PUBLIC AFFAIRS QUARTERLY
Volume 8, Number 1, January 1994

VEGETARIANISM, CAUSATION AND ETHICAL THEORY

Russ Shafer-Landau

Philosophers writing about vegetarianism have often shifted subtly from an evaluation of the practices of current factory farming, to the existence of a requirement to refrain from eating and purchasing meat. These are separable issues.[1] Even if the best moral arguments favored the abolition of the factory farm,[2] and perhaps all farms where animals are raised only to be prematurely killed, this leaves us short of a moral obligation to remain or become vegetarians. I shall argue that consequentialists cannot defend the existence of an obligation to refrain from meat consumption. Deontologists may perhaps do better, but familiar deontological defenses of vegetarianism fail, and alternative accounts face severe theoretical obstacles. I conclude the paper by briefly exploring the possibilities of constructing a defense of vegetarianism along virtue-ethical lines.

I. CONSEQUENTIALISM

IF consequentialism is correct, then the move from the wrongness of factory farming to a moral obligation to refrain from meat eating is properly resistible. It is so because of the negligible causal role that each consumer plays in the generation of the harms that make factory farming condemnable.

The argument is very simple: one cannot, in one's purchase and eating of meat, have any direct influence on the amount of cruelty and harm inflicted on the animals in a factory farm. Whether one purchases a steak, or several steaks, for personal or family consumption will have no influence whatever on the amount of cruelty perpetrated on today's farms. One's meat-purchasing habits essentially make no difference at all to the total amount of suffering experienced by the billions of animals currently maltreated. The ordinary consumer of meat is so remote in the causal nexus of animal suffering, that one cannot properly attribute to any such consumer any causal, hence moral, responsibility for the admittedly wretched fates suffered by farm animals. One is morally free to do as one likes so long as one does no harm. Meat purchases do no harm. Therefore one is morally free to make them.

Call this the *Inefficacy Argument*. It has a stronger and a weaker ver-

sion. The stronger version says that ordinary meat purchases[3] make no causal contribution at all to the harms suffered by farm animals. The weaker version claims that ordinary purchases have some deleterious impact, but one so small as to be overridden by other factors. Let's consider the stronger version first.

A. The Strong Version of the Inefficacy Argument

Underlying this argument is a specific notion of how to determine the causal importance of a given action. Call it the counterfactual test of causation. This test says that some event is causally efficacious only if, had it not occurred, an outcome would not have occurred just as it did. The argument can be understood as claiming that all relevant outcomes (i.e., the amount of suffering undergone by farm animals) would have remained the same regardless of one's ordinary meat purchases. Therefore one's purchases are not causally efficacious in bringing about the harms suffered by farm animals.

There are two general ways to criticize the argument. One could deny that meat purchases fail the counterfactual test of causation. Or one could challenge the applicability of the counterfactual test itself. Let's consider this last move first.

The criticism would be that counter-factual tests of causation do not work in cases of overdetermination, cases where, had I not done what I did, someone else would have done it and brought about the same consequences. If A doesn't shoot B, C will. A shoots B. A is not relieved of responsibility, even though, had he refrained, B would have sustained the same harm from C. If meat consumption can be assimilated to such cases, then failure to pass the counterfactual test will not relieve omnivores of moral responsibility for their purchases.

But the assimilation is questionable. In cases of blameworthy causal overdetermination, each agent's actions alone are sufficient to bring about a tangible harm. If meat purchases were relevantly like this, then each purchase must be sufficient to bring about a tangible harm to a farm animal. But that is just what is in question. Unless there are independent reasons to question the legitimacy of the counterfactual test, this criticism is inconclusive.

Alternately, one could retain the counterfactual test of causation but argue that ordinary meat purchases pass the test. One reason for thinking so is the familiar law of supply and demand: a decrease in demand should generate a decrease in supply, so that, were one to avoid meat purchases and so reduce demand for meat, the supply ought also to diminish. But this "law" seems capable of bending.

Supply is sensitive to demand, of course, but there are degrees of sensi-

tivity. Your purchases may be so insignificant, relative to the total number of purchases, that the means employed by managers of meat departments to detect demand takes no notice of them. The likelihood of your purchases fading to insignificance is increased if one shops at many different places, or if one purchases meat very infrequently. Given actual methods of assessing demand, a conversion to vegetarianism might well go unnoticed by grocers and so not be reflected in their purchases from farmers. Their failure to notice your opting out of the meat market represents a failure of your purchases to have a causal impact on farm animals. If, after your abstention, meat managers purchase no less animal carcasses, then your abstention will not alter the number of animals killed or the way in which they are treated before they are killed. Your purchases, or lack of them, are therefore relevantly causally inefficacious.[4]

It is true that were consumer demand to vanish, so too would most animal cruelty on farms. The strong version may seem to miss the mark, because it appears to overlook the consumer's absolutely crucial role in spurring farmers to ever more "efficient" means of confining and killing their animals. But while true that curtailing general consumer demand would likely yield this welcome result, it does not follow that any individual's so refraining would have any effect at all. It is this latter sort of consequence that is relevant for determining where the obligations of individual agents lie.

