

About this document

I wrote this collection of essays during 2005-2006, during roughly the year after I discovered the philosophy of utilitarianism. Looking back, I'm sometimes impressed with my former self but also very often embarrassed at how much I didn't know. My views have changed significantly in the intervening period. One big difference is that I now regard preventing suffering as overwhelmingly important in contrast to creating happiness for creatures that don't exist yet. I'm making this document available for historical interest, as a snapshot into my views at ages 18-19.

—Brian Tomasik, author of <http://utilitarian-essays.com/>

Utilitarianism

Brian Tomasik

2005 - 2006

“It appeared to me obvious that the happiness of mankind should be the aim of all action, and I discovered to my surprise that there were those who thought otherwise. Belief in happiness, I found, was called Utilitarianism, and was merely one among a number of ethical theories. I adhered to it after this discovery.”

—Bertrand Russell [39]

“The wealth required by nature is limited and is easy to procure; but the wealth required by vain ideals extends to infinity.”

—Epicurus [27]

“The day may come when the rest of the animal creation may acquire those rights which never could have been withholden from them but by the hand of tyranny. The French have already discovered that the blackness of the skin is no reason why a human being should be abandoned without redress to the caprice of a tormentor. It may one day come to be recognised that the number of the legs, the villosity of the skin, or the termination of the os sacrum, are reasons equally insufficient for abandoning a sensitive being to the same fate. What else is it that should trace the insuperable line? Is it the faculty of reason, or perhaps the faculty of discourse? But a full-grown horse or dog is beyond comparison a more rational, as well as a more conversable animal, than an infant of a day, or a week, or even a month, old. But suppose they were otherwise, what would it avail? The question is not, Can they reason? nor, Can they talk? but, Can they suffer?”

—Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation*, Chapter 17, Section 1 (1789) [4]

“Never, for any reason on earth, could you wish for an increase of pain. Of pain you could wish only one thing: that it should stop. Nothing in the world was so bad as physical pain.”

—George Orwell, *1984*, Part 3, Chapter 1 (1949) [25]

“We should not [...] think of our efforts as wasted unless they endure forever, or even for a very long time. We can make the world a better place by causing there to be less pointless suffering in one particular place, at one particular time, than there would otherwise have been. As long as we do not thereby increase suffering at some other place or time, or cause any other comparable loss of value, we will have had a positive effect on the universe.”

—Peter Singer, “The Escalator of Reason,” from *How Are We to Live?* (1995) [35]

Contents

| | | |
|----------|---|-----------|
| 1 | Preface | 4 |
| 2 | Foundations of Utilitarianism | 5 |
| 2.1 | Epistemological Foundation | 5 |
| 2.2 | Ethical Foundation | 6 |
| 2.3 | Utilitarianism as Happiness Maximization | 10 |
| 3 | Essays on the Intuitive Basis for Utilitarianism | 12 |
| 3.1 | From Materialism to Utilitarianism | 12 |
| 3.2 | Subjectivism | 14 |
| 3.3 | On the Origins of Speciesism | 16 |
| 4 | Essays on Metaethics | 19 |
| 4.1 | Intuition and Reason | 19 |
| 4.2 | The Role of Unreality in Moral Reasoning | 20 |
| 4.3 | Morality as a Stimulus-Response System | 22 |
| 4.4 | Utilitarianism as a Resolution to Ethical Disagreements | 23 |
| 4.5 | Do Implicit Preferences Matter? | 26 |
| 5 | Methods of Utilitarianism | 29 |
| 5.1 | General Procedure for Utilitarian Evaluations | 29 |
| 5.2 | Affecting Public Opinion | 32 |
| 5.3 | On the Precautionary Principle | 37 |
| 6 | Intuitive Discussion of Further Topics | 40 |
| 6.1 | On Triage | 40 |
| 6.2 | Should Nature Be Inviolable? | 41 |
| 6.3 | Is Happiness Relative? A Personal Conjecture | 43 |
| 6.4 | Why Determinism Doesn't Matter | 46 |
| 6.5 | Expectational Consequentialism | 47 |

| | | |
|----------|---|-----------|
| 7 | Essays on Applications of Utilitarianism | 51 |
| 7.1 | Is Terrorism Prevention Cost-Effective? | 51 |
| 7.2 | On Popular Opposition to Utilitarian Policies | 53 |
| 7.3 | The All-or-Nothing Fallacy | 54 |
| 7.4 | Does Veganism Make a Difference? [19] | 56 |
| 7.5 | How Much is a Dollar Worth? Some Conservative Estimates | 60 |

Chapter 1

Preface

- This document presents an assortment of essays that explore various aspects of utilitarianism.
- There may be minor inconsistencies among the pieces because I wrote them at different times.
- I have tried, as much as feasible, to state precise definitions for important terms. Of course, I am unable to do this for every word that I use, wherefore some terms will, unfortunately, remain a little vague.
- If you notice any factual, logical, or conceptual errors, please let me know. Have I omitted anything important? Have I missed an important idea or reference? What should I change? Please send along comments and suggestions.

Chapter 2

Foundations of Utilitarianism

2.1 Epistemological Foundation

In order to even consider questions of morality, one must first hold some basic assumptions about the way the universe works. Below are the epistemological assumptions on which utilitarianism depends.

Axiom: It is intelligible to think and talk about concepts, such as “existence.”

The purpose of this admittedly vague axiom is to overcome nihilist criticisms of philosophy in general. Nihilists deny basic ideas about intelligibility that we take for granted, and this axiom assumes away some of those challenges. Since there are so many ways in which a nihilist could deny the intelligibility of philosophy, I am unable to give a more precise statement of this assumption. But below is an example of the type of challenge that this axiom is intended to dismiss.

Example: Suppose I am conversing with a nihilist, and I make this statement: “Okay, so you deny lots of things. You deny any underlying structure to reality. You deny the meaningfulness of logic. But surely, you must admit that at least the following statement is true: You are thinking. How could that not be the case?”

To this the nihilist says, “What do you mean by ‘true’? That statement is unintelligible. Indeed, even the words that I’m articulating right now are unintelligible.”

In a vain attempt at refutation, I retort, “Your statement itself is a contradiction. You assert that words are unintelligible, but if words are unintelligible, then the words with which you made that statement are unintelligible. In order for your assertion to be true, it must be false!”

In reply, the nihilist merely sighs and says, “I refuse to acknowledge that anything you say is meaningful. Just stop trying.” Indeed, there is no way to convince the nihilist to acknowledge intelligibility.

Like a spacecraft that has passed through a black hole’s event horizon, the nihilist cannot be pulled back away from her position by conventional argument—because, of course, such conventional argument is unintelligible.

The purpose of this axiom is to move the discussion beyond the above dispute. If the axiom is accepted, then we can steer clear of the black hole and proceed ahead to other matters.

Axiom: The law of causation holds.

Axiom: Basic logical rules of inference like *modus poens* and *modus tollens* hold.

Axiom: Induction holds. That is, the past can help predict the future.

Axiom: Assume everything else that I’ve neglected to mention.

The purpose of outlining axioms is to make clear exactly what is being assumed. I have tried to do that here, but there are undoubtedly other assumptions that I am taking for granted because I am so accustomed to thinking within a framework in which they hold. If readers notice any important assumptions that I have left out, please let me know.

2.2 Ethical Foundation

Definition: A *preference* is any explicit or implicit emotional desire for the universe to be in one state rather than another.

Definition: An *explicit preference* is a desire that an organism thinks about abstractly as being a desire. This means that the organism speaks, writes, or internally verbalizes what it wants.

Examples of an explicit desire: A person feels hungry and thinks, “I wish I had some food!”

An adult sees a violent cartoon show and says, “It’s wrong that people make these types of shows for kids to watch.”

Definition: An *implicit desire* is a desire that an organism feels but does not think about as abstractly being a desire.

Examples of an implicit desire: A person has a stomachache and endures the pain. Even though the person doesn’t have an abstract thought that “I wish this pain would stop,” she nonetheless dislikes the experience. Thus, she has a preference for it to stop even though she doesn’t put that preference into words.

A battery-cage laying hen feels hungry, agitated, and miserable.
A baby mouse looks forward to the return of its mother to the burrow.

Definition: *Sentients* are things capable of having preferences.

I denote a sentient that exists at time t as $s(t)$. The set of all sentient organisms that exist at t is $S(t)$.

Examples of sentients: People, monkeys, cows, chickens, bluejays, and (potentially) computers of the future [17].

Examples of nonsentients: Amoebas, bacteria, trees, flowers, rocks, water molecules, air molecules, and computers of the present.

I should add that the above classifications are not axiomatic; rather, they depend upon our scientific understanding of the internal mental states of other organisms. If we were to discover, for instance, that amoebas feel pain, I would put them in the “sentients” category.

Axiom: Sentients exist.

By definition, this also means that it is possible for preferences to exist.

Axiom: Preference strength is cardinalizable.

Different preferences may be held with different degrees of intensity. The strength of a preference has a cardinal value (*e.g.*, 2 or 14).

In reality, we don’t know what an organism’s preferences are to any precise degree, so we can’t assign numbers like this in practice. I do so here and elsewhere merely for the purpose of clarity. This does not mean, however, that preference strength does not have *some* value in principle, even though we can never measure it exactly. In any event, we can at least get a good sense of relative magnitudes of preferences, and that’s generally sufficient for a utilitarian analysis.

Axiom: Intersubjective and intertemporal comparisons of preference strength are possible.

Suppose Person A wants a blueberry muffin, Person B wants a chocolate-chip cookie, and Person C wants orange juice. This axiom says that its possible to describe the preferences of A, B, and C in a common unit of preference strength. Suppose A has a preference strength of 5, B has a preference strength of 7, and C has a preference strength of 5. Then A and C have preferences of the same strength and B has a stronger preference. This

axiom also asserts that the same would be true across time. To see an example, simply let A, B, and C above represent a single person at three moments in time.

Definition: *Unconsciousness* is the state in which an organism does not know that it exists.

Think of unconsciousness this way: time stops, everything in your life stands still, and you feel (or, well, don't feel) the sensations of nonexistence.

Definition: *Utility, u* , is a variable representing the strength of fulfilled or abridged preferences. The unit of utility is a *util/unit time*.

To make comparison of utilities easier, I define preferences to be relative to unconsciousness.¹ Thus,

$u > 0 \iff$ an organism is in a state that it prefers over unconsciousness, and

$u < 0 \iff$ an organism is in a state where it would prefer to be unconscious.

The magnitude of u is proportional to the strength of the organism's preference.

Examples: A cow wants to find a patch of fresh grass. It would prefer the state of finding the grass over the state of unconsciousness by a magnitude of 5. Hence, finding grass will furnish 5 utils/minute.

A turkey in a factory farm had its beak trimmed two days ago, and it is now enduring chronic pain. Moreover, it feels cramped, stressed, and sick due to overcrowding. If it could imagine the abstract state of unconsciousness and compare that to its present condition, it would prefer unconsciousness with a strength of 150. Since the turkey is not in its preferred state, its preference is abridged, and it feels -150 utils/minute.

¹I should emphasize that choosing unconsciousness as a reference point is arbitrary. I could just as easily have chosen utility to be defined relative to the state of stubbing one's toe. All that really matters for utilitarianism are changes in utility that moral agents can bring about, and choosing a reference state simply makes it easier to determine what those changes in utility will be. There is no such thing as "absolute utility"—since "fulfillment of a preference" only makes sense relative to some other way things might have been—but I think the definition that I used comes pretty close to any intuitive sense of absolute utility that people might hold.

Definition: *Aggregated utility*, U , is the sum total of the utility in the universe over all of time, where utility is counted equally regardless of which organism happens to experience it or when it happens to be experienced. The unit is the *util*.

$$U \equiv \int_{-\infty}^{\infty} \left(\sum_{s(t) \in S(t)} u(s(t)) \right) dt. \quad (2.1)$$

Definition: *Utilitarianism* is an ethical doctrine whose goal is to maximize U .

Definition: A *moral agent* is any thing that can think about and act upon moral ideas.

Axiom: Moral agents exist.

Definition: A *utilitarian* is a moral agent who adheres to utilitarianism.

She aims to effect the greatest possible change in U through her life's work, $\Delta U(L)$:

$$\Delta U(L) \equiv \int_{t=now}^{t=death} \frac{dU}{dt}(t) dt, \quad (2.2)$$

where $\frac{dU}{dt}(t) \equiv$ the change in U that the utilitarian brings about at t .²

Fundamental Ethical Axiom: Moral agents ought to be utilitarians.

Thus, in spite of my early aversion to Intuitional Ethics [...], I was forced to recognise the need of a fundamental ethical intuition.

The utilitarian method [...] could not, it seemed to me, be made coherent and harmonious without this fundamental intuition. [...]

And I had myself become, as I had to admit to myself, an Intuitionist to a certain extent. For the supreme rule of aiming at the general happiness, as I had come to see, must rest on a fundamental moral intuition, if I was to recognise it as binding at all. And in reading the writings of the earlier English Intuitionists, More and Clarke, I found the axiom I required for my Utilitarianism [That a rational agent is bound to aim at Universal Happiness], in one form or another, holding a prominent place [...].

²By this statement I mean that the change in U at t includes all future impacts that the utilitarian's action at t will eventually have, even those that take place after her death. If instead $\frac{dU}{dt}(t)$ were interpreted to mean just the change in U that actually occurs at t , then the upper bound of the integral would have to be infinity or, at least, the time at which the utilitarian's actions ceased to make any difference in the universe.

—Henry Sidgwick, *Methods of Ethics*, “Preface to the Sixth Edition,” 1901 [29]

2.3 Utilitarianism as Happiness Maximization

At its most basic level, utilitarianism is about preference fulfillment; that’s the core value which motivates the philosophy. But in this section, I pursue an alternate, and perhaps more intuitive, formulation: utilitarianism as maximization of positive emotions—or, for shorthand, “happiness.” What I am trying to do here is show that the precise definition of utility that I gave in section 2.2 coincides closely with what people generally think of as positive emotion.

Good Examples: A bunny comes across a small garden and looks inside. It sees row upon row of fresh lettuce. The bunny desires the lettuce with a strength of 10. It scoots under the fence, approaches the lettuce patch, and begins nibbling. In so doing, it has satisfied its preference and so derives 10 utils/minute of utility. But clearly, the bunny also derives happiness from the lettuce. How much happiness? Intuitively, a state that an organism prefers twice as much as another state should furnish twice as much happiness as the other state. So happiness should be proportional to preference strength. Well then, why not just say the bunny also experiences 10 utils/minute of happiness? Since I’m arguing that the conventional concept of happiness closely approximates the specific definition of utility, why not use the same unit for both?

A factory-farmed pig lies uncomfortably in an overcrowded truck on its way to slaughter. The pig’s leg was broken upon being shoved into the truck. In addition, the pig feels crushed by the others, sick from the journey, and anxious about being thrown into a completely new situation [31]. On account of these factors, the pig wishes that it were unconscious with a strength of 500; thus, the pig feels -500 utils/minute. Why does the pig wish it were unconscious? Because it wants to end the negative emotions that it feels. The strength of the pig’s preference comes from the strength of its emotions. Therefore, let’s say the pig feels -500 utils/minute of negative emotion.

There are some situations in which utility does not correspond perfectly with happiness—at least not according to the way that happiness is generally understood. In these cases, I think utilitarianism gives the correct moral evaluation of the situation.

Not-So-Good Examples: A monk denies himself food, sex, entertainment, and access to the outside world. He prefers living the way he does and so derives positive utility from his asceticism. Many people would probably not characterize his lifestyle as a pleasurable one, however.³

³As I said before, I’m using “happiness” and “pleasure” in the sense in which they’re usually understood. I could also have chosen to *define* happiness and pleasure as utility, in which case the definitions would

In an alternate universe, there exists a frog with one and only one wish: to be brutally tortured. This is not a passing desire of madness that disappears once the pain starts; rather, this frog wants the pain to continue for as long as possible. The people in the alternate universe consider the situation. If the frog's preference is strong enough, the utilitarians point out, it might be not only permissible but morally obligatory to torture the frog.

This last example seems troubling, if not viscerally revolting. "How can you even think of inflicting pain on another creature?" screams one's intuition. "That's horrible!" Yet, I do actually think that the utilitarian position in this case is the right one. It's hard to wrap one's mind around exactly how it would feel to be the frog. We—by virtue of being evolutionarily selected organisms—have such an ingrained reaction against pain that it's impossible for us to imagine wanting more. But this is a limitation of our capacity for empathy, not a genuine reason against torture. We oughtn't deny the frog its desire simply because we can't conceptualize its internal state of mind. Perhaps the best thing that the hypothetical humans could do would be to compare the torture in their minds with something that they could imagine desiring, such as extreme pleasure. That way, they could feel better about doing what really would be the right thing.

have never diverged. But then, I would have just pushed the problem back a step, for then, my definitions of pleasure and happiness wouldn't have corresponded to popularly understood ones.

Chapter 3

Essays on the Intuitive Basis for Utilitarianism

3.1 From Materialism to Utilitarianism

Assumption: Materialism can be regarded as “true” within the context of conventional human views of reality.

According to this assumption, all matter consists of groups of atoms that move around, recombine, and occasionally change form. Every once in a great while, atomic interactions happen in such a way that a group of atoms forms a self-replicating combination. Such incidents are not somehow “favored” by the universe any more than any other series of chemical reactions is favored; it is only because of the very nature of the combination itself that it “survives” when other combinations do not. As these combinations continue to reproduce, they inevitably mutate into a variety of different forms, and those that are best able to reproduce themselves are those that remain. Every once in a while, this process of adaptation leads to a more complex combination of atoms—not because complexity is an inevitable result of evolution but simply because a particular set of environmental factors at one particular time happens to favor it. By fortuity, the intricacy of the combinations will sometimes continue to expand and may eventually—in very rare circumstances—lead to groups of atoms that can cognize their own existence. This cognizance, along with any other thoughts, sentiments, and sensations that the atomic collection experiences, consists entirely of bunches of atoms moving and interacting in certain ways. Such is the materialist view.

