

# **How a Utilitarian Should Think About Factory Farming: A Defense of the Utilitarian Case for Vegetarianism**

**Kevin Shin  
Macalester College**

*Submitted as a senior capstone  
in partial fulfillment of a Philosophy Major*

Advisors  
Geoffrey Gorham  
Samuel Asarnow

## ***Introduction***

Modern society is marked by the emergence and domination of large corporations and systems in which the average individual participates but holds no discernible and effective power. A moral problem arises when we consider the damage and suffering that these structures create. For instance, the factory farming industry is responsible for the suffering and torture of billions of animals every year and contributes to a significant percentage of carbon emissions which harms the environment. However, our participation in this system presents a complicated situation when assessing our *individual* moral position as participants of the industry—although we might all agree that factory farming in and of itself is objectionable, it's unclear whether it necessarily follows that our participation and usage of the resources it provides is objectionable as well in light of our lack of direct causal power to effect it.

This moral predicament is *prima facie* perhaps most embarrassing for the consequentialist, who broadly judges the moral correctness of an action as a function of its effect on the world. For example, act consequentialism is the principle that an action is morally right if and only if it maximizes good, and wrong if and only if it does not. On one view, act consequentialism seems unable to prohibit purchasing factory farmed meat if an individual's actions do not minimize or maximize anything in a morally relevant way. In fact, on alternative interpretations of act-consequentialism, one could even argue that the right action is to divert one's energy towards more effective causes and continue to purchase factory-farmed products. The *inefficacy objection to consequentialism* precisely exploits this unintuitive viewpoint, and states that consequentialism is untenable because it recommends courses of action contradictory to our basic moral principles. This objection has been widely discussed as demonstrating the falsity of consequentialism, and recent literature has benefited from the numerous examples

provided by consumerism as a direct example of a realistic causal situation for which consequentialism provides unacceptable or unclear answers. In short, the existence and intractability of these large-scale collective action problems seems to be evidence that consequentialism is false.

In response to these doctrines, several writers have defended consequentialism and offered arguments to show why and how consequentialists can respond to these problems. Most notably, Singer and Kagan point to the responsiveness of the economy to consumer demand as a fact which contradicts inefficacy; the observation here is that with enough consumer input, the market must respond, and so the calculation can be done strictly as an expected utility calculation assessing the individual's ability to contribute to and influence the group's total effect. In response, others have replied by highlighting other aspects of the economy, stating that it seems unlikely that this characterization of the economy describes the relevant causal effects, or lack thereof, of individual purchases.

Broadly, it is my view that current economic research on this problem is inconclusive, and that it is rational for the consumer to be unsure. My reason for thinking this is simply that authors and economists across the field offer different accounts and theories regarding the significance of an individual's actions, and I am skeptical that a single definition or characterization of the economy will always hold—consumers often find themselves in varying degrees of purchasing power. With this uncertainty, consequentialism actually initially appears to give no answer, or certainly not the morally objectionable one attributed to it. This means that the inefficacy objection is inconclusive at best but gives us important reasons to understand how the consequentialist really would respond.

In response to this observation, I propose that the consequentialist ought to pursue a virtue-based approach when thinking about the factory farm problem. Parsing out uncertain economic or statistical evidence is only one possible approach to the consequentialist picture. I argue that even if a consequentialist is not absolutely certain about the direct economic effects of a given action, they still have good reason to refrain from purchasing meat by appealing to a set of virtues which prohibit this action. Specifically, the virtues that she should consider are those which are more likely to lead to less animal suffering if followed correctly. To this end, I offer three virtues which I claim the consequentialist should accept because adopting them plausibly increases the likelihood of performing actions which produce good consequences and show that these support vegetarianism. Thus, this paper concludes with a defense for the utilitarian case for vegetarianism.

## **PART I. AN ANALYSIS OF THE PROBLEM**

### ***Why factory farming is a collective action problem***

Collective action problems are instances where the solution to a social problem requires the participation of a certain percentage of the public acting together, but each person has no rational reason to act individually because of the insignificance of their own actions. Without collective and coordinated effort by acting alone, the individual is wasting their time and resources which could perhaps be diverted to more useful or productive activities. For example, suppose that cleaning up a beach requires the active participation of 60% of the nearby community and trash accumulates in such a way that any less participation will not clean things up.<sup>1</sup> If we observe that no one is cleaning up the beach, or even that only 50% of people are

---

<sup>1</sup> I am indebted to Samuel Asarnow for this helpful characterization.

cleaning up the beach, it is rational for us not to clean up the beach by ourselves; it is simply not a productive solution to the problem or a usage of our time. The standard characterization of factory farming is precisely in this light. The problem is too large or too complex to plausibly state that a single individual's purchase at a given supplier will yield significant effects that decrease the suffering of the animals involved or detriment the factory farming industry, and so it is rational for us to purchase. For example, my decision not to purchase beef at Whole Foods will likely not carry any weight or affect the revenue of Whole Foods.

### ***The current factory farming problem and our relationship to meat***

Factory farming is defined by Merriam-Webster as “a farm on which large numbers of livestock are raised indoors in conditions intended to maximize production at minimal cost.” Factory farms contain somewhere close to 99% of U.S. farmed animals (Sentience Institute) which amounts to around 9 billion land animals a year (“Factory Farms”). Industry standards are to treat farm animals with a strict profit-seeking lens, and this usually translates to animal mutilation justified by a cost-benefit analysis. For example, a common practice is to “debeak” chickens by searing off their beaks to minimize the cannibalism and violence which occurs when chickens are put under high-stress, cramped environments (Capps).

Aside from the obvious suffering inflicted by these practices on the animals, factory farming is notorious for its adverse effects to the environment and human health. Factory farms produce unmanageable amounts of manure, polluting nearby water reservoirs with antibiotics and disease, and livestock produce methane gas which is estimated to contribute almost 37% of total methane emissions (Good). Furthermore, a 2015 review of studies on the health consequences of eating meat found a link between the consumption of red meat and increased

risk of cardiovascular disease and type 2 diabetes and recommended that red meat should be restricted more than current standards (Richi et al.).