There are two related problems for the strong version. First, if an ordinary consumer's purchases contribute absolutely nothing to the harms suffered by farm animals, then the aggregate harms caused by all such consumers' purchases should be zero. Suppose there are fifty million ordinary meat consumers, each of whose harmful causal contribution is zero. Zero added to itself fifty million times is still zero. Surely this conclusion—that all ordinary meat purchases taken together have no negative impact on farm animals—is absurd. Rejecting the absurd implication requires altering the presuppositions that got us there. The only thing to do seems to be to attribute some minimal causal influence to each ordinary consumer's purchases.

A second ground for suspecting the strong version stems from examples in the utilitarian literature.[5] Suppose a child is drowning and six people go to its aid. Only three were necessary for the rescue. Thus each person is such that if he or she were to have refrained from contributing, the outcome would have remained the same. On the counterfactual test, this indicates that each person made no causally relevant contribution. The aggregate of six nonexistent contributions is no contribution at all. Yet there was substantial assistance rendered. Something in the strong version has gone wrong.

B. Identifying the Error

Three diagnoses of the error spring immediately to mind, but are themselves erroneous. The first involves what Jonathan Glover has called a "context illusion."[6] In these sorts of cases, the harm I do to a single individual or small set of individuals is very great, though this represents a minute fraction of the total relevant harm being suffered. Though a Nazi prison guard kills just a single inmate, he has nevertheless done great wrong, even though this represents a minute fraction of the total harm rendered by others like him. One might claim that meat purchasing is relevantly like this situation. Though my purchases may seem a drop in the bucket, they nevertheless inflict the worst sort of harm on a small number of sentient beings, and so are condemnable.

If anyone is suffering from a context illusion, it is the farmer who claims that he is not doing anything wrong, since his slaughter represents only a tiny portion of that committed each year. The farmer is doing something wrong, or so I will allow; but his wrong is relevantly different from the ordinary consumer's. The purchase of a steak is not analogous to the direct killing of an inmate or to the deliberate failure to easily rescue. These latter actions or omissions lead directly to great harms. Individual meat purchases do not.

A second diagnosis focuses on causal features of aggregations of actions. Omnivores do not simply make one or two negligible purchases during their lives. They exhibit a purchasing pattern that must be considered as a whole. Over the course of a lifetime, an omnivore may spend thousands of dollars on meat purchases. One's repeated purchases do add up, and so do affect the overall amount of animal suffering.

This line of argument is clearly problematic because it incorporates a fallacy similar to one that Samuel Clarke committed in his formulation of the cosmological argument three centuries ago. Clarke treated as a distinct entity the set of all past, present and future events, and asked of this set what might have caused it. He found only God. His error lay in thinking of this agglomeration as a distinct entity and investigating it as if it had a peculiar causal structure. A similar error underlies the argument that treats a temporally very extended series of purchases as a single thing and attributes causal powers to the agglomeration.

The critic is claiming of the entire series of an individual's meat purchases that it, as an entity unto itself, is wrong. I doubt that a plausible metaphysics will allow this sort of agglomeration. Even if it does, such an entity still fails to be causally efficacious. Were the entire series of a consumer's ordinary purchases removed from the world's history, the amount of animal suffering produced on factory farms would not have diminished. To think otherwise is to assume that one can add up the costs of each

purchase and then ask of the final sum whether it would have any causal efficacy. The fact that one has spent ten thousand dollars over a lifetime does not entail that one has exercised during one's lifetime the causal influence attainable by a one-time ten thousand dollar purchase. Just the opposite is true.

Rule consequentialists offer a third attempt to identify the error of the strong version. But theirs is more a capitulation than a diagnosis of error. This family of theories tells us that moral obligations are fixed by rules general conformity to which would maximize good consequences (or minimize bad). This allows an action to be obligatory even though, in isolated cases, it is does not yield optimal consequences. The rule consequentialist can allow that ordinary purchases have no deleterious effects, while continuing to condemn such purchases. Given a choice between two rules regulating meat purchases—forbidding them, or retaining the status quo—one might argue that the former option is optimific, and so dictates to all a moral requirement to refrain from meat purchases.