To many people, this conception of the universe would lead to what I shall call

Materialistic nihilism: If all feelings and emotions can be reduced to atoms moving in certain ways, why do they matter? Who cares whether certain groups of atoms move in one way rather than another way?

In response to these discomfoting questions, most people turn to belief in God, a spirit, or some other form of meaning that transcends matter. I obviously cannot dismiss such beliefs as definitely false (speaking, of course, within the conventional human view of “reality” that considers it possible to define “truth” and “falsehood”) for the same reason that prevents me from dismissing a belief that little green aliens will colonize Earth within the next ten years: in both cases, I simply lack evidence to disprove the beliefs with certainty—however unlikely they may be. But in choosing between a model of the universe for which there is substantial evidence and one simply that cannot be disproved, I consider it sensible to select the former. And it is for this reason that I shall maintain the assumption set out in the first paragraph that materialism can be regarded as “true.” Does this not, then, lead back to nihilism? I shall argue that it does not: to me, a purely materialistic view of existence leads not to a rejection of concern for others but to a defense of it.

Let us first examine the perspective of materialistic nihilism. In doing so, it will be helpful to focus on one particular example of a sensation—say the pain of torture. The nihilist is indifferent to the suffering of the victim because she sees the experience of pain merely as the movement of certain atoms in the victim’s brain and nervous system. And if pain is nothing more than the movement of atoms in one particular way, what makes it any different from the movement of atoms in any other way, such as one that an organism experiences as pleasure? If the whole field of ethics concerns itself with making sure that a few atoms in the universe move in one way rather than another, what’s the point? For that matter, who cares about art, literature, music, or anything else, because all of these activities accomplish nothing other than moving atoms in certain ways?

The flaw that runs throughout this argument is the assumption that “it makes no difference whether groups of atoms move in one way or another.” Such a statement may seem valid *prima facie* because people tend to view atoms as abstract objects like rocks—and few people care how a pile of rocks is oriented. But some atoms are not just meaningless pieces of matter: those that enable a person to live, feel, and think are as valuable as the sensations themselves—indeed, the atoms of emotion are the sensations themselves. If a certain arrangement of atoms is a *sine qua non* for the experience of love, then that arrangement genuinely is love, and those tiny bits of matter should be viewed with the same reverence as we generally give to the abstract notion of love itself.

In the preceding paragraph, I attempted to demolish what I defined as the materialistic-nihilist argument, which is the equation of thoughts and sensations with specific movements of atoms in an attempt to transfer the putative dryness of the latter onto the former. It is possible, however, for the nihilist to go further, by denying not only that atoms as an abstraction can have value but that the emotions that they create do not matter either. I think not that this revised nihilist approach can be refuted in the sense of regular argument, for it reaches down to one of the most basic axioms of human existence: namely, that sensations of pain and pleasure, fear and comfort, frustration and satisfaction all do matter. Such an axiom cannot be “proved” to someone who refuses to accept it, inasmuch as the process of argument requires shared assumptions of value on both sides, and this

dispute is over the existence of value itself. I cannot prove that it is better not to burn my finger, break my arm, or otherwise needlessly inflict pain on myself, and I cannot prove that there is value in fulfilling my preferences. I simply take these feelings for granted as the singular ethical axiom on which moral reasoning rests.

If we accept the assumption, as many people do, that emotions have value, then we see that the atomic interactions that create sensations have the same value—for indeed, they are two sides of the same coin. So if my sensation of joy matters, then the atomic interactions that create it matter equally. From this materialistic standpoint, empathy follows quite naturally. I care about moving atoms in my own body in ways that I prefer, so why shouldn't I want to also move atoms in preferable ways in other bodies? If there is no higher power, no spirit, and no God, then the value of the universe derives wholly from these desirable atomic movements, and we ought to create as many of them as we can; it makes no difference in which collections of atoms the positive atomic movements take place.

3.2 Subjectivism

A person can view a subject from a number of angles and understand how all of them might be accepted; math illuminates the world differently than poetry, for example, and both approaches can be meaningful at the same time. With two or three people, the possible approaches to a subject are even greater. Humanity as a whole has contrived myriad ways of viewing reality, even though human beings are nearly identical to each other biologically, genetically, behaviorally, and environmentally. Consider, then, how many more possible conceptions could exist that humans either have not yet discovered or that the human mind is not even capable of imagining. The most fundamental philosophical ideas—existence, time, cause and effect, et cetera—might be incomprehensible to an alien creature capable of thought. And even with this comparison, we are speaking only of organisms that can “think”—a restriction that in and of itself favors a certain perspective.

So how is it, then, that one can claim the existence of an “objective” or “universal” truth? Doing so requires acceptance of one particular mode of reality and rejection of all others. Even if we considered it possible that one form of thinking can know “truth” while all of the rest languish in delusion, how do we know that it is our form of truth which is right? We have no basis for asserting that our basic philosophical view of the world is more likely to be correct than that of another person, another animal, or a hypothetical alien—for those beings could say the same about their perspective. Furthermore, the idea of “truth” itself assumes one particular perspective (or at least a small subset of the total), and even to use language and logic favors certain views of reality. The probability that these specific views of reality are “correct” is likely close to zero, in which case it is meaningless to speak of a “true” reality. But even to make such statements is to assume certain ideas, wherefore the entire process of reasoning ends in failure (a concept that is itself partial to

certain views).

As a result, I would turn to a more pragmatic point: Even if there is absolute truth, why should we care? What is wrong with muddling around in delusion?¹ Readers might retort that there is value in seeking truth, that objective knowledge is the purpose of human existence. In response, I would simply point out that the ideas of “value” and “purpose” are defined by the person holding them—that is to say, they are subjective in and of themselves. So the objectivist position that people should strive to uncover universal truth might actually be considered a subjective position in disguise. Any other “absolute” values or principles that one may hold are also subjective opinions.

Imagine that an asteroid hits a desolate, lifeless planet ten thousand galaxies away from us, and no being observes it. Whether the event actually happened is debatable (perhaps it’s not even clear what it would mean for a nonobservable event to “actually happen”). As we saw above, such a controversy cannot be resolved. But what is clear is that the occurrence of the event doesn’t matter, inasmuch as no being’s experience was in any way altered by it. Similarly, when a person holds a religious belief, the “objective truth” of the religion doesn’t actually matter so long as the person does not question it, because its effect on the thoughts, feelings, and actions of the believer are exactly the same whether it is “true” or a delusion. In certain cases, “thinking something to be true” is indistinguishable from “knowing the truth.”

Applying this idea to other examples yields highly untraditional conclusions. For example, consider the farmed humans in *The Matrix*. The thoughts and feelings that most people experienced in virtual reality were exactly the same as they would have been had those people existed in an identical “real” world. Along the same lines, whether or not $2 + 2$ actually equals 5 is irrelevant as long as the citizens of Oceania believe it, because the consequences are precisely the same either way. Readers may remonstrate that “truth” does actually matter, that it is inherently wrong for people to be deceived. But this “inherent value of reality” is itself a subjective position in the minds of those who hold it. If one does not know that one is delusional, one cannot feel this subjective pang, wherefore the “inherent value of reality” no longer matters—except, possibly, to outside observers.²

¹Readers at this point may be slightly confused by my stance. In “Foundations of Utilitarianism” (chapter 2), I devoted an entire section to establishing epistemological axioms so that I could eschew nihilism. But now I seem to be giving in to the legitimacy of the nihilist position. If I don’t care about objective truth, why did I establish those axioms?

I think the answer to that question is that my axioms were not concerned with absolute truth in the sense that I discuss here of “discovering the underlying structure of reality.” Rather, I was assuming the bare minimum needed to make such statements as “Why should we care?” even intelligible. Even just to ask “Why should we care?” assumes certain ideas (intelligibility of language, meaningfulness of the word *should*, etc.), and the position I take here relies on the same axioms as those that I already established.

²I acknowledge that I’m being a little aggressive in pushing for my position here. When I say that the “inherent value of reality” no longer matters, I must be holding some stance as to what it means for something to “matter” in the first place. The stance that I hold is my own: the utilitarian belief that things only matter to the extent that sentient organisms care about them. If someone else felt that “truth” was important even when nobody knew the difference, my statement would not apply.

One final objection must be considered: If it is acceptable to deceive people in certain hypothetical situations on the basis that “they do not know the difference,” why is it not also acceptable for politicians, scientists, and the media to lie in order to pacify the public? I respond firstly that fooling everyone all of the time would be very difficult for anything of major significance; of course it is possible, for example, to change the date on a letter without anyone knowing otherwise, but it is not so easy to hide massive suffering. And even if the government did succeed in deceiving everyone else—for example, if the government had a secret torture chamber about which no one else knew and whose existence no one else even suspected—the truth would still matter to the organisms being tortured, and it would, in fact, be better for the public to know what was going on. By the same token, the falsity of the Matrix would matter if there were some who knew that it was false, because the discrepancy between fiction and reality would cause a difference in the subjective feelings of those people. (Whether the uneasiness felt by Morpheus and Neo would outweigh the presumable increase in happiness for those inside the Matrix is another question, but the point is that their uneasiness would at least have to be taken into account.) In other words, the “inherent value of truth” only disappears completely when deception is universal.

The type of subjective consequentialism expounded above, when combined with empathy, leads naturally to utilitarianism. Just as subjectivists care only about the attitudes, thoughts, and emotions that individuals happen to have, utilitarians care only about the utility and disutility that individuals experience. Any other “objective” values that people might hold—such as truth, rationality, freedom, equality, and the preservation of nonsentient nature—only matter to the extent that people derive utility from their fulfillment and feel disutility at their abridgement. If, for example, people did not value freedom and preferred to be locked up in cages, such would be the utilitarian policy (assuming the absence of a third option that would provide more utility than either freedom or incarceration). The additional component of empathy is important only because it determines which individuals deserve to have their subjective emotions counted. For egoistic consequentialists who lack empathy, the only utility that needs to be calculated is that of oneself. For utilitarians, in contrast, empathy requires consideration of the emotions of all beings capable of feeling them.

“[T]here is nothing / either good or bad, but thinking makes it so [...].”
—Hamlet, *Hamlet* [28]

3.3 On the Origins of Speciesism

None of the below discussion is novel or original; all of the ideas mentioned have been explicated at length by many other authors, particularly Peter Singer [see 34, 31]. I wrote the following passage merely in order to develop the established argument for animal welfare with a slightly different emphasis.

Speciesism—discrimination on the basis of species—has been the dominant attitude of human beings ever since they developed thousands of years ago. Indeed, a disposition of apathy toward other animals is probably favorable from an evolutionary standpoint: those early human beings who cared for animals were probably not likely to survive easily, particularly before the advent of agriculture. Of course, that something is natural does not necessarily mean that it is ethical. It is natural for men to repress women, for the wealthy to exploit the poor, and for the strong to dominate the weak. But most people would agree that these practices are not moral. To quote Stephen Jay Gould, “The factual state of the world does not teach us how we, with our powers for good and evil, should alter or preserve it in the most ethical manner” [15]. (For more on this, see section 6.2, “Should Nature Be Inviolable?”) Thus, it is also possible that our current treatment of animals is not moral—however natural it may be.

Concern only for members of one’s own species seems generally favorable for survival in the wild. But speciesist attitudes have also been reinforced by prevailing cultural standards—particularly the Judeo-Christian ideas that God made man in His image and gave man dominion over the animals. Just as virtually every civilization throughout history has seen itself as the center of the universe, so too humanity has viewed itself as supreme and fundamentally different from every other living creature on the planet. It is this view that has caused humans throughout the centuries to care more about the welfare of their peers than their cows, and it is because of this view that we, for example, shudder at Abraham’s determination to kill Isaac while thinking nothing of his actual slaughter of the lamb.³

This Judeo-Christian view of humanity was slowly eroded by science—notably the discovery by Copernicus and Galileo that earth was not the center of the universe—before being completely shattered by Charles Darwin, whose 1871 *Descent of Man* proclaimed the following:

The main conclusion here arrived at, and now held by many naturalists who are well competent to form a sound judgment is that man is descended from some less highly organised form. The grounds upon which this conclusion rests will never be shaken, for the close similarity between man and the lower animals in embryonic development, as well as in innumerable points of structure and constitution, both of high and of the most trifling importance—the rudiments which he retains, and the abnormal reversions to which he is occasionally liable—are facts which cannot be disputed. [...] He who is not content to look, like a savage, at the phenomena of nature as disconnected, cannot any longer believe that man is the work of a separate act of creation. He will be forced to

³As Singer and I recognize, there may be morally relevant differences between killing a lamb and killing an acutely self-aware human being. The point of this statement was not to suggest that killing a lamb is necessarily as bad as killing a person but merely to convey the point that in hearing of the slaughter of the lamb, many of us give little consideration whatsoever to the suffering that it would experience.

admit that the close resemblance of the embryo of man to that, for instance, of a dog—the construction of his skull, limbs, and whole frame on the same plan with that of other mammals, independently of the uses to which the parts may be put—the occasional re-appearance of various structures, for instance of several muscles, which man does not normally possess, but which are common to the *Quadrupedia*—and a crowd of analogous facts—all point in the plainest manner to the conclusion that man is the co-descendant with other mammals of a common progenitor. [8]

In other words, *Homo sapiens* are simply biological entities like any other organism; they may possess unique capabilities—particularly writing, farming, and manufacturing, not to mention having the most intricate language of any species—but they are not, at a basic level, different from other advanced animals. Distinguishing members of the human species from members of the chimpanzee species represents not a fundamental division but merely one more arbitrary way to classify various organisms. And while this separation may be useful, it is not necessarily morally relevant. In the same way, it is helpful in many ways to distinguish between males and females, but such a distinction does not necessarily have moral significance.

This is where concern for the welfare of other species takes hold. When we stop viewing humans as members of one species and pigs as members of another, but instead view them both as biological entities fulfilling their own needs and desires, we remove the inveterate conceptual roadblock that inhibits tantamount concern for the feelings of the pig. We can then concentrate on the ethically relevant characteristics of each biological unit—to wit, the capacity to experience fear, stress, and pain, and the capacity to feel joy, satisfaction, and fulfillment. In short, the Darwinian approach allows us to understand that division on the basis of species is but one possible way of viewing animals, and—furthermore—it is not a morally important one. We might similarly classify organisms based on their ages, their colors, or the speeds at which their hearts beat; it is because we are so accustomed species-based division that we consider it more than one arbitrary way to arrange biological entities.

Are we to extend our concern to all the beings capable of pleasure and pain whose feelings are affected by our conduct? or are we to confine our view to human happiness? The former view is the one adopted by Bentham and Mill, and (I believe) by the Utilitarian school generally: and is obviously most in accordance with the universality that is characteristic of their principle. It is the Good Universal, interpreted and defined as ‘happiness’ or ‘pleasure,’ at which a Utilitarian considers it his duty to aim: and it seems arbitrary and unreasonable to exclude from the end, as so conceived, any pleasure of any sentient being.

—Henry Sidgwick, *Methods of Ethics*, Book IV, Chapter 1, 1907 [29]

Chapter 4

Essays on Metaethics

4.1 Intuition and Reason

In math and other fields of rational inquiry there are two degrees of understanding. The first is intuitive knowledge. Whether because we experience something directly or because we conceptualize it through example, intuitive facts just feel right. For example, it seems obvious that the angles opposite the equal sides of an isosceles triangle will be the same. At times, concepts are not immediately intuitive but can become so through analogy, such as the realization that the total resistance of two wires connected in parallel is *less than* the resistance of each wire individually because one can draw a comparison to the amount of water that can flow simultaneously through two pipes rather than one. The second mode of understanding is formal proof. We may intuitively know that all of the angles of an equilateral triangle are sixty degrees, but we also must prove it to be sure. Math and science abound with examples in which the common-sense answer is utterly wrong: heliocentrism, relativity, and quantum mechanics, just to name a few. Complete understanding of something generally requires both degrees of thought.

The same process of understanding applies to morality. Ethical questions should not be solely decided on the basis of pure intuition or on the basis of pure reason.¹

The basis of morality must always be intuition: this grounds and gives meaning to any reasoning that one hopes to undertake. For example, one must intuitively establish that “suffering is bad” before the rational realization that “this action will reduce suffering” has any moral significance. To use the terminology of argument, one must establish a *warrant* before an audience will make the leap from one’s *reason* to one’s *stated conclusion*. At the foundational level, such warrants can only arise from intuitive emotion.

In the same way, mathematics can only arise from certain core axioms that one takes for

¹Indeed, it’s hard to even imagine what it would look like for something to be decided purely on the basis of reason. At some point, one needs to establish why things matter, and it’s hard to imagine how this could be done without emotion.

granted. Like basic intuitions in ethics, these fundamental postulates cannot be debated because doing so would require that both sides share some set of even deeper assumptions, which is not the case once one has gone down to the foundation of the theory. In the same way that two people can disagree about the core values of ethics, two people may take for granted different core axioms in math and arrive at two different theories (take, for example, Euclidean and non-Euclidean geometry).

Once basic assumptions have been established, however, it is best to let reason take its course, both in mathematics and ethics. While intuitive mathematical conclusions do often turn out to be correct by logical proof, there are some instances in which that is not the case; sometimes theorems that are clearly proved from basic axioms turn out to be counterintuitive. When it comes time to apply mathematics in practical situations, we trust knowledge derived from logic, not “what feels right.”