Conversely, vegetarianism has well-documented health benefits, and a significant proportion of vegetarians are often motivated by their own health concerns or hopes that switching to a plant-based diet will enhance athletic performance. Vegetarian diets have been shown to reduce the risk of heart disease, certain types of cancer, and diabetes, and promote weight regulation (Harvard Health Publishing). A body of research has also shown that vegetarian diets have the potential to increase performance of endurance athletes (Nieman), and some of the most elite athletes in the world practice veganism or vegetarianism specifically for this reason.<sup>2</sup> In addition, a significant proportion of the population has access to high-quality alternatives, which provide comparable levels of pleasure and adequate micro- and macronutrients. Vegetarian diets have been shown to provide more than adequate amounts of protein (Marsh et al.) and many restaurants almost always have alternatives which are often less expensive.

The philosophical importance of these statistics is to demonstrate both the severity of the problem and the fact that as an industry, the amount of pain and torture imparted on the animals is cruel, intentional, and unnecessary. Few of us have ever visited factory farms and perceived the conditions of the animals, and it's worth briefly noting that our moral sensitivities on this issue are likely to be sharply different if information about these practices was broadcast more faithfully than presented to us as nicely packaged products at grocery stores. In another sense, establishing the viability and accessibility of alternative non-meat options is crucial when assessing the role meat plays in our diet. These health benefits alone may provide adequate

---

<sup>2</sup> Venus Williams, Colin Kaepernick, and Novak Djokovic are some of the most well-known.

utility for an individual to reduce their meat consumption, and furthermore meat cannot be claimed to be a nutritional requirement which provides an irreplaceable utility of good health. In other words, abstaining from meat does not imply a sacrifice in personal well-being—in fact, given the documented benefits of vegetarianism, the opposite seems true.

### ***Two Classic Act Consequentialist Viewpoints: Singer and Kagan***

Recognizing the amount of suffering that the industry causes, act consequentialist arguments against eating meat builds off the fundamental belief that a utility calculation, when comparing our pleasure from eating meat to the lifelong suffering inflicted on that animal, makes the claim that we ought not to consume it. For example, Peter Singer's *Animal Liberation* begins with the premise that we ought to extend our consideration and utilitarian calculus to include non-human beings, simply for the reason that they are capable of feeling pain (Singer, 2). Combined with the premise that our gustatory pleasure is outweighed by their suffering, Singer's claim is simply that one ought to be a vegetarian through a straightforward welfare-calculation.

From one perspective, these observations seem unobjectionable.<sup>3</sup> Additionally, the appeal of this argument is captured by a general consensus that factory farming harms animals in an unnecessarily cruel way. However, the claim against this view is that between my action of purchasing the meat and the suffering of the animal from which it came is the economy, which complexifies and distorts the causal chain between my purchase and the specific death and suffering of that animal such that it there is no longer a meaningful relationship. As one

---

<sup>3</sup> For the purposes of my argument, I assume that this is an uncontroversial opinion, although a number of neuroscientists and philosophers may hold the position that animals do not feel pain in the morally relevant sense or naturalist arguments that reject this view on the grounds that animals are “meant to be eaten.”

individual consumer in an industry that produces and tortures billions of animals and services millions of customers, my decision to purchase or refrain amounts to practically nothing.

What does it mean to say that the economy is complex in such a way that it changes our moral status within the broader picture of the industry? Consider the following thought experiment:

(T1) Suppose that Ann comes across a pig, and she would like to consume its meat. The only tool available for her is a special kind of weapon that imparts unimaginable pain to the pig for an extended period of time, before finally ending its life. Uncomfortable with using this weapon, she employs the help of a friend Ben, who seems to have less reservations and concern for animal welfare than she does, and explains that while she would like to eat the pig, she does not want to be the person that inflicts this suffering. Ben happily complies, Ann closes her eyes, and later she ends up with the meat.

In this circumstance, Ann's desire to consume the pig meat and her subsequent employment of Ben is directly the cause of the suffering of the animal. Surely, Ben is also guilty of causing the suffering directly by his actions, but Ann harbors a significant amount of the moral responsibility as well. Singer's welfare calculation is applicable here with its full force—it is morally objectionable to kill the animal because the suffering of the animal heavily outweighs the enjoyment Ann would receive from consuming it.

Now suppose that when Ann asks Ben to use this weapon to kill the animal, Ben claims that he is too preoccupied with schoolwork, and is unable to help Ann by literally killing the animal. However, Ben himself has a friend, Ben<sub>2</sub>, who he knows can help Ann—the best Ben can do is to ask Ben<sub>2</sub> on Ann's behalf, and bring her the pig meat at a time convenient for him. Ann appreciates Ben and Ben<sub>2</sub>'s help, and with some coordination, Ben<sub>2</sub> uses the weapon on the pig and Ann receives the pig meat.

Has anything about the situation changed in terms of moral responsibility from a consequentialist perspective? Ben has clearly changed his role, but his facilitation of this process



seems to be morally important. The causal situation seems not to have changed, because it really is still Ann's desire and employment of Ben which is the reason Ben<sub>2</sub> uses this weapon on the pig. We should still judge this situation with the same moral framework—it is morally objectionable to procure the pig meat in this way.

In fact, if we abstract this case to the broader perspective by supposing that Ben<sub>2</sub> serves as the communicator of Ben<sub>3</sub>, who asks Ben<sub>4</sub>, who emails Ben<sub>5</sub>, and so on, creating a chain of  $n$  Bens, each of which serve as one facilitator of this transaction (and nothing else), nothing about the moral and causal situation ever changes. Ann's role in this causal change does not shift because the number of Bens increased from  $n$  to  $n+1$ , and inductively this means that Ann's action is still morally objectionable regardless of the number of Bens.