There are two things to be said here. First, this is not an argument against the strong version's causal thesis, since rule consequentialists need not show that isolated purchases do have some measurable causal impact. Rather, they claim that the causal issue is irrelevant—an act can be wrong even if causally inefficacious. This yields the counterintuitive conclusion that one can be morally responsible for certain outcomes even if not causally responsible for them.

Second, the ordinary rationale for moving away from act consequentialism is absent in this context. Rule consequentialisms sometimes require the performance of nonoptimific acts, and forbid the performance of optimific acts. This is usually justified by citing ignorance factors and the costs of deliberation for each fresh situation. But these problems don't arise for those convinced of the negligible or nonexistent causal contribution of an ordinary meat purchase. If omnivores know in advance that each meat purchase is causally inefficacious, then there is no good reason to adopt a rule forbidding them from doing so, since on *every* such occasion refraining from the desired food will be nonoptimific. Any rule consequentialism generating *that* outcome is surely implausible.[7] It's better to stick with a conception of consequentialism that makes the moral evaluation of actions dependent on the outcomes they actually produce, rather than on the hypothetical ones generated by hypothetical conformity to certain rules. If that is so, we are no closer to an analysis of the error of the strong version.

C. The Weaker Version of the Inefficacy Argument

If the stronger version of the inefficacy argument is wrong, it is so because of its incorporation of the counterfactual test of causation. Identi-

fying the error of the strong version thus requires replacing the counterfactual test with another measure of causal efficacy. The likeliest accounts come from theories of collective action. However, any such plausible theory will still leave the consequentialist short of justifying vegetarian obligations.

The argument would run as follows. First we aggregate all the harm done by all ordinary omnivores. Even if each omnivore's contribution is negligible, that done by the millions of ordinary consumers is, when aggregated, very substantial. We assign causal responsibility to each individual consumer by identifying a method of apportioning responsibility in a derivative fashion, first citing the harm caused by the aggregate, and then distributing it to its members.

The plausibility of this argument depends on a defense of two controversial assumptions. First, one must defend the claim that individuals can be derivatively responsible for harmful consequences brought about by the action or inaction of some group of which the individual is a member. Many think that individual responsibility is logically prior to collective responsibility, and so reject the method of first inquiring into the responsibility of a group and subsequently apportioning individual responsibility. Second, one must be able to assign derivative responsibility to individuals who are members of loosely structured groups,[8] i.e., unorganized groups that lack effective decision mechanisms. On one view we cannot. This view has it that a group assumes responsibility for its actions only if it can act, and it can act only if it has a collective will or intention, and it can have such an intention or will only if it has some deliberative mechanism that aggregates individual intentions. So there can be no group responsibility assigned to loosely structured groups. For now, however, let us grant that this argument can be met. It then becomes our task to identify a mechanism for apportioning derivative responsibility to ordinary omnivores in order to sustain the challenge to the stronger version of the inefficacy argument.

Probably the best way to allocate individual responsibility would be to divide the total sum of money spent for everyone's purchases by that amount spent on one's own. In most cases, the responsibility assigned will be extraordinarily miniscule. For instance, in 1989, individual U.S. consumers spent 122.7 billion dollars on meat "products."[9] Even if a person spent fifty dollars a week on meat, $2500 a year, her portion of responsibility for farm animal cruelty would be .00000002. In other words, our imaginary consumer would have contributed two dollars for every hundred million dollars spent on meat.

Still, though we have this tidy figure, how does it translate into an assignment of responsibility? Shall we say that, for every hundred million increments of harm caused to farm animals, this consumer is responsible for precisely two such increments? That doesn't sound very plausible to

me, but suppose I am wrong. Even so, the appropriate responsibility "dividend" should be even less, since we have thus far failed to factor into the calculus the contributory responsibility of the farmers who directly perpetrate the cruelty. How we might combine these two indices is a nice question. No matter how it is done, if doable at all, it will greatly diminish the already infinitesimally small degree of responsibility each individual consumer bears for farm cruelty. If farmers are at least equally responsible for farm cruelty, then (dividing individual consumer responsibility by half in the previous example) our consumer would now be responsible for one unit of harm for every hundred million such units generated.