I assert that we should do the same in ethics. Beyond the most basic, axiomatic expressions of value, pure intuition becomes fuzzy, capricious, and self-contradictory [see 32]. It is here that reason must take hold and guide our emotional energies toward those decisions that genuinely are most in line with our fundamental intuitions.

4.2 The Role of Unreality in Moral Reasoning

The field of ethics is often supposed to be a sort of empirical science—or at least it appears so from the way philosophers develop arguments. Many ethicists begin by examining a variety of intuitive moral reactions to assorted situations and, treating those observations as “facts,” develop theoretical models to “fit the data.” I have the utmost respect for this type of empiricism in natural- and social-science realms. Indeed, it is only through evidence that humans acquire useful knowledge about the world.

Yet, even in science, there is a difference between believing everything that one sees and employing prudent use of reason to understand underlying mechanisms. When a pencil is stuck into a glass of water, we do not automatically presume that it has actually broken at an angle. Instead, we are able to explain this example of light refraction through our basic understanding of physics. In the same way, we should not take for granted an intuitive moral judgment of a complex situation simply because we feel it; rather, we should evaluate that intuition within the framework of our foundational moral principles. Thereby, we can decide whether our visceral reaction made sense or whether it was illogical (like the bent pencil) given our underlying ethical values.

If we reject the intuition-based empirical approach to ethics except at the core level of establishing fundamental moral beliefs, then we can see that we don’t need to have an actual, real-life moral dilemma before we can contemplate ethical ideas. In statistical regression analysis, one ought to beware of “extrapolation” (calculating points outside the domain of input values on which one collected results). But if ethical judgments are made *a priori*—as I am arguing they should be—then this restriction does not apply, for we

create the regression curve ourselves.²

To put it more concretely, contemplation of complex moral situations serves not the purpose of collecting data but, rather, the purpose of applying one's reasoning skills and refining one's thinking on a point. Often, imagining thought experiments can force us to challenge existing notions that we would otherwise not be willing to question. A nice example hereof might be Peter Singer's famous argument that if self-awareness and mental understanding are to determine the wrongness of inflicting pain, then we ought to support testing of new chemicals on severely intellectually disabled adults who lack significant family ties; after all, they would provide far more reliable data than lab animals, and we have to do something with them anyway [31]. The abhorrence wherewith most people (including Singer and myself) meet this proposition serves a purpose that might not otherwise have been accomplished: compelling the opponent of ethical consideration of animals to critically examine his preconceived belief.³

But must we only to utilize real-world situations? I assert that we do not. Since the purpose of considering examples is to challenge and exercise our minds, what difference does it make whether the situation actually exists? Our fundamental moral values should apply just as well to situations that, had the universe created itself differently, might have been true. Granted, an exception would have to be made for a supposition so completely and abstractly foreign that regular ethical argument would not apply, but by and large, most fantasied conditions that people might contrive should work appropriately.

I shall give an example of an instance in which the previous point would apply. Very often in debates over veganism, the argument comes up that, because meat consumption is natural, it must be acceptable. I counter the warrant upon which this conclusion is drawn—that what is natural must be acceptable—with the assumption-challenging point that male domination of women and the forceful rule of the strong over the weak are also natural, but they are not automatically “right.” Generally, this suffices to settle the matter, but a particularly persistent contrarian might respond that repression of one person by another has not been demonstrably shown to be genetically predisposed, while omnivorism certainly has been (one need look only at our sharp front teeth). It is here that the thesis of this essay becomes relevant. One might acknowledge that current scientific evidence has not irrefragably confirmed the genetic inevitability of repression, but what difference makes

²Strictly speaking, this does not have to be true. One might decide to create a regression curve that's defined on only a subset of all of the values on which it could be defined. In the same way, even an *a priori* moral view could intentionally limit the set of circumstances to which it applies. I'm just suggesting that doing so is less natural than it is when moral theories are constructed empirically.

³“But wait a second,” skeptical readers might reply. “Here's an example where you're taking your intuitive reaction to a given situation and adjusting your regression curve accordingly. Isn't that the same thing as what the ‘empiricists’ are guilty of?” I suppose this is a good point. I *am* adjusting my theory on the basis of a new datum. There is, however, a difference between my approach and the strictly empirical one. The latter, upon discovering a new data point, merely adds a bend to an existing curve. The former, in contrast, treats the new data point not as something to fit into the old curve but, rather, an opportunity to reexamine the slope and intercept of its line (*i.e.*, the original core axioms).

that? It is conceivable that such a determination could be made (one work that argues such a position is [14]), and in that case, what would be evolutionarily natural would not be moral. Why need the scientific fact actually be shown before the theoretical point can be accepted? General philosophical questions, *ipso facto*, rely not on detailed technical information; situations which one might imagine should be just as helpful in challenging one's mind as those that scientists currently believe to be veridical.

4.3 Morality as a Stimulus-Response System

[Certain] reactions, brought upon by evolution, instruct the user to avoid things that drain utility and come into more frequent contact with those that provide it. With the reaction to the inputs, the system can then make conclusions which are tested with more data, what we have termed experience. [23]

I agree with the above quote. One can picture this phenomenon easily in simple organisms: for example, one might picture unicellular creatures that avoid adverse stimuli and seek attractive stimuli. I should note that these actions do not automatically indicate the presence of “pain” and “pleasure” sensations as we know them, for that would require some sort of nervous system. The responses of plants and unicellular animals can best be considered as reflexes whereof the organism is unaware. However, the sensations of pain and pleasure in more complex animals are only different forms of the same process, the principal distinction being that the organisms in question are cognizant of adverse sensations and thence can consciously choose means to escape them, instead of relying on instantaneous reflex.⁴ Especially in the case of birds and mammals, “adverse” and “attractive” stimuli become very complex, as they may involve intricate social relationships and intellectually derived (not just corporally derived) forms of pleasure—but these things still count as stimuli.

Morality is not “on a higher plane of existence” than observation of the physical world; rather, science and morality are two distinct mental operations that the stimulus-response units that we call human beings undertake (along with any other organisms that are moral agents). Both science and morality require observation of the world and application of deductive logic. The difference is that science is descriptive and morality is prescriptive. Morality does not transcend nature; indeed, it has no meaning without some physical reality for it to consider. Rather, it merely processes the material world (what “is”) and decides what to do about it (what “ought” to be done).

To put the concept crudely and in a light favorable to the utilitarian, observation of the world (which, in advanced forms, is thought of as science) is the stimulus, and it leads to a response in the form of a desire for action on the basis of that observation.

⁴I should add that some behaviors of complex animals can still be considered reflex. For instance, upon touching a hot surface, we reflexively take our hand away even before our body produces the sensation of pain.

As a simple example, consider a mouse that becomes stuck in a sticky mousetrap. In noting the situation, the mouse sends a stimulus to its brain, which then produces a response in the form of an urge to escape (which can lead the mouse to struggle, whine, and even chew off its own limbs) [9]. The mouse’s impulse to free itself—however basic or uncalculated it may be—constitutes its morality at that point in time.⁵ Human morality, I assert, is just a more intricate and empathetic form of the mouse’s impulse to escape. We, too, observe the world and feel an impulse to act somehow in response. Granted, most of us care about other people (whereas the mouse at that moment is only concerned with its own situation), but this simply means that the pain of others is, to some extent, internalized so as to evoke a response within us as if we ourselves were in pain. Human morality, then, consists of those responses that we have to adverse and attractive stimuli in ourselves and in others.

I said that this description of ethics favors the utilitarian because what I just expounded is the utilitarian basis for morality, that is, following the natural impulses to reduce adverse stimuli and promote attractive stimuli to the greatest extent possible. Ethicists who hold an alternate moral perspective might remonstrate at my characterization of human morality.

4.4 Utilitarianism as a Resolution to Ethical Disagreements

How ought one to resolve seemingly irreconcilable ethical disagreements? For instance, suppose that Person A supports needless infliction of pain, while Person B opposes it. Beliefs such as “pain is good” or “pain is bad” are, to many people, the most basic principles from which rational argument derives meaning; hence, the principles themselves are unarguable. How can ethical disputes be resolved once they have reached the level of undebatable sentiments?

One approach that appeals to many people is to claim objectivity in ethics. “There are some things,” they say, “that are absolutely right and absolutely wrong. It is absolutely wrong to inflict needless pain, wherefore I am correct and you are immoral.” Two problems arise from this, however. The first, obviously, is that either side can claim to be on the side of objectivity (in much the same way that both sides of a religious war can claim to be on the side of God). The second objection arises from the concept of materialistic subjectivism described in section 3.1. Living organisms are nothing more than fortuitous combinations of atoms that have organized themselves in such a way as to be able to cognize various aspects of their environments. Any emotions, beliefs, and principles that such organisms hold are products of the ways in which their atoms are arranged. It thence seems very unlikely (although not impossible) that the perceptions of these beings could represent

⁵I contend that all sentient organisms possess this “aversion to suffering” form of morality and, hence, that all sentient organisms hold legitimate moral positions. Many philosophers, though, would probably object to my characterization of ethics in this way, claiming that morality encompasses only those prescriptive thoughts made by rational, self-aware creatures who have taken a universalizing point of view. For more on this, see section 4.5, “Do Implicit Preferences Matter?”

anything “absolute” or “universal.” Emotions, tastes, and fundamental moral values are simply *created* by those combinations of atoms that are capable of having them. In this way, materialism leads to subjectivism.

So the universal-moral-absolutes approach is unsatisfactory, and we’re back where we started: How do we resolve a conflict between two subjective moral values held by two different groups of atoms? I shall propose one method, but—as it turns out—it is simply a restatement of utilitarianism.

All sentient beings (except maybe nihilists) hold fundamental ethical values. If the organism is a normal developed human or (potentially) an advanced animal, it may express some of its moral premises in thoughts and words (*i.e.*, explicitly), while others will be implicit in its behavior. If the being lacks the capacity for words, the sensations that it experiences can be considered by those who do possess thoughts and language to constitute an equally valid expression of moral opinion. To use the terminology of chapter 2, I am merely asserting here that some organisms that lack abstract thought are still sentient (that is, capable of having preferences). Readers objecting to this claim should see section 4.5, “Do Implicit Preferences Matter?”

Now imagine a universe in which just one of those sentient creatures exists. Because the being is the only combination of atoms capable of conscious experience, it is the only organism for whom anything can matter. So the ethical standpoints that this being holds are the only ones that exist and, hence, are the only ones that ought to be fulfilled. In this single-being universe, “moral actions” are those actions and only those actions that conform to the being’s ethical framework.

What happens if we add a second sentient creature with different notions of morality, some of which are mutually exclusive with those of the first creature? A given action could now be moral and immoral at the same time, depending on the perspective. How do we reconcile this antinomy? I suggest that we consider the conviction with which each creature holds its own belief; whichever position is held more fervently is the one that ought to be followed. But how can we say that it is better to appease one standard of morality than another, since, to each organism, the standard is absolute? Recall that any given moral standard only starts to matter once there exists a being that holds it. So if one being holds its ethical belief more passionately, then in some sense, that belief “exists more.” As a result, accordance with that belief matters more, too.

The same principle, it is readily seen, applies to any number of beings. Imagine a three-organism universe in which each creature maintains the belief that it holds with equal ardor. Suppose further that two of the organisms agree with each other but disagree with the third. Then, their belief “exists” twice as much—and thence matters twice as much. Extension of these ideas to a universe full of sentient beings will, it is readily apparent, lead to maximization of preference fulfillment—in other words, utilitarianism.⁶

⁶Not everyone will accept this proposal for dealing with divergent ethical principles, but even those who reject it will at least have their sentiments and beliefs counted by it!

The stance of this essay is well summed up in the following passage by Peter Singer:

I am not defending the objectivity of ethics in the traditional sense. Ethical truths are not written into the fabric of the universe: to that extent the subjectivist is correct. If there were no beings with desires or preferences of any kind, nothing would be of value, and ethics would lack all content. On the other hand, once there are beings with desires, there are values that are not only the subjective values of each individual being. The possibility of being led, by reasoning, to the point of view of the universe provides as much “objectivity” as there can be. [35]

I should address one objection that readers might have.

Critical Readers: “Earlier you said that ‘the universal-moral-absolutes approach is unsatisfactory.’ But now you’re saying that there is an objective way—utilitarianism—to resolve ethical disagreements. Isn’t that being inconsistent?”

I acknowledge that this criticism is valid. Indeed, I am taking an absolute ethical position, in violation of my previous statement. By holding my Fundamental Ethical Axiom (section 2.2), I do maintain an absolute moral stance. In some sense, I am just one more person who thinks he’s right and everyone else is wrong. How do I justify this?

Consider the arguments that I put forward earlier in this essay in contradiction to a universal morality. The second of them was that, in view of materialism, it seems unlikely that the universe would have an underlying objective structure that would be consistent with any moral position. However, as Singer pointed out, this objection does not apply to utilitarianism because that philosophy does not require “Ethical truths [to be . . .] written into the fabric of the universe.” It just requires that there be sentient organisms with meaningful subjective experiences.

The first objection that I gave is more difficult: Everyone claims that her own moral theory is the right one. How can you claim to be on the side of objectivity when everyone else says the same thing?

In response to this fact, a person might logically decide to adopt one of three positions:

1. Nihilism: There’s no such thing as objectivity in ethics. It’s wrong for Person A to tell Person B that B is acting immorally. But indeed, even the statement that I’m making right now is an absolute moral stance: I’m asserting that one person shouldn’t tell another what to do, but in so doing, I’m telling you what to do. So even my position is meaningless. In fact, everything is meaningless. Even the words with which I say this are unintelligible.

2. Pure Relativism: I can’t say that Person A’s belief is better or worse than Person B’s. A is right, but B is also right; it just depends on the perspective.

3. Absolute Relativism: From Person A’s perspective, A’s belief is right. From Person B’s perspective, B is right. The overall “right” answer is some combination of the perspectives of A and B. Morality is relative, yes, but this relativist position is absolute. All moral beliefs are relative except for this one.

Utilitarianism is a form of absolute relativism. Specifically, it is the form for which “combination of the perspectives of A and B” takes the form of aggregation weighted in proportion to the relative strengths of the beliefs.

Of course, I still need an axiom to establish utilitarianism, and I still am holding an absolute position. Intuitively, though, I find that utilitarianism is somehow preferable to other absolute moral beliefs because it “includes” all of them. But that’s just my opinion.

4.5 Do Implicit Preferences Matter?

Utilitarianism prescribes maximizing fulfillment of preferences, where “preferences” is defined to include both explicit and implicit desires. But why should we care about implicit desires? If an organism can’t think of its preference in abstract terms, why should we count it? This essay presents an intuitive answer to that question.

Definition: *Explicit-preference utilitarianism is the form of utilitarianism in which only explicit preferences are counted. My own position—that both explicit and implicit preferences matter—might be referred to by the ungainly term *explicit-and-implicit-preference utilitarianism*.*

At one time, I argued in favor of explicit-preference utilitarianism, although I didn’t think in terms of this framework at the time. I didn’t really hold the belief very passionately, but I did at least make reference to it in one essay that I wrote in November 2003:

The human species is capable of a degree of self-awareness [and abstract thought] far exceeding that of any other creature on earth. [...]

It is, moreover, that same human awareness that creates and defines morality [*i.e.*, explicit preferences]. The idea that a person is of more value than any other animal is not a universal truth but, rather, a human moral standard. The anthropocentricity of my belief [in the acceptability of slaughtering chickens] is therefore inconsequential, for it is humans who decide the morality of actions. [Why do only humans define the morality of actions? Because morality consists in fulfillment of explicit preferences, and (according to my ignorant view of animal capabilities) only humans have explicit preferences.] If it is generally considered acceptable to kill animals for meat, then to do so is not “wrong.” At the same time, commonly held notions of what is “right” forbid the murder of a person, almost without exception. Not everyone agrees, of course, on the

guidelines and boundaries of morality, but the most passionate beliefs of the greatest number of people affected by a decision should generally militate most heavily on the final course of action; it is my perception that the slaughter of chickens for food conforms to this standard.

Note that the method I prescribed in the final sentence for deciding what moral position should be adopted—namely, considering “the most passionate beliefs of the greatest number [...] affected by a decision”—is the same as the utilitarian approach outlined in section 4.4. The essential difference was that, in the above passage, I only included explicit preferences in my assessment.

According to explicit-preference utilitarianism, animals—including human babies and severely intellectually disabled adults—have nothing to be counted in the utilitarian ethics calculation. This does not mean, of course, that these animals will necessarily be treated poorly: if most of those whose preferences *are* counted feel concern for those beings, then it will be moral to take the interests of those beings into account. In the present day, that is largely true for human infants, severely intellectually disabled human adults, and pets; for virtually all other animals, it is not.

In response to explicit-preference utilitarianism, I give an intuitive appeal. It is unfair to limit the set of organisms whose preferences count to only those that have abstract reasoning ability. Why? Well, why do we care about preferences at all? What’s the intuitive emotion that underlies the Fundamental Ethical Axiom? It’s hard to articulate exactly what the emotion is that motivates the axiom, but it’s something like this: Organisms are combinations of atoms constructed in such a way that they want certain things to happen. Preferences are sort of an intangible “will” that a combination of atoms feels. I care about preferences because I am one of those combinations of atoms, and I know what it’s like to have a “will” for the universe to be in one state instead of another. That “will” does not take form only once I speak or write words to express it; rather, it’s always there, as part of my being. It is the intangible will that matters; words are just tools for communicating information about the will. Preferences don’t matter because they’ve been made explicit. Preferences matter when they’re implicit. That is, explicit preferences only matter *because* implicit preferences matter.