The key observation here is that Singer's argument can be interpreted as persistence through one specific chain of the economy. The sheer size of the problem does not change the fact that if my dietary habits hurts just one animal of this population and I am the catalyst of this chain, we ought not to do it on consequentialist grounds. The idea that the economy is *too complex* for an individual to have significant power is vague at best, and upon this analysis, cannot appeal to the number of suppliers in between, nor their role in the chain or own motivations, if the primary and original impetus for the chain reaction was consumer demand. If one dimension to appealing to complexity to show why the individual is causally irrelevant, this argument fails or at least heavily underdescribes the case.

Another strategy for the consequentialist is to argue that the economy, although complex, is necessarily responsive to consumer demand if it is big enough. The line of reasoning is that even if one purchase might not have a global consequence, enough purchases do; this is sometimes referred to as the *threshold argument*. The specific situation that Kagan uses is

described roughly as follows. Suppose that in a small economy, a given order of chickens is 25 chickens, and this means that when the butcher notices that she has sold all her chickens, she orders 25 more. You, as a consumer, have no special information into the inventory, and so when you go to the counter, deliberating on whether to make a purchase, you have to assume that there is a 1 in 25 chance that you purchase the chicken that specifically runs out the supply and triggers the butcher to order 25 additional chickens. Therefore, the expected disutility of the purchase is  $(1/25)*25 = 1$  chicken,<sup>4</sup> and along classic consequentialist lines, we ought not to purchase the chicken (Kagan, 122). The threshold argument can be seen as the act-consequentialist's strategy of appealing to expected utility by arguing that the welfare-calculation still yields a negative result, and so we ought not to purchase factory farmed meat.

### ***Nefsky's Argument Against Thresholds***

Thus, consequentialists can capitalize on observations about responsiveness to consumer demand to illustrate that there is still space in the production chain for the individual to have real impact. However, this characterization of the economy has been contested by philosophers who argue against consequentialism and offer alternative viewpoints. Most notably, Kagan's argument has been criticized by Nefsky, who interprets Kagan as suggesting that "... in all triggering collective harm cases the relevant acts are *guaranteed* to have a negative expected utility" (369). Nefsky claims that this guarantee cannot be true, and the calculation provided in Kagan's argument only works because he has described the situation to fit the mathematical result. Without empirical evidence that these constants are reasonable approximations of real-world situations, the thought experiment is a contrived fiction. For example, if the threshold were

---

<sup>4</sup>  $1/25 =$  probability you are the difference maker, and 25 chickens being the disutility associated with this action.

instead  $t = 1000$  and the additional order  $n = 50$ , the utility calculation would be  $50 \cdot (1/1000) = 0.05$ . Nefsky also casts doubt that thresholds in producer supply would mean the *factory farm* would alter their production scale, stating that it's likely that even if the grocer did not order their expected amount, the factory farm would shift its focus into finding another grocery store rather than accept the reduction in profits and slow their operation (370).

### ***Budolfson's Inefficacy Objection***

Budolfson's arguments also rely on a different perspective of the economy. His broad goal is to demonstrate that the economy is structurally designed for the individual not to matter, so that *any* argument that relies on probabilistic accounts for expected utility in contributing to a threshold severely overestimate this impact. He points out that at each step of the supply chain, providers institute "buffer zones"; in order to ensure that demand is not missed due to a lack of supply on their part or to protect against losses from faulty products, providers will always order more than they need (1718). Let  $h$  be the number of chickens currently purchased by a supplier, and  $c$  represent the size of the buffer instituted in their ordering strategy. As long as  $h + 1 < h + c$ , whether I buy a chicken literally does not have an impact because  $h + c$  chickens have been killed whether I buy it or not. Rudolfson's argument is that the expected utility calculation cannot simply be  $(1/n) \cdot E$  as Kagan would have it, where  $n$  is the number of people involved and  $E$  is the total effect of the group of  $n$  actors. In his own words,

"Problem with the reasoning is that it overlooks the fact that we can know enough about the supply chains in both cases to know that threshold effects are not sufficiently likely and are not of sufficient magnitude to drive the expected effect of consumption anywhere close to the average effect."

Thus, Rudolfson's argument has a similar structure as Nefsky's—based on certain features of the economy, the consequentialist's calculations are dismissed as unrealistic and unlikely to yield the

correct answer. For instance, suppose I go to Whole Foods to pick up chicken for dinner, but I am struck with the realization that there might be a significant chance that my purchase contributes to the suffering of an animal, so I opt instead for a vegetarian option. In some iteration of the supply-calculation, Whole Foods has already accounted for my decision here—trying to accommodate the variability in the purchasing patterns of its patrons, it has ordered the  $c$  more chickens. In doing so, I don't make any change in their ordering strategy.

To what extent are these arguments conclusive? It may be the case that Kagan factually has the incorrect numbers, but the concern for the utilitarian calculation is only that the number which captures the disutility of their action outweighs the pleasure that one derives from purchasing meat. Budolfson is correct that if the existence of buffer zones really do economically render the individual impotent, that is, the probability that their effect on the supply chain is zero, consequentialism would necessarily have to conclude that there is no negative in their expected utility calculation, regardless of the potentially large magnitude of disutility.

However, other philosophers and economists have criticized whether *this* characterization of the economy really introduces this zero term. For example, Halteman and McMullen have written persuasively on the role of the individual consumer in the factory farming industry and argue that there are characteristics of the economy which indicates that the individual consumer *can* really make a difference. First, they notice that Budolfson's argument about buffer zones does not necessarily change the application of the threshold provided by Kagan—buffers can just be thought of as increasing the threshold, since they are just additional resources that are used up before being replaced. Alternatively, Halteman and McMullen also observe that buffers are likely to be constants in the production system, meaning that consumer demand does not have any effect on the magnitude of the buffer. The key point is that “the mere knowledge that there is

waste in the system overall... does not give the customer any information about whether their purchase will influence the retailer to order more or fewer products. Similarly, the knowledge that buffers exist does not tell a consumer anything about their location relative to those buffers” (12). Thus, Budolfson’s argument that the knowledge of past thresholds and values renders the calculation as predictably insignificant fails—the consumer does not know, and the magnitude and efficiency of the buffer zone is ultimately more dependent on the particular production chain.