Even if there were this measurable reduction in animal suffering attributable to one's abstaining, this would be insufficient to show that such forbearance is obligatory on consequentialist grounds. One would have to measure the sufferings averted against the gustatory satisfactions and economic benefits gained. While possible that abstaining will win out, the details of this sort of moral calculation are generally not well understood, and resolution may be a matter of appeal to intuition about the comparative moral weighting. Given the diversity of the relevant intuitions, such appeal would effectively provide no resolution at all. In any event, I can think of no other instance in which harms attributable to a loosely structured group are such that its members are obligated to exclude themselves from the group if their contribution is as slim as 1/100,000,000.

On a consequentialist view, our responsibility as moral agents is to perform actions that maximize utility or minimize disutility. It is hard to see how abstaining from meat can maximize positive utility. The relevant question is whether it will minimize disutility. If my claims about the lack of non-negligible causal effects are correct, then abstention from meat purchases does not minimize disutility. So even if we can attribute responsibility to loosely structured groups, and even if we can, in a principled fashion, divide it up amongst its members, we have yet to identify a secure consequentialist justification for imposing a moral prohibition on the purchase of slaughtered animals.

Thus far I have been speaking only of meat purchases. Meat eating is a different matter. Even if an individual's meat purchasing did have some imperceptibly small causal relation to farm cruelty, meat eating need not. Conversely, even if meat purchases have no deleterious consequences, meat eating may. As is the case for any consequentialist calculation, it all depends on the circumstances. The president of PETA may do no wrong if clandestinely purchasing a steak, but he may do great damage to his cause if seen publicly indulging in a burger with members of the Beef Board. Alternately, even if individual purchases have some deleterious consequences, isolated meat consumption may not. If your host has prepared an elaborate dinner, with Chateaubriand as the main course, it may do more

harm than good to announce your abstention. It all depends on the circumstances of the particular case. Since meat *eating* does not directly cause animals to suffer (you have not caused the cow whose carcass you are eating to suffer), it can be wrong, on consequentialist grounds, only when it has deleterious third-party effects that override one's enjoyment in partaking in it. This may be the case only infrequently, especially if one eats alone.

II. Deontology

The sort of arcane calculations and fairly generous argumentative concessions made in the last section might be enough to force a recognition that consequentialism and vegetarian obligations are unstable philosophical partners. Those who hope for a stronger, more reliable justification for the obligation to refrain from carnivorism may turn to deontology. This may be the best move for those sharing this hope. But the best may not be good enough.

A. Rights

Things might look more promising for a deontologist, since acts or refrainings can be obligatory on such theories even though they do not lead directly to the minimization of harm. But one of the traditional worries surrounding the plausibility of deontological theories immediately surfaces. The criticism is that in severing the foundational connection between obligation and harm and benefit, the deontologist leaves a justificatory vacuum that cannot be filled. In other words, if we do not essentially tie moral imperatives to the production of benefit or prevention of harm, what are they tied to?

A familiar response locates the ground of obligation in an agent's possession of rights. For instance, stealing your unwanted rug is wrong, not because doing so will diminish your welfare, but because I fail to respect your claims on me and so violate your rights. Perhaps animals have rights. And perhaps an ordinary consumer violates those rights when she makes her purchases. If so, then we have a straightforward route to justifying an obligation to become a vegetarian.

I will assume, for purposes of argument, that farm animals have rights. Such rights can be grounded ultimately in one of two sorts of consideration. On one view, they exist in order to protect certain interests of the right-holder. On another, they exist in order to confer a sphere of protection that enables the right-holder to choose to be or do as she pleases.[10] Call the former the interest view; the latter, the autonomy view. On the former view, rights violations presuppose setbacks to interests. On the latter view, rights violations presuppose infringements of autonomy.

If one thinks that the permanently comatose or the pre-toddler has rights, then one has reason for doubting the autonomy view. Certainly such a view is less attractive than the interests view for the animal rights advocate. It is far from clear that farm animals have the requisite cognitive skills to enable us to speak of their making important choices that structure an autonomous life. Still, let us grant that some such attribution is plausible and can ground a right. It does not follow that a rights view fares any better than consequentialism in generating the desired obligation to refrain from meat consumption.

On a rights view, one is violating an animal's rights only if one is setting back its important interests or importantly restricting its ability to make free choices. I don't believe that most individual restaurant patrons and supermarket customers do this. They are too removed in the causal nexus to have such an effect. If animals have rights at all, it is the factory farmer and slaughterer who violate them. I might put the point crudely. An animal is either dead or alive. If dead, it hasn't any rights, so one's purchase does not violate its rights.[11] If alive, the animal may be mistreated, but by the farmer, not the ordinary consumer. So a purchase does not violate the rights of the animal purchased. Nor does it seem to violate the rights of an animal not yet killed; the causal connection linking a purchase today with the death of an animal on a farm tomorrow is either nonexistent (the strong version) or extremely tenuous (the weak version). If in all other contexts of assessing blame, we exclude from culpability those whose contributions to the regretted outcome were as small as 1/100,000,000, then consistency requires us to do the same here. *De minimis non curat lex.* This is as effective a maxim for applying the moral law as it is for the positive law. Thus ordinary purchases of animal remains do not violate animal rights. Granting rights to animals makes it easy to justify the prohibition on their killing. But it will not assist the vegetarian in justifying the obligation to avoid ordinary meat consumption.

