does a great job of bringing this advantage to the fore and promises more of the same regarding the mushy notion of sustainability. If the book currently under review is any indication, maybe he can pull it off. If not he has, at least, left us with this excellent and original contribution to the debate about animal welfare—or, at least, vertebrate welfare.