Additionally, Halteman and McCullen argue against Budolfson’s characterization of the market’s response to decrease in consumer demand. To reiterate, Budolfson’s claim is that the priorities of the producers are to maximize their profit, so their strategy in the face of a decrease in consumer demand is not necessarily correlated to decreasing their production. Halteman and McCullen offer the following economic observations:

“When the demand for chickens drops, the result is not that a particular farmer chooses to hatch one less chicken. Instead, the threshold-triggered event is a particular grower’s failure to get a contract to raise birds at all, or a delay in the next shipment of birds, a switch to a different type of agriculture, or a rancher’s choice to sell her land to a developer... Standard economic models of these industries would predict that in the long run, a competitive market will be much more responsive than in the short run... In a large industry, moreover, competition will always drive the price to a point where the marginal producer is right on the edge of non-participation in the market” (15–16).

Halteman and McCullen add a different perspective to the causal picture and threshold breakers. The “rancher’s choice to sell her land to a developer,” for instance, means that the consumer demand has destroyed the potential for future farms to produce more cattle to torture, and this is an incredibly significant disutility that is prevented if the operation is shut down. Another argument proposed by Halteman and McCullen is that the significance of the individual in the economy merits a broader economic analysis than merely isolating the effects of one particular purchase. The thought experiments of Kagan and Nefsky have largely focused on one individual

and their reason to make one specific purchase and have used *only* the expected utility of this decision to evaluate whether to purchase or abstain. If the real economic power of an individual's actions is substantial taken as an aggregate in the long run, these thought experiments may be misleading and a more accurate picture of the consequences of an individual's purchasing habits must incorporate their net contribution over a longer period of time.

## PART 2. A PROPOSED SOLUTION

### *Reflections on the limits of consequentialism*

We ought to observe that economic theories can diverge and give different characterizations of the probability that a consumer makes an impact. Additionally, the actual effects of the purchases can be thought to trigger different types of events, ranging from being an insignificant blimp in a supplier's ordering strategy or shutting down an entire factory, all with different assignments of disutility.

Given the discussion above, it's doubtful that this situation can be meaningfully captured by the classic utilitarian calculation which is offered by Singer and Kagan and criticized by Nefsky and Budolfson:<sup>5</sup>

$$EU(\textit{purchase}) = p * n + \alpha$$

where  $p$  is the probability of triggering a threshold,  $n$  is the number of chickens that the trigger would affect, and  $\alpha$  is the positive utility that a consumer gets from purchasing the meat. Rather, trying to incorporate all the economic elements of the calculation into one expected utility value might look something like:

$$EU(\textit{purchase}) = p_1n_1 + p_2n_2 + p_3n_3 + p_4n_4 + \dots + p_kn_k + \alpha$$

---

<sup>5</sup> *EU represents the expected utility function.*

where  $p_a n_a$  refers to an economic event  $a$ , with probability  $p_a$  and utility  $n_a$ .<sup>6</sup> Singer and Kagan can be thought to simply refer to one  $p_m n_m$ , where  $p_m$  represents the probability of being the threshold purchase and  $n_m$  the associated disutility of crossing the threshold, while Budolfson argues that buffers render  $p_x$  for all  $x$  too close to zero. Halteman and McCullen argue that economic theory suggests some  $p_x$  are not close to zero, and so on and so forth.

The implications of this diversity of opinion for the consequentialist is that each framework of the problem seems plausible, and it is unclear which economic theory more accurately describes an individual consumer's actual predicted expected utility calculation. It would be arbitrary to accept one characterization of the economy over another without a reasonable amount of certainty in its correctness or superiority, and the philosophical debate surrounding the factory farming problem clearly demonstrates a lack of consensus on which interpretation is more likely to be true. The role of the consumer in the factory farming industry is an economic problem for which it is clear we are missing a large amount of data and research.

This appears to turn the consequentialist threshold argument and its rejection into contingent arguments based on contested empirical statements about the economy. If theory A describing the economy is true, then consequentialism may state that we ought to refrain from purchasing meat; if theory B describing the economy is true, then consequentialism may state that we ought to purchase the meat. However, deciding between theory A and theory B is not an empirical fact that is available to the average consumer.

Therefore, it's rational for the consumer to be unsure. This might not be true of all cases—certain situations, like small closed economies, may lead to more predictive answers, but

---

<sup>6</sup>  $k$  is assumed because the equation needs to be finite (or more accurately, the number of terms with positive probability need to be finite).

for the average consumer participating in the large American meat industry, it's completely unclear whether we actually make a difference.

This uncertainty renders the inefficacy objection inconclusive. In light of the empirical uncertainty regarding the economy, it appears that consequentialism does not give an answer contradictory to our moral principles, but rather fails to give an answer at all. Recall the structure of the objection, formally presented here:

- (1) If consequentialism is true, then if a consumer's purchase of factory farmed meat yields no disutility, the consumer ought to purchase it.
- (2) A consumer's purchase of factory farmed meat yields no disutility.
- (3) Assume consequentialism is true.
- (4) Then, if a consumer's purchase of factory farmed meat yields no disutility, the consumer ought to purchase it.
- (5) The consumer ought to purchase it (2 and 4).
- (6) But the consumer ought not to purchase it (because it's wrong).
- (7) Contradiction. Consequentialism is false.

Since (2) fails to be demonstrable, the argument is invalid, or at least must be framed to make (2) a conditional premise or an added assumption, rather than presented as an empirical fact.

However, it would be a mistake to interpret this analysis as an automatic vindication for the consequentialist. Consequentialists may boast that their moral evaluation relies impartially and objectively on data, but the factory farming problem presents a case in which this data is incredibly difficult to find, measure, and interpret. It may be the case that certain individuals in society, because of their celebrity status or influence, really do have consequentialist reasons to refrain from purchasing meat, simply because of the high probability that a large number of people will follow suit.<sup>7</sup> This is only a small subset of the population, and does not answer any question with regards to the average consumer.