B. Universalization

I think a rights-based view shares with consequentialism an attachment to what I shall call the *causal requirement for wrongness.* This requirement stipulates that an action is wrong only if it causes infringements of the rights or interests of a moral subject. If the causal arguments in section I were correct, then they explain why consequentialism and rights-based ethics cannot justify an obligation to refrain from meat consumption. Since I will assume those arguments were correct, the obvious alternative for the ethical vegetarian is to abandon the causal requirement. Some deontological views do just that.

Perhaps the most prominent such view is a Kantian one, in which the rightness and wrongness of action depends on the universalizability of its governing maxim. In Kantian and neo-Kantian theories, one is forbidden from doing actions whose governing principles one cannot universalize. Since one concerned to end animal suffering would not want everyone to purchase and eat meat, but instead to refrain, a universalization requirement would appear to morally prohibit such consumption. Continuing to purchase meat would be to make an exception of oneself, to hope all others abstain while allowing oneself special privileges morally denied to others. On a generalization view, one's actions could be wrong even if casually impotent. The structure of such a justification looks promising as a means of getting the vegetarian what she wants.

However, there is a perennial problem with this sort of program, and one that seems to me insuperable. Maxims may be of greater or lesser specificity, and there appears to be no principled account of how to precisely restrict the sorts of considerations deemed relevant for maxim-construction. This raises problems, for one of the relevant circumstances attending an agent's decision to consume meat may be the failure of enough others to cooperate in reducing the amount of animal suffering. It is not clear why one would be prohibited from introducing the actions of others into one's considerations when formulating a maxim for action. If one is allowed to take such things into account, then there seems no bar to endorsing a maxim that allows meat consumption in those cases where one's consumption would have no detrimental impact on farm animals. And this is just to reintroduce the causal requirement for wrongness. Those who intend to do good and prevent harm may insist that they be allowed, in the formulation of their maxims, to incorporate relevant facts about the potential cooperation of others in the pursuit of their aims.

In a world where everyone would in fact abide by the principles one endorses, it seems we should opt for a maxim requiring us to refrain from meat purchases. If everyone were to do that, then such universal action would have the most dramatic impact on reducing animal suffering. But in our non-ideal world, it is not clear that we are bound by such principles, at least in cases involving collective action problems. It seems we may incorporate into our maxims a rider that makes our endorsement of courses of action dependent on the cooperation of enough others to achieve our goal. When one's goals cannot be attained through individual effort alone, there is special justification for including such a rider. Without it, morality would often require virtual heroism or utter self-abnegation.[12] If we assume with most nonconsequentialists that morality is not so demanding (they use the stringency of consequentialist demands as an important critical weapon[13]), then the relevant maxim would tell us, not to refrain from meat

purchases *simpliciter*, but to so refrain only if enough others were acting in concert so as to effectively reduce the amount of animal suffering perpetrated on farms. If a scheme of cooperative action is not yet in effect, one would not be bound to refrain. Short of an account defending restrictions on information about the actions of others in maxim construction, a universalization test will not generate the desired vegetarian obligation.

C. An Attractive Deontological Principle

Deontological considerations are not exhausted by rights-based considerations and universalization tests. In what follows, I shall offer what I think is the strongest deontological principle defending the existence of an obligation to refrain from purchasing and eating animal remains. Given what has gone before, the principle must show how we can be morally required to refrain from certain actions, even if those actions do not directly infringe the autonomy or set back the important interests of animals. It must be strong enough to overcome the presumption (in law and morality) that we are free to do what we please so long as we do no (nonnegligible) harm.

A first, natural way to express the operative deontological principle would be the following: one must refrain from supporting or otherwise contributing to cruel practices. Since factory farming is cruel, we are therefore obligated to refrain from supporting it, and so must avoid meat purchases. However, this is too quick. Against this, one might note that our government supports cruel regimes overseas, and our university psychology departments and medical schools often allow cruel treatment of lab animals. Application of this principle would thus require us to abstain from paying taxes or tuitions. These consequences are so extreme as to call the principle that generated them into question.