Here’s another intuitive way of looking at the situation. The decision as to which organisms should have their preferences counted is analogous to the decision of which organisms ought to be able to vote.⁷ Except when discrimination takes place, restrictions on suffrage arise from necessity: it is simply not possible for a frog to vote. In the same way, one might want only to consider explicit preferences because those are the only types of preferences that one can hear or read about. Only organisms that have voices understandable

⁷Suffrage is not exactly the same as utilitarian aggregation of preferences, forasmuch as the former counts only the “location” of an organism’s desire while the latter also considers the strength of the desire. But this distinction is unimportant to the present example, which concerns itself only with the question of which creatures are counted in aggregating preferences.

to humans are able to participate in the discussion, so to speak.

As a matter of practical necessity, we exclude children and future generations from being able to vote. But that doesn't mean it's *fair* to do so, because those people are or will be just as affected by politics as voting-age adults. In response to this unfairness, adults can choose to correct the situation to some degree by giving surrogate representation to "those who cannot speak for themselves." This approach is far from perfect in its application; indeed, one of the most common complaints against democracy is that voters care only about their own generation and burden posterity with massive fiscal and ecological debts. But this is all the more reason for people to strive to correct the system's flaws.

The problem with explicit-preference utilitarianism is that it confuses what is what is convenient—considering only preferences that can be put into words—with what is right—considering all preferences that exist.

Chapter 5

Methods of Utilitarianism

5.1 General Procedure for Utilitarian Evaluations

In the following exposition, I frequently make use of example ΔU values for the sake of clarity. In reality, though, it is impossible to assign a precise numerical value in utils to the consequences of an action. Sometimes numbers given in more concrete units can help to approximate magnitudes of different effects—for example, one might say that refraining from the purchase of a store-bought chicken prevents around six weeks of suffering in a factory farm, among other things [18]. And in some situations, no numbers are necessary at all. Thus, utilitarianism is not just an abstract nice idea; often, it can be approximated in practice.

The procedure for doing so is rather elementary. One simply estimates the net ΔU that would result from each possible action that one might undertake. While utilitarianism implies that the choice of one's action must be made from among all possible options, one can usually simplify the consideration to just two or three of them. Similarly, while it would be theoretically desirable to know the ΔU for every tiny effect that one's action might have, it is generally sufficient to consider the most significant impacts.¹

Resource Equalization

Estimation of ΔU for each possible option is the first step, but one must also consider the potential ΔU s that one will give up in the process of undertaking each action. All actions require some amount of resources (even if only a few seconds of time), and in using resources to undertake Action A, one forgoes resources that might have been applied to

¹This is not to say that one should never consider the other possible choices for action or the smaller effects of any given action; indeed, there may be actions which one has never before contemplated that would prove better than the rest, and there may be seemingly negligible details that turn out to be important. One ought to devote some time to exploring these possibilities. However, doing so for every decision would be a waste of time that could be better spent on other pursuits.

Action B. The easiest way to factor this into one’s deliberations is to consider options for which resource use has been equalized. To take an example, suppose that Utilitarian C is unsure whether he should walk to a nearby protest or stay at home and write a letter to Congress. The former option would take five hours and would result in 60 utils, while the latter would take one hour and produce 15 utils. Naively, Utilitarian C might say, “Oh, well if I go to the protest, I can effect a bigger ΔU . I ought to do that.” But such an analysis is flawed because the two options being considered utilize different amounts of resources. Assuming that time is the only resource involved, Utilitarian E can set equal the resource use of the two possibilities by imagining that he wrote five letters in five hours. Now the utility comparison for the two applications of equal resources are 60 utils for the protest and 75 utils for the letters. In this case, using only the options and numbers given, Utilitarian C should write to Congress.

Marginal Aggregated Utility

Resource equalization is generally the easiest way to do back-of-the-envelope comparisons. However, one can also factor in resource use with the following approach.

Definition: *Marginal aggregated utility is the change in aggregated utility that results from application of a a unit of resources.* If r represents the amount of resources used, marginal aggregated utility $\equiv \frac{dU}{dr}$.

Again naively, Utilitarian C might think, “Marginal aggregated utility is the bang that I get for the buck, so to speak. So clearly, I should always choose the option for which marginal aggregated utility is highest.”

In general, this thinking is correct. Consider, for instance, the protest-versus-letters scenario above. Assuming that $\frac{dU}{dr}$ is constant at least over small values of r , $\frac{dU}{dr}$ for the protest = (60 utils)/(5 hours) = 12 utils/hour, while $\frac{dU}{dr}$ for writing letters = (15 utils)/(1 hour) = 15 utils/hour. Choosing the option with higher $\frac{dU}{dr}$ does indeed give the right answer.

But consider this situation. Suppose that, in rummaging around an old stack of papers, one finds an already-written advocacy letter to McDonald’s that will effect an expected value of 370 utils. However, the letter still needs a 39-cent stamp in order to be sent. Let r be in units of cents. Then, $\frac{dU}{dr}(1) = 0$; that is, the marginal net utility of the first cent used is nothing. The same is true, of course, for the next 37 cents: $\frac{dU}{dr}(r) = 0$ when $0 \leq r \leq 38$.² But $\frac{dU}{dr}(39) = 370$ utils/cent. Presumably, one ought to spend all 39 cents to achieve this ΔU , but one would not reach such a conclusion simply by looking at the initial value of $\frac{dU}{dr}$.

²Perhaps this is bad notation. I’m using $\frac{dU}{dr}$ to represent the rate of change of a discrete variable: number of cents. However, in general, other components of r like time *are* continuous. I didn’t want to add confusion here by changing notation to a form more appropriate for discrete rates of change.

In fact, this result makes perfect sense. Recall from equation (2.2) that the utilitarian aims to maximize the integral of marginal aggregated utility. There, I wrote marginal net utility with respect to time, but one could equivalently write the equation thus:

$$\Delta U(L) \equiv \int_{r=0}^{r=last} \frac{dU}{dr}(r)dr, \quad (5.1)$$

where “ $r = 0$ ” signifies the beginning of resource use, and “ $r = last$ ” signifies use of the last resource that one has available.

Looking at this integral all of the time might be cumbersome, though. A shortcut that works in most situations is to observe that if marginal aggregated utility remains constant or decreases over all values of r , then it is sufficient in maximizing ΔU to choose the option whose immediate marginal aggregated utility is highest. Here’s a formal statement of that idea.

Definitions: A naive utilitarian is at $r = a$ and is deciding how to apply his next unit of resource. He must choose among a finite number n of options, denoted $1, 2, \dots, n$. Since n depends on r , I use the notation $n(a)$ to represent the number of options available at $r = a$. Let

$$\left(\frac{dU}{dr}\right)_m(a) \equiv \max\left\{\left(\frac{dU}{dr}\right)_1(a), \left(\frac{dU}{dr}\right)_2(a), \dots, \left(\frac{dU}{dr}\right)_{n(a)}(a)\right\}. \quad (5.2)$$

The *naive utilitarian approach* is to choose option m in every instance of a future resource application, from $r = 0$ until $r = last$. Call N the state of fact in which the naive utilitarian approach would actually succeed in maximizing ΔU .

Rule:

$$\forall r \in (0, last) \forall i \in \{1, 2, \dots, n(r)\}, \left(\frac{d^2U}{dr^2}\right)_i(r) \leq 0 \implies N. \quad (5.3)$$

Resource Use

Finally, I shall make two clarifying points on resource use.

- What if different options require utilization of different types of resources? For instance, how does one compare donating money to a wilderness-preservation organization and spending time handing out literature on factory farms? Fortunately, resources can be converted into one another given the specifics of the situation. In the example, one acquires the money that one wants to donate by using time. Assuming that the most efficient way to convert time into money is to work at one’s job

(say the effective pay rate after taxes is \$10/hour), then one can compare spending two hours handing out brochures against two hours of working at one’s job, the latter being equivalent to donating \$20 to a wilderness-preservation organization. This explains why I was able to use a single variable r to encompass all different types of resources.

- What if application of a resource results not just in a change in aggregated utility but also the creation of more resources? For instance, time spent exercising not only enhances one’s own immediate utility; it also extends one’s life expectancy and thereby “creates” more time. As before, one can interconvert units. If exercising for 30 minutes furnishes 5 utils of immediate personal aggregated utility and adds 45 minutes to one’s life, then one simply estimates the ΔU that those extra 45 minutes will produce—say 10,000 utils—and takes the sum: 10,005 utils for 30 minutes. Assuming linear relationships between variables, the marginal net utility of exercise would be 333.5 utils/minute.

5.2 Affecting Public Opinion

Introduction

There is much that utilitarians can accomplish by themselves or with other utilitarians. But sometimes it may be helpful to change the way other people see things, too. This essay identifies the three ways in which changing popular attitudes can increase ΔU .

Definition: *Individual utility, $i(t)$* , is the utility that a single organism experiences at t .³

Definition: *Individual aggregated utility, I* , is the sum total of utility that a single individual experiences over its lifetime:

$$I \equiv \int_{t=birth}^{t=death} i(t)dt. \tag{5.4}$$

Definition: $t = 0 \iff$ the present time.

Definition: $i(t = 0) \equiv i^0 \equiv$ immediate individual utility.

³This is the same thing as $u(t)$ for an individual organism, but I use a separate letter to avoid confusion.

Definition: A *maxi organism*⁴ is one who, when she makes a deliberate decision,⁵ always does what will maximize her i^0 .

Example: Maxi Person A is considering whether to go jogging or biking. On this particular day, i^0 for jogging is 15 utils/minute, while i^0 for biking is only 10 utils/minute. A will go jogging.

Assumption: All sentient organisms that make decisions are maxi organisms. That is, every decision-making sentient organism, when it makes deliberate decisions, does the action for which it has greatest immediate individual preference.

An Objection to This Assumption: “What about altruistic decisions in which Sentient Organism B sacrifices for the benefit of Sentient Organism C? For example, imagine that Person B decides to adopt a vegan diet in order to avoid causing harm to Chicken C and countless other animals. In so doing, B denies herself the short-term pleasure of eating chicken nuggets. Is this not a case in which B deliberately chooses an option that will entail a lower i^0 ?”

Response: Consider, though, that there is a reason for which B made her decision: taking this action provided her the satisfaction of knowing that she did not cause terrible suffering to another organism—which is to say, taking this action gave her positive utility. My assumption that all sentient, decision-making organisms are maxi simply means in this case that, in order for B to choose tofu over chicken nuggets, it must be true that i^0 for tofu is higher than i^0 for chicken nuggets (as well as higher than i^0 for any third options). This can happen if B feels enough empathy and concern for Chicken C.

This discussion should not be construed as predicating that humans (and sentient, decision-making organisms) are incorrigibly egoistic and that they cannot have genuinely selfless concerns. Rather, in showing that altruistic action occurs only when it furnishes maximum immediate individual utility, I intend to illustrate the mechanism by which genuine selflessness is expressed. That beings act to maximize i^0 is merely a reflection on the internal decision-making process within an organism’s mind. Pure altruism can still exist; it merely takes form of altering the organism’s perceptions of i^0 for various options.

This leads us to the first way in which utilitarians can increase ΔU by influencing others.

⁴“maxi” is short for “maximum i^0 .”

⁵This clause is important. When a doctor taps your knee, you don’t have a reflex reaction because you think that doing so will maximize fulfillment of your immediate preferences. Similarly, when a person gets stage fright and is unable to give a speech, the person doesn’t think that he will thereby maximize his immediate individual utility, even though getting stage fright is an action that he performs. Organisms are maxi only with regard to actions that they decide to take.

Utilitarian Approach 1: Aligning Individual Utility with Social Utility

Continuation of the Previous Example: Call G the state in which B gives up chicken nuggets and $\neg G$ the state in which she buys the nuggets. Suppose that the immediate individual utility that B derives from buying the nuggets $i_B^0(\neg G) = 3$ utils/minute. If she doesn't buy the nuggets, B will feel mildly disappointed: $i_B^0(G) = -2$ utils/minute. Assume that each of these levels of utility would stay constant for 2 minutes. Then, given only the numbers so far, $\Delta I_B(G) = (-2 \text{ utils/minute})(2 \text{ minutes}) = -4$ utils, while $\Delta I_B(\neg G) = (3 \text{ utils/minute})(2 \text{ minutes}) = 6$ utils.

Now suppose the following values for the the effect of B's decision on Chicken C: $\Delta I_C(G) = 0$ utils and $\Delta I_C(\neg G) = -1,000,000$ utils. Ignoring other associated changes in utility (such as the environmental benefits G), $\Delta U(G) = \Delta I_B + \Delta I_C = -4 + 0 = -4$ utils, and $\Delta U(\neg G) = 6 + -1,000,000 = -999,994$ utils. Clearly B ought to choose G because $\Delta U(G) \gg \Delta U(\neg G)$ (and presumably $\Delta U(G) > \Delta U(X)$ for all other options X). However, my assumption that B is maxi implies that she will only choose G when $i_B^0(G) > i_B^0(\neg G)$ (and, of course, $i_B^0(G) > i_B^0(X)$ for all other X).

What I haven't considered so far is the additional component of $i_B^0(G)$ that B might feel due to moral satisfaction at not causing suffering to Chicken C. Say that B cares a little bit about Chicken C, so that she derives, say, 3 utils/minute from knowing that she did the right thing. Combining this component of $i_B^0(G)$ with the frustration component that I mentioned previously, $i_B^0(G) = -2 + 3 = 1$ util/minute. Unfortunately, that's still less than the 3 utils/minute that she derives from buying the chicken nuggets, so B doesn't take the utilitarian action.

Now suppose that B cares a lot about Chicken C, so that the empathetic component of $i_B^0(G)$ is 20 utils/minute. Now, total $i_B^0(G) = -2 + 20 = 18$ utils/minute > 3 utils/minute $= i_B^0(\neg G)$, and B *does* refrain from eating chicken nuggets.

This example illustrates the first approach that utilitarians can take toward others: fostering empathy. Utilitarians can increase ΔU by making individuals happiest when they do things that make others happiest.

A Coercive Example In some cases, it may not be possible or effective to persuade people to do *proprio motu* what is optimal for society as a whole. In these cases, the best option may be a coercive one. Coercive utilitarianism works in the opposite way from its voluntary counterpart. Instead of increasing the utility that an individual derives from choosing the best option, coercion reduces the utility that the individual would derive from any alternate option.

One example is taxation of the rich to pay for environmental programs. Suppose that Rich Person D experiences -50 utils/minute if she pays \$1 million in income taxes but would feel 0 utils/minute if she successfully evaded those taxes. The social benefit resulting from \$1 million of extra tax revenue might be 100,000,000 utils, making the net social benefit of taxing D exceedingly high. But right now, D would experience more individual

utility by not paying taxes (0 utils/minute) than by paying taxes (-50 utils/minute). So if the government cannot persuade D on the basis of altruism, it must resort to coercion: namely, threatening to bring charges serious enough that D would feel more negative utility from not paying his or her taxes (say, -200 utils/minute) than from paying them (still -50 utils/minute). Thus, paying the \$1 million has become the option of maximum individual utility—even though that utility is negative.⁶

Economics Parallel: Encouraging people to consider the impacts of their actions on all sentient beings affected rather than merely on their own personal happiness has a standard name in economics: *internalizing the externalities*. The process simply involves changing the individual cost-benefit calculation that a decision maker faces in such a way that maximization of the individual’s profit also leads to maximization of social profit. This can be accomplished in two ways: (1) subsidizing decision makers to encourage them to take or not take a certain action and (2) taxing decision makers to encourage them to take or not take a certain action. The parallels to voluntary and coercive utilitarianism are obvious.

Utilitarian Approach 2: Aligning Immediate Utility with Long-Term Utility

A Second Objection to My Assumption: “Your assumption that all sentient, decision-making organisms are maxi means that all organisms maximize *immediate* individual utility. How, then, do you explain the fact that people plan ahead for the future? How do you explain allocation of resources toward investment? And how do you explain the fact that people don’t always engage in reckless behaviors?”

Response: This situation is very similar to the previous one. Just as it’s possible for organisms to incorporate empathy into their i^0 , so too it’s possible for them to incorporate concern for their own futures into it.⁷ In the preceding section, I outlined the mechanisms behind voluntary and coercive approaches to correcting selfishness; the same principles can also serve to minimize personal shortsightedness. Sometimes the most effectual option is to change people’s perceptions of the worth of long-term utility so that they will derive

⁶**Reader:** “How is it possible for i_D^0 to be negative? If $i_D^0 = 0$ utils/minute corresponds to the desire for experiencing the state of unconsciousness, wouldn’t unconsciousness be the option of maximum i^0 ?”

Response: It’s true that $i_D^0 = 0$ utils/minute for experiencing the state of unconsciousness, but this is different from the i_D^0 that D would actually have for *becoming unconscious*. $i_D^0 = 0$ describes one’s desire for the emotional sensations (or, in fact, lack of sensations) of unconsciousness, but one’s desire to actually become unconscious could be very different. In D’s case, unconsciousness would take the form of suicide or drugs, and D would find neither of those desirable because she cares about her future. That is, i_D^0 for actually becoming unconscious would be very negative, far more so than -50 utils/minute.

⁷Indeed, concern for one’s future self is in some sense a specific form of concern for other organisms, if one imagines one’s future selves as being a number of unique organisms that are different from one’s current self.

more short-term utility from choosing the foresighted option (think, *e.g.*, of DWI-impact panels, which, incidentally, may also accentuate the value that people place on the utility of others). And in other cases, coercion will be more effective in forcing people to choose the long-term decision by reducing the utility of more immediately gratifying options (think, *e.g.*, of the seat-belt law). Or sometimes a combination of both would work best.