---

<sup>7</sup> Consider what would happen if Peter Singer were caught eating a hamburger.



In summary, consequentialism is only able to yield a moral judgment when there is a clear empirical understanding of the causal consequences of a given action. We do not know the empirical claims when it comes to our participation in the factory farming industry and the effects of our market presence, and so both consequentialism and the inefficacy objection suffer from the same unclarity.

### ***Examining the full range of consequences***

Both the threshold argument and its objections take for granted the measurability of the economic effects of an individual's purchase and in doing so, actually create a narrow scope of what supposedly informs the consequentialist's perspective. This is why the problem feels so intractable in light of our lack of insight into the economy. However, this is not the only way to frame the consequentialist's approach to the factory farming problem. Recall the formulation of the inefficacy objection's first premise:

(1) If consequentialism is true, then if a consumer's purchase of factory farmed meat yields no disutility, the consumer ought to purchase it.

This premise in and of itself is problematic—consequentialism is *not* committed to the position that a lack of disutility yields the answer that one *ought* to purchase meat. In the case of factory farming, it is rational to be confused, doubtful, or even skeptical of one's impact, but from this it does not necessarily follow that the consequentialist concludes the permissibility of all actions. Consequentialism instead instructs us to align our actions to produce the best consequences possible. If the inefficacy objection is really to succeed by demonstrating that consequentialism gives us the wrong answer, it must actually show that purchasing meat is the possible world with the *most* utility, not just demonstrate that there *might* be one positive utility.

The measure problem for consequentialism is a familiar one, one so central to the calculation and project of utilitarianism that Mill foresaw it and gave a response. Mill specifically objects to is the claim that consequentialism unreasonably demands an agent run a supposedly incredibly difficult and complex calculation before each of action. In *Utilitarianism*, Mill provides an answer worth quoting in full:

“The answer to the objection is that there has been ample time, namely, the whole past duration of the human species. During all that time mankind have been learning by experience the tendencies of actions; on which experience all the prudence, as well as all the morality, of life is dependent. People talk as if the commencement of this course of experience had hitherto been put off, and as if, at the moment when some man feels tempted to meddle with the property or life of another, he had to begin considering for the first time whether murder and theft are injurious to human happiness” (34-35).

Mill’s analysis of the calculation problem is simply that our intuitions guide us to recognize certain behaviors as consistently or even probabilistically leading to the correct moral decisions on consequentialist grounds, and that an individual can appeal to this when assessing the consequences of our actions. In the case of murder, it’s quite clear that refraining from murder is something that, in an effectively total number of cases, is one that does not merit sophisticated calculation to understand as wrong on consequentialist grounds. Put simply, a consequentialist does not have to redo utility calculations of their actions each time.

Mill’s approach is applicable to the factory farming case as well. In light of the uncertainty of the expected utility calculation when regarding an individual’s power in the economy, a consequentialist can—and should—draw on other considerations which have a more transparent utility. Analyzing these provides a fruitful avenue for the consequentialist to overcome a state of ambiguity in assessing, on consequentialist terms, the correct moral action and respond to the inefficacy objection.<sup>8</sup> Thus, a different approach for the consequentialist

---

<sup>8</sup> Of course, Peter Singer is not blind to this consideration. Chapter 4 of *Animal Liberation* is devoted to describing the personal benefits one can gain from a vegetarian lifestyle.

confronted with the economy is to rely on something like virtues or patterns of behavior. They can acknowledge that there is a strong utilitarian case for adopting a certain set of virtues because they are more likely to produce action that leads to better consequences. This means that the consequentialist answer does not require the reacquiring and reassessment of data every time, since we rely on our moral intuitions to instead judge for ourselves which actions are likely to produce more utility.

### ***A Virtue Based Approach***

The virtue-based approach I offer here builds off the assertion that certain practices are likely to be more productive if our goal is to produce the best utility. We may not have the insight into how our isolated purchases affect the economy, but we can be reasonably certain that a life of refraining from engaging with this industry is more likely to decrease suffering to animals than a life full of consumption. In other words, our uncertainty regarding the consequences of our individual purchases does not imply an uncertainty regarding which behaviors and dietary preferences are *more likely to produce* better consequences. If a person has the virtue that callous treatment of nonhuman animals is wrong, then they are more likely to refrain from purchasing factory farmed meat, and even more likely to refrain from abusing pets, supporting zoos, etc.

This nuanced approach thus sees virtue as having consequentialist merit, and recognizes the compatibility of consequentialism with moral intuitions, principles, and rules. Note that this is a different position than virtue ethics; instead of appealing to virtue as the source of ethics, we evaluate and accept virtues insofar as they increase the probability that a given individual's actions will produce beneficial consequences.

In fact, a stronger view might be to define virtue fully in consequentialist terms. For instance, Julia Driver in “Uneasy Virtues” observes an inadequacy in Aristotelian conceptions of virtue because “not only is knowledge not essential to virtue, but other intentional states...are not necessary.” She argues that a more plausible approach to defining virtues is to state that “virtues are character traits (these include dispositions to be motivated in various ways conjoined with yet other dispositions) that produce good effects” (67). In other words, Driver aims to show that a consequentialist framework can fully and more adequately accommodate virtues. Driver also presents a sophisticated account of the epistemological considerations surrounding a consequentialist-approach to virtues. In demonstrating the strength of this definition, she notices, “...when we do see that we have misjudged the consequences of a trait, we change our judgment of the trait’s status as a virtue” (84). Her main example of this is the virtue (or former virtue) of chastity—society was once in a state in which women practicing chastity supposedly had a wide range of benevolent social consequences. However, as society evolved, and the value chastity held in society deteriorated or indeed flipped, chastity was abandoned as a virtue (84–85). On this account, we consider things virtuous when they have utility—and this can be for personal or social reasons—and we can abandon them if through the evolution of society or circumstances the utility of this virtue is no longer existent or apparent.

This theory of virtue reveals the compatibility of consequentialism and virtue ethics. The inefficacy objectionists have sometimes described consequentialism as being defeated by superior moral theories, but upon closer examination, an analysis of how consequentialists should assess virtues and principles provides an avenue for the consequentialist to draw on the strengths of other approaches and step beyond their uncertainty.