We should envision the amended principle as follows: one must refuse (even symbolic) support of essentially cruel practices, if a comparably costly alternative that is not tied to essentially cruel practices is readily available. Factory farming is essentially cruel because the benefits reaped by farmers are directly dependent on the cruelty they impose on their animals. The factory farm industry, with its current profit margins, can benefit farmers only by ensuring the most unnatural and uncomfortable conditions for its inhabitants. Alternatives to meat eating are readily available and often end up costing less than factory farmed food. This principle would also allow us to pay taxes and tuitions, so long as the relevant institutions only incidentally, and not by nature, support cruelty. Thus given the cruelty on today's farms, and the easy availability of vegetarian alternatives, there would be a general obligation to abstain from even symbolically supporting meat "producers."

Endorsing this principle allows us to defend satisfying responses to the following cases. There does, for instance, seem something morally repugnant about a willingness to utilize or purchase soap made from the bodies of concentration camp victims. This is so even though one's present refusal to use the soap could not contribute to any reduction of cruelty.[14] Likewise, voicing one's support for a racist dictator overseas is condemnable, even if one's expression has no causal impact whatever on the treatment of the oppressed. Or suppose one is already well-attired but receives a fur coat as a gift. There seems something objectionable about wearing it, even if doing so in no way contributes to the suffering of animals on "fur farms." By wearing the coat, one exhibits a certain callousness concerning the suffering of the animals whose lives were taken to produce it. The same line of reasoning would support condemnation of ordinary meat consumption. That an omnivore's individual actions are not responsible for the death of any particular animal may nonetheless leave her open to strong criticism.

The abovementioned principle offers the best hope for those concerned to defend the existence of an obligation to refrain from animal consumption. But it faces two familiar problems that beset any principle specifying deontological constraints. The first is to identify the sorts of considerations that can ground such a principle. Familiar aproaches to deontological justification—appeals to self-evidence, rationality, human nature, divine commands—are each subject to exceptionally strong criticism and in need of substantial defense and articulation. The second challenge is how to deal with a paradox recently noted by Scheffler and Nagel.[15] Deontological constraints make it impermissible to engage in certain activities, even when doing so is necessary to prevent an even greater number of the same sort of objectionable activities. Since the prohibition on certain behaviors presupposes their undesirability, it is hard to see how such constraints could be justified, since respecting them would sometimes allow an even greater amount of similarly undesirable actions to take place.

III. Virute Ethics

> Were the generous friend or disinterested patriot to stand alone in the practice of beneficence, this would rather enhance his value in our eyes, and join the praise of rarity and novelty to his other more exalted merits.
>
> —Hume, Appendix III
> *Enquiry concerning the Principles of Morals*

As with many deontological constraints, the one described above seems quite deeply plausible. And yet, with many philosophers, I am uneasy about the justification of such constraints. I cannot talk here about whether the two problems I just enumerated are solvable. But I would say one word in conclusion about how ethical vegetarianism would fare if the deon-

tological program were to run aground. At such a point, one may investigate the possibilities of a virtue-based approach to ethics. There are three ways one could go.

One might *abandon* deontic concepts altogether.[16] Wrongness, obligation and impermissibility would no longer be appropriate terms of ethical assessment. Instead, one could focus solely on the exemplification of certain traits of character, and make action assessments only in terms of virtue attributions. For instance, one might talk, in our previous examples, of an agent's callousness without needing to introduce considerations of whether her behavior was morally right or wrong.

An interesting alternative would allow for the retention of deontic concepts, so long as they played a *coordinate* role to the virtues. One wouldn't abandon moral concepts, but supplement them with a second "track" along which aretaic assessments could be made. One way to do this involves retention of the causal requirement for wrongness. In that case, one could claim that ordinary omnivores do no moral wrong, since their meat consumption is relevantly causally inefficacious. Nevertheless, they may be condemnable to the extent that they display an indifference to the cruelty that went into the "production" of their "goods." If such a case could be made out, it would have quite interesting implications for an ethics of virtue that seeks to make the attribution of virtue or vice independent of determinations of moral rightness or wrongness. This independence comes at a cost: one must allow for the possibility of virtuous actions done in the service of immorality, or moral actions undermining virtue. But this might be a cost acceptable to those who can't bear to do without moral and aretaic concepts, and who believe that the moral or aretaic domain would be trivialized by deriving its edicts wholessale from the other.

