

Expected Utility, Contributory Causation, and Vegetarianism

GAVERICK MATHENY

ABSTRACT *Several authors have argued that act-utilitarianism cannot provide an adequate critique of buying meat because a single meat purchase will not actually cause more farm animals to be raised or slaughtered. Thus, regardless of whether or not the production of meat is inhumane to animals, someone who buys meat is doing nothing wrong. This argument fails to show that meat purchases are morally permissible, however, because it assumes that act-utilitarians would use actual utility in their decision to buy or not to buy meat. I show that act-utilitarians cannot use actual utility as a decision procedure and must instead use expected utility to prescribe or proscribe actions. I then demonstrate how expected utility can be applied to cases of contributory causation, where many people seem morally responsible for causing something to happen. Buying meat is one case of contributory causation where the probability of any single individual's affecting meat production is slight, but the expected disutility of affecting that production is substantial. Thus, in its expected utility form, act-utilitarianism defeats the 'causal inefficacy' defence of buying meat [1].*

Several authors have argued that act-utilitarianism cannot provide an adequate critique of buying meat because a single meat purchase does not cause more farm animals to be raised or slaughtered [2]. This argument is based on the plausible assumption that most meat purchases are causally inefficacious — a single meat purchase is too insignificant, relative to the vast number of other meat purchases, to be noticed by the manager of a factory farm. If the manager cannot perceive any increase in demand caused by a single meat purchase, no additional animals will be raised or slaughtered, and thus no harm will have been done to animals by the purchase. If this is so, then the act-utilitarian cannot condemn meat purchases on animal welfare grounds.

This causal inefficacy argument would, indeed, refute a common utilitarian argument for vegetarianism — if we accepted the premise that an action can be prescribed or prohibited only if it will *actually* cause something good or bad to happen. However, this argument can be rebutted by a version of act-utilitarianism which has an effective decision procedure, since any such version must prescribe or prohibit actions according to their *expected* utility, not their *actual* utility. I will contrast expected and actual utility, and then apply expected utility to two cases of contributory causation — *The 100 Bandits* and *The 200 Million Consumers*. I conclude that a version of act-utilitarianism employing expected utility defeats the 'causal inefficacy' defence of buying meat [3].

When we make a decision about how to act, we never know for certain the *actual* utilities that will result from all our possible actions. We may, after making a decision to act in a particular way, come to know the actual utilities that resulted from the one action we decided upon. However, this knowledge is not helpful in making the original

decision, since it is not only reached after the fact, but also limited to only one of the many possible actions we may have had to choose from. Consequently, act-utilitarians have typically argued that we should make decisions, not on the basis of actual utility, but on the basis of expected utility — the product of the utility resulting from an action and the probability of that utility resulting — that one might reasonably predict given the available evidence [4]. Since expected utility, not actual utility, can be known when making a decision, only expected utility can help an act-utilitarian to decide what course of action to take.

I have suggested that an act-utilitarian must use expected utility, rather than actual utility, when making a decision. I will now apply expected utility to problems of contributory causation, where many people seem responsible for causing something to happen. Jonathan Glover provides an example of contributory causation called *The 100 Bandits* [5]. 100 bandits descend on a village and find 100 villagers, each villager with one bowl, each bowl containing 100 baked beans. Each bandit takes one bean from each bowl, so that each bandit ends up with a bowl of 100 beans. Now, no villager can perceive the difference made by one bean being stolen from his bowl (either at the moment or later, due to malnutrition). Thus none of the bandits would seem to have individually harmed any of the villagers and so none of the villagers should have been harmed. Yet 100 villagers are without lunch and hungry. So something went wrong here.

Glover suggests we approach contributory problems like *The 100 Bandits* by employing a ‘divisibility principle’ — a single agent is causally responsible for the total utility of a contributory result divided by the number of contributing agents. In this case, the hunger of 100 lunch-less villagers is divided over 100 bandits. Glover would thus say that each bandit is responsible for the hunger of one lunch-less villager. If we accept Glover’s divisibility principle, each bandit ought not to steal 100 beans because he would then be causally responsible for the disutility of one lunch-less villager.

There may be a more compelling solution to contributory problems such as this one, however, that does not attempt to reconcile actual causal responsibility with our intuitions about moral responsibility. For in the case of the Bandits, it is not true that none of the bandits is *actually* causally responsible for harming the villagers. At the very least a handful of the bandits are causally responsible for the villagers’ hunger — those bandits who complete *threshold units*. While it is true that no villager can perceive the difference made by one bean stolen from their bowl, each can clearly perceive the difference made by 100 beans stolen from their bowl. Thus there must be some number of beans between one and 100 that is the smallest number of beans a villager can perceive. Call this number the threshold unit. Say, for instance, the threshold unit is 20. Any number of beans stolen below 20 cannot be perceived. Any number of beans stolen between 20 and 39 is perceived only as 20 beans being stolen; between 40 and 59, only as 40 beans being stolen; and so on, up to 100 beans. Thus bandits who cause a 20th bean to be stolen are responsible for the disutility of 20 beans being stolen. For instance, bandits who cause the 100th bean to be stolen from a bowl are responsible for the disutility of 20 beans being stolen, since had they not caused the 100th bean to be stolen, only 80 beans would have been perceived as stolen.