For the remainder of the paper, I seek to incorporate a similar account of virtues to the context of the factory farming, and my project is simply to delineate why a virtue-consequentialist approach is appropriate, intuitive, and productive by providing examples of virtues I defend as relevant to the problem. My intention is not to provide a decision procedure for evaluating or isolating a *particular* set of virtues. However, if we can demonstrate that there *are* virtues which have transparent reasons for a consequentialist to endorse and practice, we effectively provide a way for the consequentialist to accept this approach and fully respond to the inefficacy objection.

### ***The Virtues***

The virtues that I have in mind are altruism, mindfulness, and compassion. While these virtues are non-exhaustive, when combined, they make a strong case for vegetarianism *on consequentialist grounds*. My focus for these virtues is largely to extrapolate on how these virtues might guide a consumer at a grocery store to make ethically responsible purchases.

#### *Altruism or Equal Consideration*

I define altruism as the practice or intention of giving weight to the concerns of others. By consciously incorporating and considering the presence of others in our lives and acknowledging their preferences and values alongside ours, altruism respects the moral agency and importance of others and counterbalances selfishness and self-aggrandizement.

Altruism is an important virtue to hold for the consequentialist because our commitment to consider others means that we are more likely to incorporate their welfare when assessing the total consequences of our actions. Simply put, altruism encourages us to step outside giving a

false sense of importance and priority to our own pleasures and give equal consideration to all beings who are potentially affected by our actions. I suspect that this is the central intuition behind Peter Singer's *Animal Liberation*, as he states, "If a being suffers there can be no moral justification for refusing to take that suffering into consideration" (Singer, 8). In the context of the factory farm problem, altruism then serves an ethically important psychological function, particularly by defeating the socially held notion that food is *merely* a product for personal gain and pleasure and forcing us to consider the animals behind our food.

This view is indirectly endorsed by Jonathan Baron, who seeks to provide a framework where utilitarians can still act on probability in the face of uncertainty. Baron acknowledges that in practical problems, the true probability is not something always available. However, he then claims that this may be symptomatic of the utilitarian's own concerns. If a utilitarian is exclusively interested in self-interest, voting is irrational strictly because of the very low probability that their vote actually makes a difference *to them*. But when regarding the well-being of themselves *and* other and future people, altruism gives us reason to consider a different set of consequences:

"For a utilitarian voter with some altruism toward people in general, it is thus rationally worth the effort to vote, and to be sufficiently informed to vote for the better side, if the ballot contains proposals (or candidates) that affect very large numbers of people. It is not rational to vote at all if the voter thinks in terms of narrow self-interest alone... We should see this for what it is, a form of immoral behavior, something that results from a way of thinking that we, at least those of us who do care about people in general, should want to discourage" (Baron, 22).

Thus, altruism is a virtue worth pursuing for the utilitarian. The analogy holds well when applied to the factory farming problem—if a consumer is concerned only with their own pleasure and happiness, it may seem that the utility of refraining from this purchase is negligible or not worth pursuing because it has no benefit for themselves. However, an altruist, or anyone with a concern

for animals and their suffering resists this pattern of thinking. By choosing to put the considerations of animals on higher footing than presented by grocers when purchasing food at a grocery store, consumers are more likely to prefer other options and refrain from purchasing meat, as well as other products associated with objectionable practices.

One counterargument to the consequentialist's acceptance of altruism is to claim that altruism as a general conduct may create instances in which the utility afforded to others comes at a cost to oneself. For instance, while donating a proportion of one's income to a cause is morally praiseworthy, altruism may appear untenable if it forces us to sacrifice too much of our own resources for the sake of others.<sup>9</sup>

However, it seems like in the case of dietary modifications, altruism is incredibly unlikely to yield cases where we are sacrificing something of such moral importance or incurring personal detriment. As discussed previously, modern medical research has consistently found that vegetarian diets provide a robust profile of nutrition, and so there is very little to lose or be sacrificed, as evidenced by the millions of healthy practicing vegetarians in the world.

### Mindfulness

Mindfulness is the characteristic of critically thinking, analyzing, and remaining conscious and aware of the reality around us. Spoken in these terms, mindfulness appears to be an ephemeral, perhaps idealistic state of mind, but *practicing* mindfulness in modern consumer culture presents a significant challenge. Capitalist institutions rely on the lack of transparency into their practices to sustain their public image, and the industry of factory farming operates merrily on the ignorance of the average consumer to research the source and evil of the products

---

<sup>9</sup> See Peter Singer, *Famine, Affluence, and Morality* for a possible endorsement of this type of altruism from a consequentialist.

they choose to buy. Furthermore, the culture of celebrity cooking, indulgence, and viral food trends further dulls our awareness into the circumstances which make such products possible. As a society, we tend not to think of food beyond our gustatory preferences or vehicle for personal pleasure.

Mindfulness is then both a method of resisting this social effect of ignorance and a personal commitment to recognize the story behind our purchases. Consider the following endorsement of mindfulness by Dale Jamieson. In “When Utilitarians Should Be Virtue Ethicists,” Jamieson analyzes climate-change as a problem for the consequentialist in which they should similarly adopt a virtue theoretic approach. In delineating some virtues we might include, he states:

“Much of our environmentally destructive behavior is unthinking, even mechanical. In order to improve our behavior we need to appreciate the consequences of our actions that are remote in time and space... Making decisions in this way would be encouraged by the recognition of a morally admirable trait that is rarely exemplified and hardly ever noticed in our society” (Jamieson 23).

Thus, Jamieson too recognizes the importance of mindfulness in light of the current social conditions of consumer habits. The simple intuition is that awareness of the source of our everyday food, the trace of our commercial products, and the social conditions which make those possible will yield to more informed and more accurate assessments of which products we ought to prefer. All things being equal, if option A produced more disutility than option B, the consequentialist prefers option B. However, a significant barrier to assessing whether things are equal and whether option A is better than option B is a total ignorance of the histories of those options in the first place. This is not to suggest that calculations like these are necessarily easy or covers all cases, but in a large part of our consumer lives, we can do better to seek out information. The reason we ought to do this is fully explainable on consequentialist grounds—



more information allows us to make better judgments about which products we should purchase and those from which we should refrain.