A final alternative is to allow attributions of moral rightness or wrongness, but make them *derivative* of virtue attributions. An act would be wrong to the extent that it exemplifies a vicious character, and right if demonstrative of a virtuous one. This would allow one to abandon the causal requirement for wrongness. Meat purchasing and eating would be wrong, even if causally inefficacious, because they demonstrate a disregard for the suffering experienced by the animals whose remains one is wearing or eating.

A virtue ethics run along any of these three lines would allow us to make sense of Hume's wonderful remark, quoted at the beginning of this section. We feel justified in praising the person who stands for worthy ideals, even if their execution is stymied by the absence of cooperation by one's fellow citizens. Cynics might say that vegetarians cannot "stand alone in the practice of beneficence," since they confer no benefits by their abstention. However, even if true, the reply is only partly effective. Any plausible virtue ethic would be equipped to make both positive and negative aretaic

assessments. Even if vegetarians are not virtuous for their abstinence, omnivores may be less virtuous for their indulgence. Seeking and deriving satisfaction from "products" that are known to result from cruel practices diminishes one's admirability. This is so even if the practical impact of one's indulgence is nonexistent or negligible.

Of course, my suggestions here are only programmatic. Full defense awaits further refinements in virtue ethics. In any event, one important avenue for either deontological or aretaic research programmes is a direct assault on the causal requirement for wrongness. This is of crucial importance for the ethical vegetarian, if my causal arguments in section I were correct. Failure to undermine this requirement generates the paradoxical result that, although the practices reigning on farms today are condemnable, ordinary purchases made from such farms are not. I do not believe that consequentialists can avoid this paradox. Whether their theoretical opponents fare any better remains to be seen.[17]

University of Kansas

Received June 25, 1993

NOTES

1. For prominent arguments that seem to slide from the wrongness of contemporary farming practices to an obligation to avoid meat consumption, see (inter alia) P. Singer, *Animal Liberation*, 2d ed., (New York: Avon, 1990); *Practical Ethics* (Oxford: Oxford University Press, 1979), chs. 3 & 5; "All Animals are Equal," in P. Singer, ed., *Applied Ethics* (Oxford: Oxford University Press, 1986), pp. 214-228; T. Regan, "Ethical Vegetarianism and Commercial Animal Farming," in R. Wasserstrom, ed., *Today's Moral Problems* (New York: MacMillan, 1985), pp. 459-78, and "The Moral Basis of Vegetarianism," *Canadian Journal of Philosophy*, vol. 5 (1975), pp. 205-13.

2. I shall take as background for the following discussions the conditions that obtain on today's "factory farms." I am convinced that these are most always quite bad places for animals to be, not (only) because they are there to be killed, but because of the suffering they must undergo before slaughter. I do not want to take a stand on whether the painless killing of an animal that has been well treated, and perhaps allowed to live out its natural lifespan, is morally permissible. Instead, I shall simply assume the going factory farm conditions as the context for the arguments to come. For detailed information about the state of such farms, see Jim Mason and Peter Singer, *Animal Factories* (New York: Harmony Books, 1990).

3. The remainder of this essay will focus exclusively on *ordinary* meat purchases. I allow that there will be unusual circumstances where one's meat purchases are capable of having a dramatic impact on the treatment of animals. Generally, this will occur if one's purchases represent a sizeable portion of the market. The CEO of McDonald's clearly has a substantial impact on the lives of farm animals. Of course, the consequentialist isn't in the clear even in these cases, because the interests of such

people (and their employees) in retaining current practices must be balanced against those of the animals whose remains they process and serve.

One might argue that the workers have the option of finding employment elsewhere, while the farm animals have no such choice. However, vocational alternatives will not in fact exist for all currently employed in meat-related businesses. The closure of every meat-related industry would result in the loss of thousands of jobs at all income levels. In addition, the move undertaken by those able to alter their vocations would not be costless. Reductions in pay and retraining costs must be balanced against the (admittedly huge) costs borne by farm animals. The calculus may very well tilt in favor of the animals. But not obviously.

To know whether this is so, we need a procedure for performing interspecies comparisons of welfare. One might hope to satisfy this need by importing a well-accepted mechanism for making inter*personal* comparisons of well-being, but there is no such mechanism. Though we feel intuitively certain that A suffers less harm in having her nose tweaked than B does in having her leg amputated, we have no widely endorsed method for saying why this is so, much less one that gives us assistance in the hard cases (where we really need it). See generally J. Elster, ed., *Interpersonal Comparisons of Well-Being* (Oxford: Oxford University Press, 1992).