This, I think, is the approach to take in describing the causal responsibility, after the fact, of agents in similar problems of contributory causation [6]. However, as I suggested above, this retrospective description of actual utility does not help us to *prescribe* action.

For this, act-utilitarians must use expected utility. Imagine that the bandits are contemplating stealing beans again. This time, each bandit knows villagers can perceive only threshold units of 20, but each bandit does not know whether he will be stealing a 20th bean from each bowl. Under this uncertainty, each bandit ought to calculate the expected disutility of stealing 100 beans as the probability of completing a threshold unit in each bowl ($1/20$) times the disutility of perceiving that threshold unit (20) times the number of bowls (100), which equals the disutility (100) of one hungry villager. Even if each bandit knows neither the size of the threshold unit nor which bean he is stealing, he can still calculate the expected disutility. In each case he will know that the disutility of a threshold unit times the probability of completing a unit in each bowl is one. (This is so because the size of the threshold unit and the probability of completing it always vary inversely.) Hence the expected disutility of stealing 100 beans will always be 100. The only condition under which the expected disutility will be less than 100 is when the Bandit has information about both the exact size of the threshold unit and the exact position of a particular bean within that unit. In most cases of contributory causation, this kind of information will not be available.

As a decision procedure, expected utility yields the same prescription as Glover's divisibility principle, don't steal beans. We should have anticipated this, since the sum of all the bandits' chances of completing a perceptible unit is one and the product of each of these probabilities is also one. One virtue of the expected utility view, then, is that it provides the same prescriptions as Glover's divisibility principle but without a questionable view of *actual* causal responsibility.

If we accept this expected utility interpretation of utilitarianism and its applications to problems of contributory causation, then the 'causal inefficacy' defence of buying meat no longer holds. There must be *some* threshold at which point a unit of meat demanded by some group of customers is perceived by the grocer. At the very most, the size of this threshold unit is the difference between the demand for no meat and the current demand for meat. Likewise, there must be some threshold at which point a unit of meat demanded by some group of grocers is perceived by the butcher. And so on, all the way to the farmer. The expected disutility of my completing a threshold unit that affects the production and slaughter of animals is thus the product of all the probabilities of completing each threshold unit [$p(\text{All}) = p(\text{Grocer}) * p(\text{Butcher}) * \dots * p(\text{Farmer})$] times the disutility of that entire threshold unit of animal production. I suspect that, as others have suggested, this probability is quite small. However, the disutility of completing the threshold unit is the disutility of the entire unit, not some portion of it. This disutility is quite large, since it involves raising and slaughtering a significant number of animals [7].

For example, take the case of *The 200 Million Consumers*. There are 200 million consumers, each of whom eats 50 farm animals each year. In this market there are only ten possible annual outputs of animals for farmers: one billion animals, two billion, and so on, up to ten billion. The difference between each of these annual outputs, one billion, is the smallest unit of demand perceivable to the farmer and is thus the threshold unit. Since there are 20 million customers per threshold unit, and only one of these customers will actually complete the unit of which his or her purchase is a part, the probability of my completing a unit is one in 20 million. That means by buying meat I have a one-in-20 million chance of affecting the production and slaughter of one billion animals. The expected disutility is then one-20-millionth times one billion, which equals 50 — that

is, the disutility associated with raising and slaughtering 50 animals per year [8]. In my view, that's substantial. And these are close to the actual numbers for meat production and consumption in the United States.

As with *The 100 Bandits*, with *The 200 Million Consumers* only a small fraction of individuals may actually *cause* harm, as determined after the fact. While at first glance this seems to weaken the act-utilitarian argument against buying meat, on closer inspection it makes no difference. Only a decision procedure can tell us whether or not we ought to buy meat. The act-utilitarian must thus use expected utility, not actual utility, to defend buying meat. It is important to note that the uncertainties in this case militate against buying meat, not for it. Although we don't know what the actual thresholds are, this does not change our expected disutility, since the probability of our completing a threshold unit times the disutility of that unit is always 50 (again, because the size of the threshold unit and the probability of completing it always vary inversely). It is only with a great deal of certainty (more than is probably possible) about thresholds and the location of our purchases within them, that we could claim eating meat is morally permissible on the grounds of 'causal inefficacy'.