Economics Parallel: Unlike efforts to address selfishness, utilitarian policies aimed at curbing myopia have no associated economics term of which I am aware.⁸ However, one might invent a parallel phrase, such as *immediatizing the eventualities*.

Utilitarian Approach 3: Aligning Ostensible Utility with Actual Utility

Even if utilitarians succeed in convincing people to care about others and their future selves, there's still a potentially significant factor preventing people from achieving the highest possible ΔU : ignorance.

Examples: Person E engages in a mystical ritual, thinking that it will bring about heaven on earth. In reality, he accomplishes nothing other than wasting time.

Person F hears about the potential of a biological terrorist attack. In response, he goes to the hardware store, buys twenty rolls of duct tape, and seals off his house.

People follow the dictates of superstition and put duct tape over their windows because, within the scope of the limited knowledge that they have, doing so is the optimal decision. If you genuinely believe that terrorists will strike any minute and that duct tape is your best protection against them, then you will actually think that you are maximizing your own utility by sealing off your house. When you learn how absurd your fears and protective response were, you stop putting duct tape over your windows because you no longer think that doing so will maximize your utility.

Reader Comment: “But even that assessment—that duct-taping one’s house is a waste of time—is itself potentially erroneous. Indeed, any factual notions that one might hold about the world are probably, to a greater or lesser degree, incorrect. You can never know for sure that what you think will maximize ΔU actually will.”

Response: Indeed, that’s correct. We can never have perfectly precise knowledge. The best that we can do is to learn enough about a topic that we can be reasonably confident that what we hope will maximize ΔU will, in actuality, get pretty close. In the same way,

⁸The reason for this is that economists misguidedly tend to characterize people as *rational*, which in my language, would be like saying that people are *maxI*—*i.e.*, all of their deliberate decisions maximize I , not i^0 . Immediatizing the eventualities is actually just an attempt to make people more rational.

scientists admit that they will never attain perfect understanding of the natural world, but with greater study, they can develop a more accurate picture of it. However, to recognize that incomplete knowledge will always interfere with actual maximization of aggregated utility is not to conclude that there are not varying degrees of ignorance. Education about sex, exercise, nutrition, factory farming, and the economic and environmental effects of government policy can vastly improve society—simply by changing people’s perceptions of which actions will really make them and others happiest. Thus, the purpose of what I shall call *educational utilitarianism* is to enhance the accuracy of people’s knowledge in ways that will augment the actual utility that results from following their own perceived utility.

Summary

This discussion explored three ways in which utilitarians might set about improving society beyond their own personal actions:

- (1) voluntary and coercive approaches to reducing adverse impacts of egoism,
- (2) voluntary and coercive approaches to addressing the perniciousness of myopia, and
- (3) education.

5.3 On the Precautionary Principle

This essay critiques the basic notion behind the precautionary principle. I acknowledge that the precise definition of the principle may indeed consider the points that I make; if that is the case, then this criticism is not directed against technical application of the actual principle but, rather, against popularized and oversimplified conceptions of it.

In looking through Google’s list of definitions for “precautionary principle,” I found a few that state the idea in the way in which I have always heard it:

“Better safe than sorry” attitude. The idea that, in the face of uncertainty, society should assume that potential problems are real and address them accordingly. [12]

Assumption of the worst-case scenario with respect to actions whose outcomes are uncertain. [13]

In other words, when one is uncertain how harmful something will be, one assumes the worst until science proves otherwise.

This idea may sound very nice initially, but consider a few cases that one might encounter in practice.

Example: Suppose that Monsanto has just developed two new pesticides, Pesticide A and Pesticide B. Each has so far been subjected to only one preliminary study, so their

actual toxicities remain uncertain. Yet, comparing only the preliminary results, it appears that A is very likely to be carcinogenic, while B shows no such signs. Following the simple precautionary approach given above, we must assume that both chemicals are seriously carcinogenic, because that is the worst-case scenario and their actual deleteriousness is not yet certain. But this would mean that regulatory agencies ought to spend just as many resources controlling and limiting the use of B as they do A and that environmental organizations should direct just as many resources toward suppression of B as A.

Example: In 1985, EPA characterized the chemical dioxin as “the most potent carcinogen ever tested in laboratory animals” [7]. One source of dioxin is the bleaching of paper with chlorine. It is true that confirmedly safe alternatives for bleaching paper have been discovered, but imagine that they had not. Suppose someone develops a new chemical that proves as effective and inexpensive in bleaching paper as chlorine. This chemical, too, creates bleaching byproducts, but they are entirely different from dioxins. Before any studies have been performed, the precautionary principle requires assuming the worst possible toxicity—perhaps one as detrimental as that of dioxin. But in all probability, the new chemical byproduct will not be quite so baleful. It makes sense, then, for paper mills to replace chlorine with the new bleaching agent, even before years of extensive study take place. Yet the precautionary principle as outlined above rejects this action.

As these examples show, the problem with the simplified precautionary principle lies in its inordinately absolute presumption of harmfulness, even when such harmfulness is not likely. Certainly we should not wait until adverse consequences are proven to take action—to that extent, the precautionary principle is right. But it can be just as shortsighted to assume great harm as it is to assume none at all.

The best approach, therefore, is to calculate an approximate expected value of deleteriousness.⁹ Granted, coming up with probability values is an inexact and somewhat arbitrary process,¹⁰ but what is the alternative? It is even more arbitrary to assume one particular result, as the simplified precautionary principle does. It is ultimately more helpful to say that A has a 0.9 chance of being carcinogenic while B has a 0.4 chance than to assume that they both have a 1.0 chance until proven otherwise, and doing so enables society to

⁹If H is a discrete random variable for harmfulness, and h_i represents the magnitude of harmfulness in the i th possible case, then the expected harmfulness H^e is given by

$$H^e \equiv \sum_i h_i P(H = h_i). \quad (5.5)$$

If H is a continuous random variable with density function $f(h)$ defined over harmfulness values h ,

$$H^e \equiv \int_{-\infty}^{\infty} hf(h)dh. \quad (5.6)$$

¹⁰Indeed, the frequentist school of probability rejects assignment of such values even in principle [11].

conserve its resources for those dangers that genuinely are most threatening.

Does this critique of the precautionary principle mean that preventive efforts are wasteful? Are environmentalists deluding themselves with a faulty doctrine that merely fosters irrational fear? In most cases, the answer is no. In fact, my modified form of precaution actually confirms most of the measures that environmentalists exhort us to take; it only does so on a sounder basis.

Example:

High-energy particle accelerators, used by physicists to investigate the fundamental laws of nature, could produce particles that create hyperdense “strange matter” that in turn might attract nearby nuclei, thus growing larger and attracting ever more nuclei, until the entire planet is compressed into a sphere no more than 100 meters in diameter. [...] The official risk-assessment team for one of these accelerators, at the Brookhaven National Laboratory, offered a series of estimates, one of which puts the annual risk of a disaster at one in five million. [33]

A precautionary-principle adherent would assume the worst-case scenario—that life on earth will come to an end—and conclude that such particle accelerators should be stopped. But given the enormous cost involved—not only death to current organisms but prevention of life for all future organisms—any sensible risk assessor using the expected-value approach would come to the same conclusion. Even though the probability of disaster is tiny, its expected value is enormous.

The same logic applies to almost all basic precautionary actions. For example, the risk of a car crash on any given car ride is quite small, yet it nonetheless remains a good idea to wear a seat belt. The reason is that the benefit of doing so—even when multiplied by a very low probability—is still so extraordinarily high that it outweighs any minimal costs that taking the time to buckle the belt might entail. One does not have to assume that the worst will happen in order for the potential consequences of the worst outcome still to militate with greatest power on a cost-benefit calculation.

Chapter 6

Intuitive Discussion of Further Topics

6.1 On Triage

Utilitarianism emphasizes efficiency: Any action that fails to effectuate the maximum possible ΔU is considered immoral. But isn't this standard for use of one's time and money too harsh? As long as one is making some sort of positive impact on the world, isn't that enough? Can't one live ethically without devoting all of one's energy and attention to helping others?

Perhaps an illustrative example will serve to weaken any intuitive sympathy that readers may hold for the above position. Imagine a military doctor who comes across a battlefield laden with hundreds of injured soldiers in severe pain. The doctor calls for assistance, but the additional medical units will not arrive for thirty minutes. However, the doctor happens to have with him a bag of pain medicine that he can use to palliate the suffering around him. Would it be acceptable for him to treat five of the soldiers and then stop to read a comic book, arguing that he has produced some positive ΔU and he needn't spend all of his effort helping others? Similarly, would we countenance his decision to spend most of his limited supply of pain killer on the mildly injured patients nearest to him, even though many of those a bit farther away are in absolute agony? I believe that the answers ought to be "no." Rather, triage—giving greatest medical attention to those who can be helped most in the least amount of time—represents the ethical imperative under these circumstances.

Yet how are other situations any different? In choosing how to spend one's time, what to do with one's money, what to pursue in one's career, and how to devote one's life, we are making the same choice as the doctor wondering whether to treat suffering patients or read a comic book; the only difference is that the consequences of the latter option are not so immediate and tangible. Similarly, decisions as to which efforts to pursue and which

actions to undertake are tantamount to the choice that the doctor faces regarding which patients to treat. Those who say, “I realize that this undertaking will not help as many sentient organisms as much as possible, but at least I’m doing something,” are in effectively the same position as the doctor who treats only those mildly injured patients nearest to him, because he is “at least doing something.”

But isn’t this focus on efficiency cold-hearted? In your pursuit of utilitarian triage, aren’t you neglecting the pain of those who don’t get preferential treatment? As an illustration of this objection, I recall one scene from *And the Band Played On*—a film about the first years of medical encounter with HIV—in which one character scoffed at a legitimate concern about a large pecuniary expenditure, saying that it was cold-hearted to care about money when lives were at stake.

In the real world, though, we can’t do everything. Resources are limited, and we inevitably face choices between helping one being or another. Questions of war and peace are not the only life-and-death decisions that a government makes; decisions of money allocation also decide who lives and who dies, who is helped and who is not. Indeed, the yearly federal budget kills vastly more people by those things on which it fails to spend money than the wars and Iraq and Afghanistan killed directly. In this real world, we must recognize scarcity and make tough decisions. It is precisely the sympathy that we feel for those animals left homeless by a hurricane that so fervently motivates us to devote our money and toward a more efficient way of ameliorating animal suffering: promotion of veganism. The deep concern that we feel over a new development in a rare ecological site near our home is precisely the source of our dedication to the amelioration of those environmental challenges in which we have the greatest potential for long-range impact, such as reducing overpopulation. Triage is not an act of harshness; rather, it represents the highest form of mercy and compassion.

6.2 Should Nature Be Inviolable?

Among philosophers and the general public alike, it is common to find a belief that nature ought to be inviolate. The view stems from several sources. One is probably the religious notion that “nature” represents the deliberate creation of one or more divine beings; since God or the gods created the Earth in this way, we should not alter it. Indeed, this is the basis on which many evangelical Christians are now beginning to demand action on global climate change and other ecological issues. It is interesting that a similar belief in nature’s irrefrangibility exists among secular environmentalists. This impulse, like the religious one, may arise from spiritual inspiration at nature’s wonder, but it may also arise from the scientific understanding that when humans try to intervene in nature, we generally end up doing more harm than good. Thus, many people make the leap from the empirical conclusion that humans rarely succeed at challenging nature’s course to the prescriptive belief that challenging nature is inherently wrong.

In this essay, I shall not attempt to challenge the religious and spiritual notions presented above. They represent fundamental intuitive premises of their respective belief systems, wherefore I cannot challenge them by means of argument. All which I can say is that I do not happen to share those core ethical principles. However, the final point that I mentioned—the “ecological” basis for nature’s sanctity—can be disputed. In fact, it’s rather easy to do so. The argument confuses innate inviolability with sensible case-by-case restraint: that is, it claims that because humans have not done well at challenging nature’s laws most of the time, nature must be innately irrefrangible. Such an argument is akin to the assertion that because wars have generally done more harm than good, war is wrong *ipso facto*. This is, of course, an illogical leap: war is wrong because we anticipate that it will have adverse consequences similar to those of past wars, not because war in and of itself is immoral. Similarly, the fact that humans should usually refrain from changing nature arises from our expectation that inimical consequences would result from future perturbations, not from an inherent principle that nature is sacred *per se*.

The sanctity of nature comes up often in ethical disputes, particularly with respect to the ethical status of animals. “Nature is harsh,” say opponents of animal welfare. “In the wild, animals eat each other, and those organisms that are born disabled or that become sick simply die in solitary misery. But that’s the way things are supposed to be. It would be arrogant/unfaithful/irreverent to think that human beings should try to change nature.”

I shall endeavor to challenge this rationalization on the following ground. Nature was not designed by an intelligent creator. It has no purpose or goal. Rather, it represents the accumulation of incidental occurrences: namely, the origination of self-replicating combinations of atoms and their subsequent expansion and diversification. Eventually, sentience emerged in those animals for which it bestowed evolutionary advantage, and in the process, things suddenly mattered for those organisms. The development of sentience just happened to occur, but because it did, the animals possessing it now had value. Yet nature took no heed of that fact and continued on as before. (Indeed, it is meaningless to talk of an incidental collection of occurrences having intention, and I do so merely for expository effect.) Nature does not care what it does to beings with feelings; if it were evolutionarily favorable for organisms to suffer excruciation for the entire duration of their lives (which surely is a reality for some wild animals), that would come about. The state of suffering or happiness in the natural world just happened, so it isn’t surprising that there is a great amount of pain.

The evolution of human beings (and perhaps other advanced animals that display empathy) was thus a very interesting occurrence, for now there existed groups of organisms capable of recognizing suffering in other animals and capable of doing something purposeful about it to a greater extent than ever before. To resist the natural (that is, incidental) course of things is thus to take advantage of the miraculous opportunity that we human beings happen to have for changing the world not according to capricious chance but in ways favorable to those beings for whom things matter—sentient organisms. It may turn out, of course, that people can best reduce suffering in nature by leaving it alone—and our

experiences with ecological systems have usually confirmed this—but such a conclusion comes from a particular evaluation of outcomes given the assumption that that humans ought to intervene or refrain from intervention in nature so as to maximize ΔU ; it does not arise from the inviolability of natural occurrences themselves.

6.3 Is Happiness Relative? A Personal Conjecture

For much of my life, I held the belief that happiness is relative, meaning that regardless of the conditions of one's life, one will experience a certain amount of joy and a certain amount of displeasure. I expounded this idea best in a book review of Winona LaDuke's *Last Standing Woman* that I wrote in 2003:

[T]he Ojibwe [Native Americans] “understood that both the making and the unmaking were essential parts of life and necessary to keep the balance. After all, what was dawn without dusk and what was life without constant change? [...] For all the pain and heartache we have felt, there has been and will be, an equal amount of joy. That is how everything works. There is always a struggle to maintain the balance” (187, 299). This statement explores ideas that I have also found to be true. The second sentence explains that life would be monotonous and meaningless if our surroundings and situations were not changing all the time. This applies even to moments of felicity, for if humans only experienced happiness (“dawn”) we would have no other emotion for comparison, and we would not be truly joyful. Changes are required for us to appreciate life's best moments, or else none of our experiences would be special, unique, or important. This relates to the ideas in the last three sentences of the above quote, which explain that, because the world has a balance, all of our emotions equal out, and we can never have more of one than another. When moments of elation subside, we are met with less pleasant emotions, just as feelings of desolation are always followed by periods of exultation.

One might, with levity, call this observation the

Law of emotion: For every feeling there is an equal and opposite feeling.

Comparisons between people in affluent nations and those in the Third world may to some degree support this claim. It is often commented that peasants and villagers in Latin America and Africa live joyfully in spite of their hardscrabble situation. At the same time, middle- and upper-class Americans, with all of their luxuries and conveniences, are very often depressed, aggravated, and discontent. While primarily an argument against relying on consumerism for happiness, this generalized comparison may also indicate that the material conditions of one's life will not, in the long term, affect the quality of one's life.

The law of emotion presents a problem for utilitarians: Why should we worry about economic justice if a tenant farmer in India is just as happy as a CEO in Japan? Why bother providing health care and clean drinking water if people who are sick will experience just as much happiness as those who are not? Why should we concern ourselves with the confinement of 10 billion animals each year [16] in stressful, unsanitary, and cramped factory farms if the animals will adjust to the harsh conditions? In short, the question of whether we should work to improve the quality of lives rests on the extent to which happiness is relative.

First, I acknowledge that happiness is relative *to some extent*, a conclusion with which few would probably argue. Certainly it would be more unpleasant for a rich corporate executive to live for a day in the life of a malaria victim than for the malaria victim to live for one more day, simply because the latter is far more used to the pain. And it is also true that the presence of health feels much better right after one has recovered from sickness than it does after being healthy for a long time. Obviously, emotions change from time to time, and diversity is essential for happiness (which is, of course, a restatement of the law of diminishing marginal utility).

If one were to draw a frequency distribution of happiness, I suspect that the mean value for most people’s lives would be positive,¹ because if asked whether they were glad that they had the chance to live their lives, most people would probably answer affirmatively. A frequency distribution of happiness is useful because it acknowledges the diversity of emotional states while simultaneously giving a single number—the mean value—with which to compare two different options.