### Compassion

I define compassion as emotional sensitivity for the circumstances and suffering of others. While altruism describes developing an awareness of consideration of others and moving beyond self-interest, I understand compassion as necessarily having a sense of sympathy or emotional connection with another being.

The value of compassion to the consequentialist has been articulated well by Nathan Nobis in “Vegetarianism and Virtue: Does Consequentialism Demand Too Little?”, in which he adopts a similar line of reasoning of incorporating the virtue-ethicist’s perspective into the consequentialist’s approach. In trying to merge these two moral theories, he writes,

“...pre-theoretically, it seems that, all else equal, a person will bring about more goodness if she has the virtue of compassion, cares about and is sensitive to unnecessary cruelty and suffering (wherever it is found, in humans or animals)... It seems exceedingly unlikely that anyone would, in general, come to treat other humans worse were she to become a vegetarian or vegan out of compassion or sympathy for animals. In fact, the opposite seems likely” (Nobis 153).

Compassion is a virtue that seems utterly lost when confronted with the statistics of factory farming and the types of practices which occur in these facilities. Nobis cites Russ Shafer Landau, who suggests the vice practiced by meat-eaters is “that they display an indifference to the cruelty that went into the ‘production’ of their ‘goods’” (Nobis 150). Compassion, in some ways, is then a corrective virtue to combat the callousness caused by modern consumer culture and the ignorance afforded by distance from the objectionable processes that make the comforts of our lives possible. By increasing our emotional sensitivity to the suffering of animals, especially in light of the vast uncertainty we face in evaluating our positions in the economy, we

are driven to more consistently vet the products we buy, more rigorously consider whether we ought to purchase products which are associated with animal suffering, and understand the gravity of the problem, the importance of our activism, and the need for change. Our lack of compassion for these animals is highlighted by comparing our treatment of our pets and the treatment of the animals in a factory farm—most of us would cringe and protest if it were our pets which were treated this way, *even if* we could do absolutely nothing about it. If we just submit to internalizing that some animals are worthy of our moral attention while others are not, we decrease the change that these animals are ever liberated.

No change for the better will be successfully implemented if the modern consumer is indifferent to the suffering of nonhuman animals and because they are emotionally unaffected, considers this as not their problem.

### *A Possible Objection*

These three virtues provide a plausible starting point for developing a consequentialist approach to the factory farming problem which seeks to identify virtues we should adopt on the grounds that they produce behavior likely to yield positive utility. A possible objection to this approach, however, is that the choices of virtues is arbitrary, and it is unclear why the uncertainty of the effect of our actions is granted but we do not mirror the same lack of knowledge to the virtues we pick. In short, if the consequences of our purchases are unknown to us, why are the consequences of our virtues so transparent?

In theory, it is possible that we can make mistakes about which virtues to incorporate, and some characteristics may not be tenable on this view. For instance, suppose that we grant abstinence as a virtue—in morally risky situations, the abstinent person avoids the potential for

disaster and categorically refuses to partake in any of these practices. Depending on the situation, this may yield positive and negative consequences. If abstinence is seen to mitigate personal impulse to purchase and consume meat and promote consistency to vegetarian practices, it seems to have beneficial consequences to support vegetarianism. However, abstinence taken as a blind truth may yield different cases as well. Consider the food waste problem, in which a meat-based dish is about to be thrown out, and a virtuous person happens upon this circumstance in a fit of hunger. A utilitarian is likely to judge this instance as a clear utility calculation. If the virtuous person abstains, the food is wasted, yielding zero utility. However, consuming the food yields in one clear positive utility—the alleviation of hunger. A more sophisticated consequentialist analysis of the situation may consider alternative feelings in the equation as well, such as the guilt the virtuous person might feel when compromising their standards, but the fundamental intuition here is that some characteristics, even when founded on good intuitions, may cause a net negative utility. We certainly don't want to practice abstinence if it increases unnecessary food waste dramatically. *How can we tell?*

While this problem is complex, and my account does not provide a decision procedure to assess each virtue, it would be a mistake to consider this problem as a refutation of the whole approach. We may not know with absolute certainty that altruism, mindfulness, and compassion, guarantee a positive utility when followed, but I find it difficult to consider this as a real objection to why, in the general case, we ought to be skeptical about these virtues. For instance, suppose that you are considering donating \$10 to Oxfam because of an altruistic desire to aid others. It is completely within the realm of possibility that this donation actually yields a negative utility—perhaps three people watched you donate, and now refuse to donate themselves because in their view someone else has already done it. A straightforward utility calculation

suggests that had you not done it, the utility would have been \$30, but your action caused it to be \$10, so this was the wrong thing to do.

Does such a possibility give us reason to reject altruism on the basis that at each altruistic instance we can actually be incurring negative utility instead? This seems like a form of radical skepticism unfairly pressed to the consequentialist when evaluating virtues. As Mill wrote, it is not incumbent to the consequentialist to consider all alternative worlds at all actions at every point in time—we can rely on our previous experiences and our developed moral practices and intuitions for guidance in shaping our conduct. Altruism, in general, produces good.

Some virtues may need further refinement in nuanced cases when the answer is not clear, but the account of virtues given here can fully accommodate a picture in which one is accepted then later rejected assessing its utility. However, to argue that consequentialists cannot adopt compassion and empathy towards animals as a good trait that yields positive utility on the view that there may be boundary cases in which practicing this trait turns out to yield the wrong consequentialist answer seems contrived. Jamieson states this position concisely, when he says:

“...there is no algorithm for designing the optimal utilitarian agent... Nevertheless, we have better and worse people and constitutions, and sometimes we know them when we see them. It might be nice to have a calculus that we could apply to constitutions and character, but absent this, we can still go forward living our lives and organizing our societies” (Jamieson, 21).