I have argued that act-utilitarianism uses expected utility, not actual utility, as a decision procedure in prescribing moral actions. Expected utility, moreover, allows us to make wise moral decisions in cases of contributory causation. Buying meat is one case of contributory causation where the probability of affecting meat production is very slight. Nevertheless, in its expected utility form, act-utilitarianism defeats the 'causal inefficacy' defence of buying meat.

*Gaverick Matheny, School of Business, Duke University, Durham NC 27701, USA.
gaverick.matheny@duke.edu*

NOTES

- [1] I am grateful to Bart Gruzalski, Peter Singer, James Rachels, Nathan Nobis, Zephyr Teachout, and an anonymous referee for their helpful comments on an earlier draft.
- [2] TOM REGAN (1983) *The Case for Animal Rights* (Berkeley, University of California Press), 224–6; RUSS SHAFER-LANDAU (1994) Vegetarianism, causation and ethical theory, *Public Affairs Quarterly* 8, 85–100; R. G. FREY (1983) *Rights, Killing, and Suffering* (Oxford, Basil Blackwell), 207–15; MICHAEL MARTIN (1976) A critique of moral vegetarianism, *Reason Papers* 3, 13–43; HUD HUDSON (1993) Collective responsibility and moral vegetarianism, *Journal of Social Philosophy* 24, 89–104; NATHAN NOBIS (Forthcoming) Vegetarianism and Virtue, Does consequentialism demand too little? *Social Theory and Practice: An International and Interdisciplinary Journal of Social Philosophy*.
- [3] The 'causal inefficacy' argument suggests that even if farm animals are raised inhumanely, a utilitarian could still defend buying meat. If the inefficacy argument is incorrect, however, the utilitarian has one less defence for buying meat, as the suffering involved in the production and slaughter of farm animals is considerable. See PETER SINGER (2001) *Animal Liberation*, 1st Ecco ed. (Hopewell, NJ Ecco Press); JIM MASON and PETER SINGER (1990) *Animal Factories*, 2nd ed. (New York, Harmony Books); and GAIL EISNITZ (1997) *Slaughterhouse* (Amherst, NY, Prometheus Books).
- [4] The expected utility interpretation is consistent with the classical formulations of utilitarianism by Bentham and Mill. It is also suggested by R. B. BRANDT (1996) *Ethical Theory* (Englewood Cliffs, N.J., Prentice Hall); and more explicitly by JOHN HARSANYI (1979) Bayesian decision theory, rule utilitarianism, and Arrow's impossibility theorem, *Theory-and-Decision* 11, 298–317; and BART GRUZALSKI (1981) Foreseeable Consequence Utilitarianism, *Australasian Journal of Philosophy* 59, 163–176.
- [5] JONATHAN GLOVER (1975) It makes no difference whether or not I do it, *Proceedings of the Aristotelian Society*, Supp. Vol. XLIX, 171–90. For more problems of contributory causation, see DAVID LYONS

(1965) *Forms and Limits of Utilitarianism* (Oxford, Clarendon Press); DONALD REGAN (1980) *Utilitarianism and Co-operation* (Oxford, Clarendon Press); and DEREK PARFIT (1984) *Reasons and Persons* (Oxford, Clarendon Press), 76–80. Despite the clear advantages of expected utility in prescribing action, Parfit uses actual utility to explore imperceptible harms and confuses matters in the process. For responses to Parfit, see EDWARD GRACEY (1989) Comment on Shrader-Frechette's 'Parfit and mistakes in moral mathematics', *Ethics* 100, 157–159; KRISTIN SHRADER-FRECHETTE (1987) Parfit and mistakes in moral mathematics, *Ethics* 97, 50–60; and BART GRUZALSKI (1986) Parfit's impact on utilitarianism, *Ethics* 96, 760–783.

- [6] For instance, in Parfit's cases *The Harmless Torturers* and *The Drops of Water*, where there must be some thresholds at which point aggregates of imperceptible effects become perceptible harms; and where there must be some number of torturers and altruists who complete threshold units.
- [7] This point has also been made by PETER SINGER (1980) Utilitarianism and vegetarianism, *Philosophy & Public Affairs* 9:4; BART GRUZALSKI (1989) The case against raising and killing animals for food in Tom Regan and Peter Singer (eds.) *Animal Rights and Human Obligations* (Englewood Cliffs, N.J., Prentice-Hall); and MICHAEL ALMEIDA and MARK BERNSTEIN (2000) Opportunistic carnivorousism, *Journal of Applied Philosophy* 17, 205–211.
- [8] As with the case of *The 100 Bandits*, Glover's divisibility principle yields the same prescription as expected utility — we divide the total disutility of factory farming per year, 10 billion animals, over the number of customers, 200 million, again giving us 50 animals per customer per year.