The debate over whether happiness is relative is not about whether emotional states vary or whether this variation is important for making certain moments special. All sides acknowledge those points. What it is about is whether the average amount of happiness that two organisms experience can have different values. To me, the answer seems clearly to be “yes.” It is clear that some people carry themselves with habitual pleasure, while others experience more-subdued emotions. This situation might result from a difference in attitude or even simply from a difference in brain composition that causes the former person to produce more “happy chemicals” than the latter. Consider, as well, clinical depression and anhedonia: these are not merely short periods of sadness from which one rebounds with an equal amount of felicity later on. And even the most basic applications of the law of emotion might be brought into question. For instance, I don’t think that the relief of convalescence is as great as the discomfort of sickness; if it were, I would have no reason to wash my hands or cover my mouth when sneezing. The same seems to apply in cases of pain, anger, frustration, and other negative emotions.²

¹For farm animals, the mean value would probably be very negative [31].

²This is not to suggest, of course, that one should not remember one’s own hardships and the suffering of others in order to appreciate one’s health and comfort. I merely intend to demonstrate that the values of an emotion and its opposite emotion need not be equal.

Critical Reader: “Yes, I acknowledge the above examples and recognize that two people can, in fact, have different mean values for their happiness distributions. But you have still not demonstrated a connection between environmental conditions and emotion. Of course two people can enjoy life to different extents, but this does not mean that the average happiness values for two populations of disparate material standards of living will differ. People who live in poverty are used to poverty: that circumstance becomes their baseline happiness value from which variation spreads on either side. In the same way, affluent people are accustomed to affluence, so their happiness values spread out from either side of that baseline. You have not demonstrated that the locations of these baseline values—which, in a symmetrically distributed frequency distribution, are the mean happiness values—are necessarily different.”

Response: I shall attempt to adduce whatever limited experience I have with this subject in demonstrating that these baseline values are indeed different and that the conditions which utilitarians seek to address do in fact militate to lower these values. First of all, consider severe depression and hopelessness. I hope that I succeeded above in demonstrating that these conditions do result in an absolute lowering of average happiness. So if certain policies increase the prevalence of depression and hopelessness, then at least from this standpoint, the policies ought to be changed. In general, I think that this charge can be leveled conservatively against torture and perhaps more arguably against disease, hunger, and general poverty. Suicide, as well, appears to be a strong indication that people’s lives are absolutely worse than normal; indeed, in order to wish to die, one must be enduring significant periods of negative happiness. So I think it fair to say, for example, that the hardships of Indian farmers brought about by Monsanto and similar agribusiness corporations, which have induced thousands of suicides, have made agricultural life worse in an absolute sense.

Finally, I shall appeal to Maslow’s hierarchy of needs. There may be certain elements of life that are absolutely essential to its enjoyment, beyond which all additional material factors are luxuries that will have no definite effect on average happiness. Consider, *e.g.*, the need for food. From my own personal experience, hunger is not only discomforting because it is less fun than satiety (*i.e.*, relative to satiety, it has a lower happiness value, but if one had never experienced true satiety, then the presence of hunger would be normal and would hence not have a lower happiness value). Rather, hunger also seems to have an absolute effect on the body’s ability to experience happiness, perhaps because it makes harder the production of “happy chemicals” and the execution of other biological processes that create enjoyment. The same applies to sickness, tiredness, and any other state of poor health. Other basic biological functions may also fit into this category. Peter Singer’s *Animal Liberation* describes the ways in which factory-farmed chickens—who have never had access to the outdoors in their entire life—nevertheless exhibit strong instinctual impulses to “walk around, scratch the ground, bathe in the dust, build nests, [. . . and] scratch their wings” [31, page 113]. Thus, even beings that have never directly experienced or even indirectly heard

of a better life can still feel frustration that is relative merely to fundamental physiological needs. In other words, it may be *biologically impossible* for happiness to be relative, at least until one's quality of life rises above a certain basic material level sufficient to maintain health.

6.4 Why Determinism Doesn't Matter

It is a common belief that the existence of “free will” (that is, the ability to make decisions that have not already been preordained by past events) is necessary in order for life to be meaningful, in much the same way that ninety percent of Americans feel the need to believe in God despite the absence of scientific evidence. But much in the way that “a spirit” need not exist in order to enjoy life, free will need not exist in order to make the best decisions: in both cases, the reality of the situation is irrelevant beyond abstract philosophy, for the effective implications are the same whether or not a spirit exists and whether or not people have free will.

Assumption: Determinism is approximately true on the macroscopic level, where quantum-mechanical randomness can be ignored.

Accordingly, it is true that all of the events of the universe are theoretically predestined, but it is also true—and more importantly true—that such a fact does not matter. The reason is that humans could never possibly predict future events by examining all of the minute details that work together to cause them; doing so may be possible in principle, but it could never be accomplished—or even approximated—in practice. Therefore, in real life, we act in the same way that we would if free will existed. The theoretical presence of knowledge is irrelevant if people will never be able to learn it; it is as if the knowledge did not exist at all. Sure, a determinist will claim that the outcome of a coin toss is predestined, but inasmuch as no one could ever possibly predict it, we must act as if the outcome were truly random—and we do this by assigning probabilities of 0.5 to both heads and tails.

Some people argue that if determinism were true, they would not have responsibility for their decisions; moreover, they would have no incentive to make the right decisions because whatever they choose would have been selected ahead of time. Such reasoning ignores the obvious fact that, if the right decision is to be carried out, one still has to do it. What does it matter whether one decides to protect the environment out of free will or predestination? In either case, one still has to make it happen. The world's molecules may have been previously ordained to move in such a way that one makes the best decision, but one still has to move further molecules in order to carry it out. If humans could somehow know the future, and if such knowledge did not alter that future (that is, if “fate” in the conventional sense of the term were true), then the mindset presented above would be acceptable. But when we cannot ever discover those future events, the effect is the same as

if such knowledge did not exist at all. Thus, whether or not free will exists is a moot topic reserved for philosophy classes; it should have no impact on the way in which we choose to lead our lives.

What if the Assumption is False? That is, what if quantum-mechanical randomness cannot be ignored? Indeed, this has largely been the resort of philosophers and scientists in modern times who intuitively yearn for the veridicality of free will and who thence adduce quantum mechanics in an attempt to salvage it [10]. But does this really save free will, or does it merely qualify pure determinism? After all, consider what the quantum model proposes: our every action is not necessarily determined by past orientations of atoms and molecules but, rather, by pure chance (at least in those rare instances when randomness on the micro scale actually affects outcomes on the macro scale). Is that really any better? Instead of choosing to help a struggling chicken because I have been predestined to do so, I help the chicken because molecules in my brain happened to orient themselves in such a way that I made that decision. In either case, “free will” was not involved; the decision just happened due to factors—deterministic or quantum—that were out of my control. Thus, even if quantum mechanics can disprove the notion of pure determinism, it cannot automatically thereby prove “free will.” Doing so would still require a metaphysical leap which materialists are unprepared to make in the absence of evidence; that’s true in the case of determinism, and it remains equally true in the case of quantum mechanics.

The practical implications of a quantum reality are the same as those of a purely deterministic one: in both instances, prediction of the future by the interactions of elementary particles is impossible (indeed, in the former case even more absolutely than in the latter). Thus, we consequently act as if free will existed and go on with our lives.

6.5 Expectational Consequentialism

The philosophy of utilitarianism predicates itself upon consequentialism—the doctrine that the morality of an action is determined by the consequences that follow from the action, rather than from the intent of the actor or any other consideration. To take an example, imagine that Person A donates money to the American Enterprise Institute (AEI) out of a naive, uninformed belief that the AEI will improve the economy and thereby reduce poverty. To a (left-wing) consequentialist, that action is immoral—notwithstanding the altruistic intentions of the donator—because the actual results that it will produce are baleful.

Consequentialism may appear reasonable enough in that context, but there are other cases in which intent, rather than result, seems to determine our response. Consider, *e.g.*, the case of an accidental homicide. Imagine that Person B is a burglar climbing up the balconies that stick out from the side of a large apartment building. Person C, unwitting that B is hanging off the bottom of his balcony, walks out, steps on B’s fingers, and

causes B to lose her grip; B falls to her death. Our intuitive reaction is that C is not culpable in any way, and the law recognizes that C has not committed murder. Yet a consequentialist—in the conventional sense of the term, at least—argues that C’s action was still immoral because it caused someone else to die; had C not walked out onto the balcony, B would still be alive.

Is this, then, not an exception to consequentialism? Are there not certain cases in which the consequences of an action do not determine its morality? Here it will be helpful to delineate two types of consequentialism.

Definition: *Actual consequentialism* refers to the conventional notion of what consequentialism means: That the morality of an action depends upon what actual effects that action brought about.

Definition: *Expectational consequentialism*, in contradistinction, holds that it is the range of possible consequences that one should reasonably expect to occur—rather than the specific result that happens to occur—that determines the morality of an action. This, obviously, is the form of consequentialism expounded in the previous essay and the type under which utilitarianism falls.

But is this not a *deus-ex-machina* attempt to save consequentialism by adding elements of intent-based morality? I shall argue that it is not and that, indeed, expectational consequentialism is itself a long-sighted form of actual consequentialism.

To begin, consider the utilitarian purpose of punishment, whether in the form of fines, incarceration, or simply scolding someone for “making a bad decision.” Punishment cannot reverse the past, nor does it provide any comfort in the present—except in serious cases wherein trauma produces temporary sadism in those close to the victim. Indeed, the only immediate effect of punitive measures is to inflict disutility upon the offender. So how can punishment ever fall within the category of a utilitarian policy? The answer is that punishment (at least in ideal cases) acts as a deterrent against future commission of similar actions. Society does not jail people for murder because doing so reverses the damage that the crime caused but because doing so serves to influence the decisions of others in the future. And in affecting future behavior, the deterrent changes the future consequences that will result over the long run. The consequentialist argument for punishment, then, involves consideration of both immediate and eventual consequences, not just those that happen within the context of a particular crime.

We have established that the purpose of punishment is to affect future occurrences, not to change the present. So the decision of which actions are deserving of punishment rests upon evaluation of which actions are likely to cause harm in the future to such an extent that it is worthwhile for society to deter commission of those types of actions. For the sake of example, consider society’s decision to impose fines on speeding. It is impossible

to tell ahead of time which specific instances of speeding will result in a crash and which ones will be harmless, wherefore society has to punish them equally. Since the point is to prevent future incidents of speeding—not to punish those hapless drivers who actually did get into an accident—the law ought to be equally harsh on all speeders, because all of them did something equally immoral. When we have no practicable way of predicting specific outcomes, we treat all actions likely to produce Adverse Consequence D as equally worthy of preventing, notwithstanding some of them would not have resulted in D had they taken place. In other words, we can most successfully obviate actual future consequences by deterring those actions that are likely to be inimical, which is done by punishing commission of the actions likely to cause adverse results, not the adverse results themselves. Thus, expectational consequentialism—evaluation (and punishment) of actions on the basis of reasonably expected outcomes—genuinely derives from long-range consideration of actual (future) consequences.

In considering the accidental-homicide example from before, expectational consequentialism leads to a perfectly reasonable conclusion. Although stepping out onto the balcony happened, in that particular case, to cause a death, the probability of such a consequence resulting in the future is so minimal that society has no reason to deter people in the future from walking out onto their balconies. The action of stepping onto the balcony should not have reasonably been expected to kill someone—nor should it in the future, unless the number of incidents of burglars climbing apartments rises dramatically—wherefore (strange as it may sound), going onto the balcony was a moral action. After the fact, of course, we can realize that it would have been better had C not stepped onto the balcony, but C had no way of knowing this at the time. Given that C should not have been expected to know of B's presence, C is not only not culpable but was actually doing the right thing.

What, then, are we to make of the example offered in the first paragraph? Did not A genuinely expect at the time of her decision that her donation would produce positive consequences? Therefore, even if the actual result of her action was to hurt the poor, was her action still not moral because she expected that it would have positive consequences? Here we must make a distinction between any given expectation of effects and a reasonable expectation of effects. Any person can have an expectation of the future and act accordingly. Person E may believe that only by killing as many squirrels as he can will he prevent the Earth from being hit by an asteroid, but this belief would not automatically make right E's corresponding faunacidal campaign. The expectations upon which one bases one's decisions should be reasonable, and this is wherein science, research, and factual information enter the picture. A reasonable expectation of the future is the expectation toward which the bulk of independent, sincere research points. Inasmuch as virtually all relevant scientists agree that human emissions of greenhouse gases are accelerating global climate change, taking action to reduce those emissions would represent an ethical action on the basis of expectational consequentialism, even if it turned out (however unlikely it may be) that climate change was not occurring at all.

But what if even science is wrong? Can we consider scientific predictions to be “rea-

sonable expectations” considering how often scientists have erred throughout history? A reasonable expectation is not, by definition of the latter word, definite knowledge of what will happen; rather, it represents the best guess with which human knowledge can come up. If E were the only organism on Earth capable of evaluating expected consequences, then his decision would be based on the best information available; hence, he would be acting per reasonable expectations and, therefore, morally. However, because it was assumed that E could consult with astronomers about his concern—and thereby realize that his expectation was not reasonable after all—his action was not ethical according to utilitarian expectational consequentialism.

As a final note, it is conceivable that there could exist a Person F who is so deranged as to fail to understand that others might hold alternate expectations of the future and who, thence, fails to consult with anyone else before executing a perverse action. In this case, F’s behavior would be moral according to my definition, but society would be obligated to restrain F’s ability to carry out this behavior.

Chapter 7

Essays on Applications of Utilitarianism

7.1 Is Terrorism Prevention Cost-Effective?

How many lives were lost in the terrorist attacks of 11 September 2001? Answer: 2,976 [1]. Approximately how many people died of hunger on that same day, September 11, 2001? Answer: 24,000 [20].

Some believe that terrorism is the greatest threat to the security of the American people. Certainly the policies of the current administration would lead to that conclusion. Our interventions in both Afghanistan and Iraq were premised on fighting terrorism. We have been expected to accept the loss of some of our civil liberties in the form of the USA PATRIOT Act and to acquiesce to greater government secrecy in the name of “national security.” It is hard to find a speech by President Bush that does not in some way refer to his paramount national goal: the War on Terrorism.

And yet, when the number of consequent deaths is our standard of comparison, terrorism is of relatively minor concern.¹ Besides the 2,976 Americans killed on 11 September 2001, only four other Americans have been killed by terrorists within US borders since 1997 [21].²

Compare this to other causes of death in the United States. In 2000, tobacco killed 435,000 Americans [26]. In the same year, 400,000 Americans were killed by poor diets and lack of exercise. Alcohol killed 85,000 Americans in 2000, not counting accidents from drunk driving [26]. More Americans—3,247—died from drowning than from terrorism in 2001 [21].

¹Strictly speaking, the relevant consideration for policy should not be the absolute magnitude of a problem but the degree to which it can be ameliorated. In this essay, I make the assumption that the latter is at least roughly proportional to the former.

²I wrote this piece in April 2004. I’m not sure if this number has changed since.

The global situation is far more staggering. 8.8 million people worldwide die every year from hunger and hunger-related causes [20]. 6.1 million people are killed each year by treatable diseases like AIDS, malaria, and tuberculosis [30].

Given these statistics, what makes terrorism such an irrational obsession for so many Americans? For one thing, its unexpectedness and the understanding that it is a deliberate slaughter of life make it more terrifying than other causes of death. Moreover, terrorism is immediate, visible, and, hence, memorable; hunger, poverty, and environmental degradation are slow, silent, and too often forgotten. Mark Twain captured the concept well in this passage from *A Connecticut Yankee in King Arthur's Court*:

There were two "Reigns of Terror," if we would but remember it and consider it; the one wrought murder in hot passion, the other in heartless cold blood; the one lasted mere months, the other had lasted a thousand years; the one inflicted death upon ten thousand persons, the other upon a hundred millions; but our shudders are all for the "horrors" of the minor Terror, the momentary Terror, so to speak; whereas, what is the horror of swift death by the axe, compared with lifelong death from hunger, cold, insult, cruelty, and heart-break? What is swift death by lightning compared with death by slow fire at the stake? A city cemetery could contain the coffins filled by that brief Terror which we have all been so diligently taught to shiver at and mourn over; but all France could hardly contain the coffins filled by that older and real Terror—that unspeakably bitter and awful Terror which none of us has been taught to see in its vastness or pity as it deserves. [38]

Twain was, of course, referring to the Reign of Terror during the French Revolution of the late 1700s, but his insight into human nature—concerning our natural emotional tendency to focus on violence that is immediate and visible rather than slow and cumulative—remains true today.

This tendency may be natural, but it does not make for sound public policy. By focusing so much of our money and attention on the singular priority of terrorism, we neglect countless other areas of preventable mortality that are in much greater need of our efforts. The current administration hopes to spend at least \$50 billion every year on homeland security [21]. This is at the same time that an increase of just \$1 billion—two percent of that amount—in our spending on worldwide AIDS programs could prevent an additional 1.6 million people from contracting HIV/AIDS within the first year of the programs and could treat 400,000 people who would otherwise die within two years [5]. Programs like these have the potential to save far more lives than domestic antiterrorism efforts ever will.

The threat of terrorism is not nonexistent; it is real, and it does result in hundreds of tragic deaths every year around the world. But it must be put in perspective. If we continue to devote so much of our energy to terrorism—both as individuals and as a nation—we will

lose sight of those other areas of human need in which we have the potential to prevent untold more suffering.

7.2 On Popular Opposition to Utilitarian Policies

I shall endeavor to answer one of the common charges against the application of utilitarian principles: If the vast majority of society rejects a certain utilitarian policy, then how can one claim to be right in advancing it?

Definition: The *direct value* of Policy X, $\Delta U_d(X) \equiv$ the change in aggregated utility ΔU that Policy X creates when one ignores the opposition of majority sentiment.