The burden of proof is on the skeptic to demonstrate why character traits like compassion and mindfulness merit such a deep-rooted epistemological concern, when our moral intuitions so obviously support them. In other words, the uncertainty we face as consumers in an economy much bigger than we are is incongruent to the uncertainty we might face when assessing which virtues to adopt. In the former, the problem is inherently epistemologically difficult, and part of the confusion lies in the difficulty of obtaining the required data and the variety of economic

theories and opinions which characterize the market differently. In contrast, our adoption of virtues seems to be a fundamentally different problem—although admittedly we cannot be absolutely certain, the plausibility of certain character traits to yield positive change has been informed by our communal moral standards, laws, and intuitions.

### ***Conclusion***

The inefficacy objection to consequentialism states that consequentialism should be abandoned because in collective-action problems, where individuals do not have reason to work towards a solution, consequentialism appears to give us an incorrect moral answer. In the case of factory farming, the inefficacy objection might state that consequentialism tells us we should buy factory-farmed meat and we really should not. I started by establishing that our relationship with the economy is complex and the wide variety of economic theories suggests that there is no one correct description of the economy that consequentialism can be said to necessarily adopt. Therefore, it is rational to be unsure. However, in light of this uncertainty, the consequentialist can still follow virtues which can be justified by their propensity of effectuating positive utility in the world, each of which would prefer vegetarianism to meat-eating. I offered altruism, mindfulness, and compassion as three examples of such virtues.

Consequentialism's project is ultimately to assign to morality the project of making the best world possible. This is a clear message that we should embrace, and sometimes our role in bringing this about is not to expect our actions to trigger total iconoclastic destructions to our social and economic structures but rather in implementing small, consistent actions guided by reasonable principles developed on consequentialist grounds.

## WORKS CONSULTED

- Baron, Jonathan. "Uncertainty and Probability within Utilitarian Theory." *Diametros*, vol. 53, Uniwersytet Jagiellonski, Sept. 2017, pp. 6–25, doi:10.13153/diam.53.0.1098.
- "Becoming a Vegetarian." *Harvard Health Publishing*, Harvard Medical School, 23 Oct. 2018, <https://www.health.harvard.edu/staying-healthy/becoming-a-vegetarian>.
- Budolfson, Mark Bryant. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response." *Philosophical Studies*, vol. 176, no. 7, Springer, 2019, pp. 1711–24, doi:10.1007/s11098-018-1087-6.
- Capps, Ashley. "Debeaking Video Shows Standard Practice on Free-Range Egg Farms." *Free From Harm*, Free From Harm, 3 Nov. 2015, <https://freefromharm.org/animal-cruelty-investigation/debeaking-video-shows-standard-practice-on-free-range-egg-farms/>.
- Driver, Julia. *Uneasy Virtue*. Cambridge University Press, 2001.
- "Factory Farm." *Merriam-Webster*, Merriam-Webster, [https://www.merriam-webster.com/dictionary/factory farm](https://www.merriam-webster.com/dictionary/factory%20farm).
- "Factory Farms." *A Well-Fed World*, A Well-Fed World, <https://awellfedworld.org/factory-farms/>.
- Good, Kate. "5 Ways Factory Farming Is Killing the Environment." *One Green Planet*, One Green Planet, 23 Jan. 2019, <https://www.onegreenplanet.org/animalsandnature/factory-farming-is-killing-the-environment/>.
- "It's Time to Try Meatless Meals." *Mayo Clinic*, Mayo Foundation for Medical Education and Research, 26 July 2017, <https://www.mayoclinic.org/healthy-lifestyle/nutrition-and-healthy-eating/in-depth/meatless-meals/art-20048193>.
- Jamieson, Dale. "When Utilitarians Should Be Virtue Theorists." *Utilitas*, vol. 19, no. 2, 2007, pp. 160–83.
- Marsh, Kate A, et al. "Protein and Vegetarian Diets." *The Medical Journal of Australia*, The Medical Journal of Australia, 29 Oct. 2013, <https://www.mja.com.au/journal/2013/199/4/protein-and-vegetarian-diets>.
- McMullen, Steven, and Matthew C. Halteman. "Against Inefficacy Objections: The Real Economic Impact of Individual Consumer Choices on Animal Agriculture." *Food Ethics*, vol. 2, no. 2–3, 2019, pp. 93–110, doi:10.1007/s41055-018-00030-4.

- Mill, John Stuart. *Utilitarianism*. Parker, Son and Bourn, 1863,  
<https://play.google.com/store/books/details?id=HCY2AQAAMAAJ&rdid=book-HCY2AQAAMAAJ&rdot=1>.
- Nefsky, Julia. “Consequentialism and the Problem of Collective Harm: A Reply to Kagan.” *Philosophy and Public Affairs*, vol. 39, no. 4, 2011.
- Nieman, D C. “Vegetarian Dietary Practices and Endurance Performance.” *The American Journal of Clinical Nutrition*, vol. 48, no. 3, 1 Sept. 1988, pp. 754–761.,  
doi:10.1093/ajcn/48.3.754.
- Nobis, Nathan. “Vegetarianism and Virtue: Does Consequentialism Demand Too Little?” *Social Theory and Practice*, vol. 28, no. 1, 2002, pp. 135–156. *JSTOR*,  
[www.jstor.org/stable/23559206](http://www.jstor.org/stable/23559206).
- Reese, Jacy. “US Factory Farming Estimates.” *Sentience Institute*, Sentience Institute, 11 Apr. 2019, <https://www.sentienceinstitute.org/us-factory-farming-estimates>.
- Richi, Evelyne Battaglia, et al. “Health Risks Associated with Meat Consumption: A Review of Epidemiological Studies.” *International Journal for Vitamin and Nutrition Research*, vol. 85, no. 1-2, 2015, pp. 70–78., doi:10.1024/0300-9831/a000224.
- Singer, Peter. *Animal Liberation*. HarperCollins Publishers, 2009.