4. Even if methods of assessing demand become so sophisticated as to measure the impact of a small meat purchase, it does not follow that abstaining from the purchase would have any causal impact. Meat managers may notice your abstention and be unable to do anything about it, i.e., may be unable to reduce the amount of animals they purchase. Butchers buy animal remains on the basis of units that are much larger than those for sale at the consumer level. A butcher may have to purchase a side of beef to satisfy customer demand, even though some of the beef goes unsold. That is, the purchase from which you abstain may represent an amount of meat less than the smallest purchasing unit a butcher can make. When that is so, your abstention, even if noticed, cannot make a difference to the amount of animal remains a butcher purchases, and so cannot make a difference to the number of animals killed or the way in which they are treated before slaughter.

5. See J. Glover, "It makes no difference whether or not I do it," *Proceedings of the Aristotelian Society*, Supplemental Volume XLIX, pp. 171-90; D. Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), pp.75-82; B. Gruzalski, "The Case against Raising and Killing Animals for Food," in P. Singer and T. Regan, *Animal Rights and Human Obligations* 2ed., (Englewood Cliffs, NJ: Prentice Hall, 1989).

6. J. Glover, "It Makes No Difference Whether or not I Do It," *op. cit.*

7. It may be that acting on a rule that prohibits meat purchases leads to best consequences. But even the best action-guiding rule may counsel the performance of actions that turn out to be wrong on consequentialist grounds. Thus there is a distinction between those standards that guide action and those that determine the rightness of an action. I am arguing against rule utilitarianism as a theory of the standard of rightness. For more along these lines, see Peter Railton, "Alienation, Consequentialism and the Demands of Morality," *Philosophy and Public Affairs*, vol. 13 (1984).

8. This term comes from Larry May, "Collective Inaction and Shared Responsibility," *Nous*, vol. 24 (1990).

9. *Statistical Abstract of the United States*, (Washington, D.C.: U.S. Government Printing Office, 1990).

10. This is obviously an oversimplification, but one that will not harm the ensuing arguments. I agree with Carl Wellman when he writes that "there is no need to choose between a theory in which having a right consists in some special status of the will of the right-holder and a theory according to which rights protect the interests of the right-holder, because we have learned that rights protect some interests of their possessors *by means of* giving dominion to the right-holder in the face of some second party whose will might be opposed to the will of the possessor of the right." [Carl Wellman, *A Theory of Rights* (Totowa, NJ: Rowman and Littlefield, 1985), p. 217.]

11. If this argument works, then shouldn't it be permissible to purchase a human corpse with the intent to eat it? No. Such purchase may well be impermissible, but its impermissibility is not explained by conferring rights to the corpse. A dead body, human or nonhuman, has no rights; its mistreatment may violate rights, but not those of the corpse. If anyone has a legitimate grievance at such behavior, it is some third party who fears for her own postmortem treatment or who is offended by what she witnesses.

12. For instance, if one weren't allowed to consider the likely actions of others, then a generalization test might require one to undertake a very dangerous rescue even though others could perform the same task without effort, to keep promises to others who are known to be cheats, or to be pacific among those known to be aggressors. It is far from clear that morality can be so strenuous in its demands on innocent agents.

13. Cf. Bernard Williams' contribution to *Utilitarianism: For and Against* (New York: Cambridge University Press, 1973).

14. There is a similar debate in the field of research ethics about the propriety of using data collected from cruel experiments. However, the two scenarios are importantly disanalogous, since in the research case there is often no low-cost, equally effective, readily available alternative to utilization of the data, whereas in the soap case one could easily purchase a cruelty-free product. I think we can construct a strong argument for using the data from cruel experiments. But because of the difference in the two cases, this would not allow us to infer the permissibility of using or purchasing the soap.

15. See Samuel Scheffler, *The Rejection of Consequentialism* (New York: Oxford University Press, 1982), and "Agent-Centred Restrictions, Rationality and the Virtues," *Mind*, vol. 94 (1985), pp. 409-19; Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), ch. 9. See also D. Parfit, "Is Common-Sense Morality Self-Defeating?," *Journal of Philosophy*, vol. 76 (1979), pp. 533-45.

16. Cf. Elizabeth Anscombe, "Modern Moral Philosophy," *Philosophy*, vol. 33 (1958), and Michael Slote, *From Morality to Virtue* (New York: Oxford University Press, 1992), for suggestions on how and why this might be done.

17. Many thanks to Tom Regan, Dave Schmidtz, Don Marquis, Richard DeGeorge, Peter Singer, Doran Smolkin and Lawry Finsen for helpful comments on earlier drafts of this paper. None were satisfied with my conclusions, and all forced me to rethink several of the arguments that got me there.