Definition: The *indirect value* of Policy X, $\Delta U_i(X) \equiv$ the change in aggregated utility that results from popular reaction.

Identity:

$$\Delta U_d(X) + \Delta U_i(X) \equiv \Delta U(X). \quad (7.1)$$

Suppose there's an ostensibly utilitarian policy A for which $\Delta U_d(A) = 100$ utils. Assume that the next-best policy that could be carried out with the same resources is Neutral Policy B with $\Delta U(B) = \Delta U_d(B) = 60$ utils. Clearly Policy A—viewed purely from the standpoint of ΔU_d —is the best option. But implementation of the unpopular policy will create public outcry, perhaps so much that $\Delta U_i(A)$ will be -150 utils. The net effect, of course, is that $\Delta U(A) \equiv \Delta U_d(A) + \Delta U_i(A) = 100 - 150 = -50$ utils and should not be carried out; we would choose Neutral Policy B instead. Is this not the end of the discussion, forasmuch as $\Delta U(A) = -50$ utils $<$ 60 utils $= \Delta U(B)$?

Not quite. The reason is that one must not only consider the actual present ΔU of an action but also the potential future ΔU . Let's say that Utilitarian Group C conducts a public-awareness campaign to demonstrate to people that Policy A itself has a ΔU_d higher than that of all other options. If Group C succeeds in convincing enough members of the general public that Policy A in and of itself should happen, then the value of ΔU_i will have been changed. Because Group C cannot convince everyone, assume that $\Delta U_i(A)$ rises merely to 0 utils, a point where the intensity of outrage that remains is balanced out by intensity of support. The result, then, is that $\Delta U(A) = 100$ utils $>$ 60 utils $= \Delta U(B)$ and should thence be implemented. Through its efforts, Group C increased ΔU by 40 utils (from 60 to 100 utils).

Does that mean, then, that utilitarians should focus on changing public opinion? Not necessarily. Group C's public-awareness drive used a certain amount of resources in achieving a 40-util enhancement to society's happiness. The effect of the campaign must be viewed alongside all other possible resource allocations. Obviously, if using the same amount of

time and money on encouraging people to eat less meat would have increased total utility by 50 utils, Group C should not have conducted the campaign.

In other words, deciding whether to change public opinion in order that society will accept a utilitarian policy must be weighed as any other use of resources is weighed. If it will not be effective (perhaps because prevailing popular sentiment results from deeply engrained biological responses), it should probably not be carried out.

7.3 The All-or-Nothing Fallacy

People too often attempt to rationalize their own failure to take a certain utilitarian action with the following statement: “One more person will not make a difference. Sure, if 10,000 other people joined me, we could really have an impact. But the effect of one person is too insignificant to measure, so it doesn’t matter if I take this action or not.” This argument may be invoked in trying to justify inaction in two distinct types of situations; in both of them, it fails.

I shall call the first type of situation one of *continuously distributed effects* or, for short, *continuous*. That is to say, the amount of effect that one’s action can have is, for all practical purposes, continuously distributed (*i.e.*, it can take any real-number value). An example is greenhouse-gas emissions from driving one’s car: for every extra meter that one drives, a slightly higher amount of carbon dioxide is emitted into the air, so that driving one meter less has a quantifiably smaller adverse impact.³

The second situation is one of *discretely distributed effects* or, abbreviated, *discrete*. Here there exist only a finite number of possible outcomes, and often there are just two: success or failure. For example, passage of a bill in its final form can be considered discrete, because it either passes or fails, and there is no in-between. A bill to which two amendments have been proposed has four possible outcomes: passage with both amendments, passage with only the first, passage with only the second, or defeat.

Application of the all-or-nothing fallacy to continuous situations is rather easily refuted. Those who accept the fallacy generally believe that, were there a critical mass of people taking the given action, the results would far outweigh the costs. For example, people who feel that it would be useless for them individually to donate \$20 to an advocacy organization would welcome a society in which everyone gives \$20 to such organizations,

³Mathematically, a discrete variable can take any finite or countably infinite number of possible values. A continuous variable can take any real number value (*i.e.*, uncountably infinite possible values). As far as I can see, all changes that one can have on the world are actually discretely distributed. For example, one can only emit 1,000,000,000,000,000,000,000,000 molecules of carbon dioxide or 1,000,000,000,000,000,000,001, and nothing in between. One can only produce 1,000,000 happy-chemical molecules or 1,000,001. So, in reality, everything has a finite number of possible outcomes. However, this is just a technicality; our perception of these variables is much closer to a continuous one. And since we cannot measure the changes with nearly this much precision anyway, we only employ mathematical models to approximate these variables, and mathematical models, of course, will assume, for the sake of simplicity, a continuous distribution.

because the enormous results of the latter situation are much more conspicuous. This feature is what makes the fallacy “all-or-nothing” rather than just “nothing”: under the former justification of inaction, people claim that if enough others took the action, the problem would be all solved, but if only they act, nothing will happen; those with the latter view do not believe that action will make a difference whether they alone or a group of ten million other people do it. The latter view may at many times be admirable; the former view is not. Because the effect distribution is continuous, the result when one person takes the action is one-one millionth of the result when a million people take the action; the burden of one person taking the action is one-one millionth of the burden of a million people taking the action. So if the costs of a million people doing something are outweighed by the benefits, then the cost of one person doing that same thing will also be outweighed by the benefits. This assumes, of course, that the relationship between number of people taking action and amount of benefit is best approximated by a line (rather than a parabola, exponential graph, logarithmic graph, et cetera), but I think that such an assumption is fairly accurate for the types of actions from which people generally try to exempt themselves by means of the continuous all-or-nothing fallacy: donating money, driving less, recycling, conserving paper, conserving electricity, and others.

The all-or-nothing argument takes on the appearance of greater sophistication in discrete situations. For instance, a single person who campaigned for John Kerry in Ohio in 2004 can tell himself or herself that his or her effort was futile, because George W. Bush would have won the state either way. Similarly, if a person wrote a letter to a wavering member of Congress whose vote determined whether or not the Clear Skies bill got out of committee and the Congressperson ultimately decided to vote contrary to the wishes of the letter writer, the letter writer might feel that his or her effort was wasted and that he or she should not have bothered. These positions seem *prima facie* defensible. But they are not. Before the event in question actually happened, the person deciding to act did not already know what the result would be. Under other circumstances, that person might have been the deciding factor. And since the wisdom of a decision is determined by the probability of the possible outcomes at the time the decision was made—not by the specific outcome that actually occurred—the decision to campaign in Ohio or to write the letter was a good one even though the desired result did not happen to come about.

This idea may be best illustrated by a simpler example. Imagine a lottery for which tickets cost \$1 each and for which the prize is \$100 million. The probability that any given ticket will be the winning ticket is 1/1,000,000. Assuming that those playing this lottery have enough superfluous money that losing \$1 will have no impact on their standard of living, the lottery is a great deal, for the expected value of buying one ticket is \$99. It is certainly better than paying \$1 to earn a guaranteed \$10. (I must immediately qualify these statements as being predicated on the assumption that winning \$1,000,000 is a million times as good as winning \$1, et cetera. While this assumption rarely holds true for money because its marginal utility drops off so sharply, it often does for public-policy gains. For example, precluding two million metric tons of carbon dioxide emissions really is almost

twice as good as precluding one million metric tons. And preventing 10,000 animals from enduring brutal lives in factory farms really is ten times as good as preventing 1,000 animals from suffering the same fate. If I had said that the money won through this lottery would go to a charitable cause, the assumption would continue to hold true.) The analogy to the discrete all-or-nothing fallacy is obvious: buying a lottery ticket represents political advocacy, such as campaigning for Kerry in Ohio, while accepting the guaranteed \$10 represents direct charity. The results of the latter may be more certain, but the expected value of the gains is far smaller. Clearly, someone who bought one of the hypothetical lottery tickets should not regret his or her decision simply because he or she did not win; buying the ticket is still the best decision even though 99.9999 percent of people do not win it.

Thus, the discrete application of the all-or-nothing argument is as fallacious as its continuous counterpart.

7.4 Does Veganism Make a Difference? [19]

One of the most common charges against conversion to veganism is that abjuring the consumption of factory-farmed animal products will not actually have a real-world impact. Indeed, this is the argument of last resort for those who come to accept the utilitarian position on factory farming yet refuse to change their diets accordingly. In this essay, I shall endeavor to refute this weak excuse.

Let us first examine the logic behind the position that abstention from factory-farmed products is not morally implied by concern for animals. Everyone in the debate agrees that, at some point, a substantial decrease in demand for meat—for instance, conversion to veganism by half of the US population—would diminish the quantity supplied of factory-farmed animals. Consider the basic supply and demand curves from introductory economics. The obstinacy with which most people maintain their right to eat meat should provide some indication as to how highly people value it: if people refuse to stop eating animal products when presented with overwhelming moral reasons, then surely they will not greatly reduce their consumption as a consequence of higher prices. In other words, I consider it reasonable to assume that demand for factory-farmed animal products is rather inelastic. At the other end of the market system are the producers of factory-farmed animals. Conjecturally, it does not seem that the marginal cost of animal production would increase very steeply, considering that economies of scale are what created factory farms in the first place. So the marginal cost curve—which is the same as the supply line—has a low slope and is rather elastic. Now, when one considers a massive exogenous decline in demand, one sees that its principal effect is to decrease quantity supplied; the slight drop in prices that accompanies the demand shift does not substantially increase quantity demanded along the demand curve, since the demand line itself is rather inelastic.

This is all very nice, argues the opponent of veganism, but it has nothing to do with

individual purchasing habits. The only decisions that will actually affect the amount of animal suffering are massive investment choices: whether to build a new factory farm or whether to put an addition on a colossal chicken house. One consumer, the critic maintains, cannot even change the bulk purchases that each individual grocery store or dining center makes, much less the collective demand that all of these retailers exert. So I can continue to eat meat, eggs, and milk without creating demand for the production of any more animals. A certain amount of meat will be produced whether I am a vegan or not, so why shouldn't I at least subsist off those excess scraps, instead of letting them go to waste?

I shall respond first with a simple *reductio ad absurdum*. The opponent of veganism asserts that her individual consumption will have no effect on the actual number of animals raised. Since the opponent's use of animal products is no different from that of any one else, it is possible for any given person to make this excuse. Eventually, it would be possible for the entire meat-consuming population to maintain the same defense of their habits, with everyone claiming individually that his or her personal diet does not affect total demand.

The reason behind the absurdity in the anti-veganism position can be understood in terms of the discrete all-or-nothing fallacy (see section 7.3, "The All-or-Nothing Fallacy"). Suppose that a supermarket currently purchases three cases per week of factory-farm milk. The store does not purchase fractions of cases, so even if several cartons of surplus milk remain on the shelves each week, the supermarket will continue to buy three cases. This is what the anti-vegan means by subsisting off of surplus animal products that would otherwise go to waste: the three cases are purchased anyway, so consuming one or two more cartons simply attenuates the surplus.

What would happen, though, if five hundred customers decided to buy soy milk instead of cow's milk? The purchasing agent who orders weekly cases of milk would probably buy two cases instead of three. So when five hundred fewer customers demand the product, the discretely distributed quantity supplied lessens. At some point between a lowered demand of zero cartons and five hundred cartons, the purchasing agent made the decision to order one less case; however, it's impossible to determine exactly where the dividing line will be. The agent's decision may be arbitrarily based on approximately how many cartons of milk the store appears to have at the moment, and it is impossible to know ahead of time whether that arbitrary judgment will be made between the 100th and 101st cartons that are not demanded, or between the 386th and 387th cartons. But somewhere, that dividing line exists.⁴

Not only does any given consumer not know the number of the carton that will determine how many cases the supermarket buys; she also does not know the number of the carton that she is purchasing, because she has no way of knowing how many other consumers like her happen—for whatever reason—to be demanding fewer cartons as well. So even if she could establish that the supermarket would not buy one fewer case of milk

⁴I do not know exactly how ordering decisions are made at supermarkets and retail stores; perhaps they are based on historical data of consumption levels at various times of the year. But regardless of the detailed procedure, the point is that they are based in some way on the actual amount of products purchased.

until consumers bought four hundred fewer cartons, she would have no way of knowing whether the individual carton that she refrains from purchasing is that four hundredth carton or—what is more likely—whether her refusal to buy the carton will cause someone else to be the four hundredth person not to buy a carton.

The same logic applies in reverse to the anti-vegan’s “subsistence from surplus” excuse. The opponent of veganism claims that the supermarket would buy three cases of milk per week anyway, regardless of whether he himself buys one carton and lowers the store’s surplus waste from 18 to 17 cartons. But at some point—at some minimum level of surplus or deficit—the store’s purchasing agent will choose to buy four cases instead of three. As before, any individual consumer has no way of knowing at which number carton that change will occur or which number carton he happens to be buying.

Now, *e.g.*, assume that it takes four hundred fewer consumers of milk in order for the supermarket to buy one less carton. Inasmuch as individual consumers have no way of telling whether their particular carton will be the one that changes the number of cases purchased, the probability of any given carton being the determining factor is one in four hundred. The expected value of an action is the probability that a benefit will result times the magnitude of the benefit if it does result, so the expected value of refraining from the purchase of any given carton of milk is $(1/400)(1 \text{ fewer case purchased})(400 \text{ cartons/case}) = 1$ fewer carton purchased. The exact expected values will of course fluctuate on account of the randomness of the purchasing agent’s decisions (if, for instance, she would not buy one fewer case until five hundred fewer consumers demanded milk, even though each case includes only four hundred cartons), but they should average out over the long run in such a way that forbearing the purchase of any given amount of an animal product will be expected to reduce bulk purchase of that amount of the product.

The logic explicated above should also apply to the rest of the factory-farmed-meat demand process: at some critical mass of fewer cartons ordered by stores, distributors will purchase less meat from farms, and that reduced demand from farms will, at some point, constrict production. By the end, the probability that any given consumer will impact animal production is miniscule, but the benefits if he does are immense. Thus, the expected value of refraining from the purchase of any given amount of an animal product is roughly equivalent to preventing the production of the portion of an animal that the product represents. Of course, it is entirely possible (and perhaps even likely) that a vegan may go through her entire life and never, by failing to purchase factory-farmed-animal products, have actually prevented any animal suffering by lowering production. But because she has no way of knowing when the special purchase that does set off the chain of significant demand reduction will be, she has to act as if every purchase does count. (For more on this idea, see section 6.5, “Expectational Consequentialism.”) And for all she knows, she may just as easily be a consumer who has more than her share of impact on demand for factory-farmed products.

While the above mechanism for the impact of veganism may apply to most cases, there are some situations in which the demand-reduction benefits of abstaining from animal

products are less likely to materialize. The most notable example might be a party or picnic to which—despite one’s best efforts—people have brought hot dogs and hamburgers. The purchasing decisions a supermarket may be imprecise and sometimes arbitrary, but they will eventually be affected if demand changes by a great enough amount. The same can not be said of those who purchase items for a picnic. In general, the purchaser will buy some overestimated amount of food beforehand, regardless of how many people actually consume those comestibles at the event. And whereas a store that purchases far too much of a product will keep records and change its behavior the next time, people buying picnic food probably will not. So it is quite unlikely that one’s decision to eat or refrain from eating a factory-farmed-animal product at an informal social gathering will make a difference in the amount of food that the organizer purchases the next time. (Perhaps the best way to limit the harm done by a picnic is to ask the purchaser ahead of time to procure some vegan products.)

However, it is not always the case that food procurement for social events is irrational and fails to consider previous experience. To cite an exception, consider my first-year combination-math-and-physics seminar. In the middle of a three-hour seminar, the class takes a break for snacks. During the first such break in my math seminar, the professor brought the snacks: chocolate-chip cookies, cheese crackers, and carrots. He asked if anyone was a vegan; when I said that I was, he informed me that only the carrots did not contain animal products. While this was a social snacking event like a picnic or party, my interest in vegan products was not forgotten by the next seminar, when students brought the snacks for physics. The student purchasing the food told me that he had looked for vegan cookies and found some that were apparently vegan, inasmuch as no animal products were listed on the ingredients label. (What is interesting to note here is that, had I adopted the attitude that the purchased food is a sunk cost and that my consumption of it would not contribute to greater demand in the future, the result actually would have been more animal products consumed in the future. It was only by adopting an ostensibly purist vegan attitude, according to which people presumed that I objected to the ingestion of animal products as such, that I was able to effect a consequentialist impact on economic demand.) The lesson from this incident is that one must be cautious in judging under which circumstances one’s consumption of animal products will genuinely have no chance of increasing demand.

Moreover, even if eating animal foods at a picnic or gathering would assuredly not change the direct amount of meat and milk purchased, there might still be other good reasons for refraining from the consumption of animal products. First, people may not realize the extensive reasoning presented above and thence may view a vegan’s consumption of a hamburger as contradictory (even if it actually is not). This misunderstanding might diminish the respect that others have for the vegan in particular and for the cause of veganism in general. Second, abstention in one’s diet—along with explicit inquiries regarding which products are vegan—can raise the topic of factory farming. This, indeed, is one of the most important consequences of being a vegan in any context, for good discussions—combined with the literature exchange or further research that might follow—have the potential to

permanently change other people's eating habits.

7.5 How Much is a Dollar Worth? Some Conservative Estimates

People do not feel in any way ashamed or guilty about spending money on new clothes or a new car instead of giving it to famine relief. (Indeed, the alternative does not occur to them.) This way of looking at the matter cannot be justified. [. . . W]e ought to give money away, rather than spend it on clothes which we do not need to keep us warm. To do so is not charitable, or generous. Nor is it the kind of act which philosophers and theologians have called “supererogatory”—an act which it would be good to do, but not wrong not to do. On the contrary, we ought to give the money away, and it is wrong not to do so.

—Peter Singer, *Famine, Affluence, and Morality* [36]

Notes

- The basic premise of this essay is nothing original. Observers throughout history have noted the misallocation of society's resources, and this essay contributes little to that concept. What this piece is intended to do, however, is to pursue the question of specific numbers, at least insofar as they can be roughly approximated. For it is easy to carry a notion of the other uses to which one's money might be put as a vague abstract idea that one should heed once in a while; it is a very different decision-making dynamic when one has some tangible estimations of how much relief of suffering is forgone by a specific frivolous expenditure. That will be the focus of this piece.
- Many of the numbers used in the following estimations are very rough (sometimes made up entirely), so I did not consider it useful to make consideration for the rules of significant figures. Numbers are carried out to several digits not because they are certain to have those values but simply because including more than one digit gave them more tangibility.

Introduction

To many in the developed world, a single dollar does not seem to have significant value. Indeed, Western consumers regularly spend twenty or fifty dollars on a concert ticket or dinner at an elegant restaurant. For those who already have their basic physiological, social, and mental needs fulfilled, these uses of a dollar furnish only minor amounts of utility, and that immediate utility—as I argue in “Is Happiness Relative?”—may actually

be illusory in the long run anyway. Let us, then, compare these typical Western-consumer uses of a dollar with other uses to which the money might be put.

Example 1: Humanitarian Aid

In its 2001 fund-raising material, the U.S. Committee for UNICEF says that a donation of \$17 will provide immunization “to protect a child for life against the six leading child-killing and maiming diseases: measles, polio, diphtheria, whooping cough, tetanus, and tuberculosis,” while a donation of \$25 will provide “over 400 packets of oral rehydration salts to help save the lives of children suffering from diarrheal dehydration.” [37]

Compare this against other uses of money. For example, the average movie ticket in the US cost \$6.41 in 2005 [3]. A new Nintendo GameCube costs \$99.99, not including accessories and games [24]. A weeklong trip to Disney World by a family of four would require \$718.06 just to pay for the tickets alone [41]. The decision to buy a GameCube or go to Disney World—or even to see a few movies—is thus a genuine life-and-death matter.

Example 2: Vegan Outreach

VO is working to promote veganism through the widespread distribution of our illustrated booklets, *Why Vegan, Even If You Like Meat*, and *Try Vegetarian* (formerly *Vegetarian Living*). [...]

Millions of hard copies have been handed out by the local members of Vegan Outreach around the world. The rate of distribution is increasing every year, limited not by demand—there are many individuals, student groups, and organizations who would like to distribute “as many as you can send”—but by availability (*i.e.*, resources for printing and distribution). [2]

Until we decide to work specifically on the long-range task of bringing about animal liberation rather than “fighting brushfires,” there will be an ever-increasing number of animals exploited, and never-ending deluge of suffering. We have limited resources, and we can choose to use these resources to help end the cycle of exploitation and cruelty—saving billions of individuals over time, with no more to follow. [...]

Often, people believe that *Why Vegan* (or *Try Vegetarian*, or *Even If You Like Meat*) is an effective tool for many purposes (*e.g.*, to use at fundraisers for other groups). However, relatively few people feel that donating to print and distribute more copies is a priority.

It is true that the work of Vegan Outreach might not seem “urgent” because we don’t use pictures of individual animals saying “donate or this animal will die.” But consider this:

Right now, it is probable that 10 percent of the US population would be willing to read and consider the information in *Why Vegan*. For each one of these people who do not receive a copy of *Why Vegan* due to a lack of resources, numerous animals will be factory farmed and slaughtered. To reach these 26 million people, we would need to print and distribute between 30 and 50 million copies (for reference, as of January 31, 2005 we have distributed 3,879,329), at a cost of \$4.5 to 5.5 million. While a huge number relative to Vegan Outreach's current budget, this is a small fraction of the collective funds brought in by large animal rights organizations each year. [22]

In order to keep my calculation conservative, I shall assume the higher number given for the cost of reaching 26 million Americans: \$5.5 million. The average cost of reaching any one person is then \$0.21.

Of course, very few of the people who glance through *Why Vegan?* are likely to actually become vegan; 26 million must be multiplied by the probability that any given person who reads the literature becomes a vegan:

We [Vegan Outreach] estimate that out of 200 brochures given to students, at least 5 become vegetarian after reading it. We cannot know how accurate this is without a complicated study. We base our numbers on this:

Often while leafleting on a campus, students read the brochure and then approach us saying they are not going to eat meat any longer. Jack Norris averaged about 1 per campus when he was traveling to campuses, averaging about 235 brochures per campus (note: the pamphlets were less powerful and in black and white at the time he was travelling in 1996-97; we expect that they are more effective now). We also get e-mails from students saying they became veg and sometimes we run into students who told us they went vegetarian after reading the brochure from a previous time we were on campus. Assuming that there are many who never contact us after becoming veg, we conclude that 5 out of 200 is a good estimate.

We also assume that many students cut back on their animal products without fully going vegetarian. And we know that many students we reach will influence other students to become vegetarian or vegan. We figure that these numbers can make up for any overestimate of the "5 out of 200" number. [6]

It is important to note several factors about this estimate that make it legitimate for the present purpose. First, Vegan Outreach's number of five in two hundred is largely a measure of converted vegetarians who wouldn't have become vegetarian otherwise. It seems unlikely that someone who was already a vegetarian or nearly convinced of vegetarianism for some other reason would report his change in diet to Vegan Outreach; only those who were greatly impacted by Vegan Outreach's specific brochure would be likely to do so.

Thus, it seems likely that the creation of five vegetarians in two hundred people represents a positive impact that can be genuinely attributed to Vegan Outreach, not primarily to other sources that would have been present anyway. Second, the list of reasons stated for increasing the figure of 1 in 235 to 5 in 200 seems compelling, particularly considering how few people would ever take the time to notify Vegan Outreach of its impact.

Yet, a third factor must be considered. The population to whom Vegan Outreach has already distributed literature was probably one of the most receptive groups available; after all, there is probably some reason that Vegan Outreach chose the places that they did for literature distribution and not others. As one continues to disseminate information, the populations become less and less ideal. So perhaps the ratio of 5 in 200 is slightly too generous, because it was based on results in the best possible places. To make the present estimate more conservative—and appease skeptics who may dispute Vegan Outreach’s other reasons for elevating their ratio—I shall cut the number in half, to 2.5 in 200. Then, $(\$0.21/\text{distribution of literature to a receptive person})(200 \text{ receptive people}/2.5 \text{ new vegetarians}) = \$16.92/\text{new vegetarian}$.

Now, it is important to remember that a significant portion of people who initially commit to vegetarianism probably do not maintain it throughout their lifetimes. A conservative estimate of the fraction that do remain vegetarians for life might be one in three: $(\$16.92/\text{new vegetarian})(3 \text{ new vegetarians}/1 \text{ permanent vegetarian}) = \$50.77/\text{permanent vegetarian}$.

Since most people decide to become vegetarians in their teenage or young-adult years and since Vegan Outreach focuses on college campuses, the expected value for their future lifespan might be around 55 years (assuming the average new vegetarian is twenty years old and the average life expectancy for Americans is 75 years). $(\$50.77/\text{permanent vegetarian})(1 \text{ permanent vegetarian}/55 \text{ years of meat consumption avoided}) = \$0.92/\text{year of meat consumption avoided}$.

[Vegan Outreach] assume[s that] 35 animals (including only mammals and birds) are eaten per year per person. This number is figured by dividing the 9,997.3 million mammals and birds slaughtered in the U.S. in 2002 (numbers provided by Farm Animal Reform Movement [...]) by 283 million people. [6]

$(\$0.92/\text{year of meat consumption avoided})(1 \text{ year of meat consumption avoided}/\text{production of 35 animals in factory farms avoided}) = \$0.02637/\text{production of an animal in a factory farm avoided}$.

The overwhelming bulk of animals in factory farms are chickens, which are generally slaughtered at six weeks of age [18]. So six weeks might be a fair estimate for the duration of suffering of the average meat animal raised in a factory farm. $(\$0.02637/\text{production of an animal in a factory farm avoided})(\text{production of one animal in a factory farm avoided}/6 \text{ weeks of abject suffering avoided})(1 \text{ week}/7 \text{ days}) = \$0.000628/\text{day of abject suffering avoided}$. *In other words, a single dollar donated to Vegan Outreach would be expected to prevent 4.4 years of abject suffering.*

Compare this with the frivolous purchases described previously. For the price of the average movie ticket, a person contributing to Vegan Outreach could prevent $(\$6.21/\text{ticket})(1 \text{ day of abject suffering avoided}/\$0.000628) = 9,889$ days of abject suffering avoided/ticket = 27 years of abject suffering avoided/ticket. By the same calculations, the money spent on the tickets for a weeklong stay in Disney World could prevent 1.1 million days or 3,131 years (or 42 average-American-lifetime equivalents) of abject suffering.

And of course, these numbers reflect only the direct harm done to the farmed animals themselves. They do not include the countless wild animals that would be given the chance to live happy lives on account of the greatly reduced ecological impacts that conversion to a vegan diet entails.

Bibliography

- [1] “9/11 Fund Deadline Passes.” *CBS News*. 23 Dec. 2003. 8 June 2006 (<http://www.cbsnews.com/stories/2004/01/16/national/main593715.shtml>).
- [2] “About Vegan Outreach.” *Vegan Outreach*. 8 June 2006 (<http://www.veganoutreach.org/about/index.html>).
- [3] “Average US Ticket Prices.” *National Association of Theater Owners*. 8 June 2006 (<http://www.natooonline.org/statisticstickets.htm>).
- [4] Bentham, Jeremy. “Chapter 17, Section 1, Note 2.” *Introduction to the Principles of Morals and Legislation*. 1789. Accessed from *Classical Utilitarianism Web Site*. 20 Feb. 2002. 8 June 2006 (<http://www.la.utexas.edu/research/poltheory/bentham/ipml/ipml.c17.s01.n02.html>).
- [5] “Health Groups: AIDS War Falters.” *Bread for the World*. 26 Oct. 2003. 26 Oct. 2004 (http://www.bread.org/media/articles/2003/detroit_oct_26.htm). This site is no longer available directly, but readers can access it through the Internet Archive: (<http://www.archive.org/index.php>).
- [6] “Calculations.” *Vegan Outreach*. 8 June 2006 (<http://www.veganhealth.org/colleges/calc>).
- [7] “Chlorine-Free Processing.” *Conservatree*. 8 June 2006 (<http://www.conservatree.com/paper/PaperTypes/CFDisc.shtml>).
- [8] Darwin, Charles. “Chapter 21 - General Summary and Conclusion.” *The Descent of Man*. 1871. Accessed from *Literature.org*. 23 May 2005. 8 June 2006 (<http://www.literature.org/authors/darwin-charles/the-descent-of-man/chapter-21.html>).
- [9] *Draft Prevention of Cruelty to Animals (Prohibition of Glue Trapping) Regulations 2005*. State of Victoria, Department of Primary Industries, 2005. 8 June 2006 (<http://www.dpi.vic.gov.au/dpi/nrenfa.nsf/93a98744f6ec41bd4a256c8e00013aa9/>

1580e4773bb74e91ca257029000074bf/\$FILE/RISandRegs.pdf).

- [10] Eddington, Arthur Stanley. “The Decline of Determinism.” This piece can be found in Gardner, Martin. *Great Essays in Science*. Buffalo, NY: Prometheus Books, 1994.
- [11] “Frequency Probability.” *Wikipedia*. 8 June 2006 (<http://en.wikipedia.org/wiki/Frequentist>).
- [12] “Glossary.” *EcoHealth: Environmental Change and Our Health*. 8 June 2006 (www.ecohealth101.org/glossary.html).
- [13] “Glossary.” *State of the Environment - South Africa*. Oct. 1999. GRID-Arendal, United Nations Environment Programme. 8 June 2006 (www.ngo.grida.no/soesa/nsoer/general/glossary.htm).
- [14] Goldberg, Steven. *Why Men Rule: A Theory of Male Dominance*. Open Court, 1993.
- [15] Gould, Stephen Jay. “Nonmoral Nature.” *Hen’s Teeth and Horse’s Toes: Further Reflections in Natural History*. New York: W. W. Norton, 1994. pp. 32-44. Accessed from *Unofficial Stephen Jay Gould Archive*. 8 June 2006 (http://www.stephenjaygould.org/library/gould_nonmoral.html).
- [16] “Humane Eating and the Three Rs.” *Humane Society of the United States*. 8 June 2006 (<http://www.hsus.org/farm/humaneeating/rrr.html>).
- [17] Kurzweil, Ray. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. Penguin, 2000.
- [18] “Life Cycle of a Broiler Chicken.” *Animal Doc Com*. University of Georgia. 8 June 2006 (<http://www.uga.edu/lam/kids/poultry/plifecycle.html>).
- [19] Matheny, Gaverick. “Expected Utility, Contributory Causation, and Vegetarianism.” *Journal of Applied Philosophy* 19.3 (2002): 293-97. (<http://www.veganoutreach.org/enewsletter/thresholds.pdf>). This piece is very similar to my own—surprisingly so, in fact, because I only came across it only after I had written mine.
- [20] “Millennium Development Goals: 1. Eradicate Extreme Poverty and Hunger.” *UNICEF*. 8 June 2006 (<http://www.unicef.org/mdg/poverty.html>).
- [21] Mitchell, Luke. “A Run on Terror: The Rising Cost of Fear Itself.” *Harper’s* 4 March 2004. 25 April 2004 (<http://www.harpers.org/ARunOnTerror.html>).

- [22] “Money or Dog?” *Vegan Outreach*. 8 June 2006 (<http://www.veganoutreach.org/about/moneyordog.html>).
- [23] Mower, Jacob. Personal email to the author. 30 Oct. 2005.
- [24] “Nintendo Game Systems.” *Nintendo*. 8 June 2006 (<http://www.nintendo.com/systems>).
- [25] Orwell, George. “Part 3, Chapter 1.” *1984*. 1949. Accessed from *The Literature Network*. 8 June 2006 (<http://www.online-literature.com/orwell/1984/18/>).
- [26] “Physical Inactivity and Poor Nutrition Catching up to Tobacco as Actual Cause of Death.” *CDC*. 9 March 2004. 8 June 2006 (<http://www.cdc.gov/OD/OC/MEDIA/pressrel/fs040309.htm>).
- [27] “Principal Doctrines: Epicurus.” *Epicurus and Epicurean Philosophy*. 8 June 2006 (<http://www.epicurus.net/en/principal.html>).
- [28] Shakespeare, William. *The Tragedy of Hamlet, Prince of Denmark*. 1601. Accessed from *The Complete Works of William Shakespeare*. 8 June 2006 (<http://www-tech.mit.edu/Shakespeare/hamlet/full.html>).
- [29] Sidgwick, Henry. *Methods of Ethics, Seventh Edition*. London: Macmillan, 1907. Accessed from *Classical Utilitarianism Web Site*. 15 Feb. 2002. 8 June 2006 (<http://www.la.utexas.edu/research/poltheory/sidgwick/me/index.html>).
- [30] Silverstein, Ken. “Millions for Viagra, Pennies for Diseases of the Poor.” *The Nation* 1 July 1999. 8 June 2006 (<http://www.thenation.com/doc/19990719/silverstein>).
- [31] Singer, Peter. *Animal Liberation*. New York: HarperCollins, 2002.
- [32] Singer, Peter. “On the Appeal to Intuitions in Ethics.” Excerpted from Jamieson, Dale, ed. *Singer and His Critics*. Oxford, 1999. pp. 316-18. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/1999- - -.htm>).
- [33] Singer, Peter. “Catastrophe: Apocalypse When?” *New York Times* 2 Jan. 2005. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/20050102.htm>).

- [34] Singer, Peter. "Equality for Animals?" Excerpted from Chapter 3, *Practical Ethics*. Cambridge, 1979. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/1979- --.htm>).
- [35] Singer, Peter. "The Escalator of Reason." Excerpted from *How Are We to Live?* New York, 1995. pp. 226-35. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/1995- --.htm>).
- [36] Singer, Peter. "Famine, Affluence, and Morality." [revised edition] *Philosophy and Public Affairs* 1.1 (Spring 1972): 229-243. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/1972- ---.htm>).
- [37] Singer, Peter. "Poverty, Facts, and Political Philosophies: Response to 'More Than Charity.'" *Ethics and International Affairs* Spring 2002: 121-24. Accessed from *Utilitarian Philosophers: Peter Singer*. 8 June 2006 (<http://www.utilitarian.net/singer/by/2002- ---.pdf>).
- [38] Twain, Mark. "Chapter 13 - Freemen." *A Connecticut Yankee in King Arthur's Court*. 1889. Accessed from *Literature.org*. 23 May 2005. 8 June 2006 (<http://www.literature.org/authors/twain-mark/connecticut/chapter-13.html>).
- [39] *Utilitarian Philosophers*. 8 June 2006 (<http://www.utilitarian.net/>).
- [40] Veenhoven, Ruut. "Is Happiness Relative?" *Social Indicators Research* 24 (1991): 1-34. Accessed from Erasmus Universiteit Rotterdam. 8 June 2006 (<http://www2.eur.nl/fsw/research/veenhoven/Pub1990s/91a-ab.htm>).
- [41] "Walt Disney World Ticket Store." *Walt Disney World Resorts*. 1 September 2005 (<http://disneyworld.disney.go.com/wdw/common/tickets/buildticketurlredirect>). For my calculation, I assumed that the family of four would consist of two adults and two children between the ages of three and nine.