

Class #26 - Consequentialism
Mill, from *Utilitarianism*
Nozick, "The Experience Machine"

I. Toward Moral Objectivity

In our previous two classes, I gave reasons to prefer moral objectivity over either nihilism or relativism. The main advantage of objectivity is that it captures our desires to see our values as transcending us. While we wish to be humble about how well we know morality, the inability to praise and criticize other cultures is a serious problem with both nihilism and relativism. Consider this description of the massacres in Rwanda, in 1995.

Encouraged by political and civic leaders, the massacring of Tutsis spread from region to region. Following the militia's example, Hutu young and old rose to the task. Neighbors hacked neighbors to death in their homes, and colleagues hacked colleagues to death in their workplaces. Priests killed their parishioners, and elementary-school teachers killed their students. Many of the largest massacres occurred in churches and stadiums where Tutsis had sought refuge - often at the invitation of local authorities, who then oversaw their execution. In mid-April, at least five thousand Tutsis were packed in the Gatwaro Stadium, in the western city of Kibuye; as the massacre there began, gunmen in the bleachers shot zigzag waves of bullets and tossed grenades to make the victims stampede back and forth before militiamen waded in to finish the job with machetes.

Throughout Rwanda, mass rape and looting accompanied the slaughter. Militia bands, fortified with potent banana beer and assorted drugs, were bused from massacre to massacre. Hutu prisoners were organized in work details to clear cadavers. Radio announcers reminded listeners to take special care to disembowel pregnant victims. As an added incentive to the killers, Tutsis' belongings were parceled out in advance - the radio, the couch, the goat, the opportunity to rape a young girl. A councilwoman in one Kigali neighborhood was reported to have offered fifty Rwandese francs apiece (about 30 cents at the time) for severed heads, a practice known as "selling cabbages" (*The New Yorker*, December 18, 1995).

It is difficult to believe that the acts described here are not clearly and decisively immoral. We respond to the Rwanda case by thinking of reasons why such behavior is wrong. These reasons push us away from relativism, on which the only reasons required are the dictates of the culture. While we surely wish to allow cultural variations in preferences regarding how to express morality, we also want to be able to communicate with others about the good. By adopting moral objectivity, we can consider culturally-independent reasons for morality. Reason-giving is essential to moral discussion. Such reasons may be used as the basis for a universal morality.

Objectivity is the claim that there are morally correct answers to ethical questions. It does not mean that everything that one might think is a moral question is in fact a moral question. The specifics of how we deal with the remains of our dead may not be a moral issue, contrary to Darius's conclusion. Objectivity should also not entail dogmatism. It is not a claim that any one knows what the right morality is. There are hard moral questions.

We should explore methods for determining the correct morality.

We are seeking the reasons which naturally constitute morality.

We need an ethical theory: a set of universal principles that tell whether an action is right or wrong.

The Ten Commandments is an ethical theory, though not a particularly good one:

It has the correct form.

It is prescriptive and it overrides non-moral prescriptions, like those found in aesthetics, law, manners, or grammar.

It is universalizable, applying to everyone.

But, it lacks the flexibility or abstractness that a good moral theory should have.

Consider the proscription against killing.

There may be cases in which killing is justified, say in self-defense or in a just war.

The Ten Commandments may be good general rules.

But those rules admit of exceptions.

Must we honor abusive parents?

Can [stealing never be justified?](#)

There is nothing in that set of rules which tells us how to determine whether a particular case is an exception or not.

We are going to look at two moral theories more plausible than the Ten Commandments: Mill's utilitarianism and Kant's deontological ethics.

II. The Three Clauses of Utilitarianism

Discussions of utilitarianism often center around the work of the nineteenth-century English philosopher John Stuart Mill.

Versions of utilitarianism predated Mill, and there are contemporary utilitarians.

But, Mill's version of the theory is classic.

Utilitarianism can be condensed into one single principle, called the Greatest Happiness Principle (GHP).

GHP	The right act is the act which produces the greatest happiness for the greatest number.
-----	---

Mill formulates GHP slightly differently.

The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure (Mill 141a).

We can analyze the utilitarian theory into three clauses:

U1: <i>Consequentialism</i>	Acts are judged by their consequences.
U2: <i>Hedonism</i>	Consequences are evaluated by the amount of total happiness they bring.
U3: <i>Egalitarianism</i>	Each person counts as one.

Considering U1, utilitarianism captures our bare intuition about the moral relevance of consequences.

By focusing on consequences rather than categorical rules, utilitarianism can be flexible.

Moral theories which include specific rules like "Don't lie," and "Don't kill," or any other absolute

proscription may be refuted by simple counterexamples.

There are times when any such specific prohibition should be violated.

Consider Danish fishing boats being stopped by the SS when transporting Jews out of areas of Nazi control in World War II.

In such a case, a lie seems morally required.

Consider another universal maxim, like “You must fulfill your promises.”

Imagine you have promised to meet some one for lunch, but find yourself in a position in which you could save a life, say passing a traffic accident.

It is surely morally permissible to break your promise.

Utilitarianism is thus a more honest theory than one which provides universal moral rules.

It builds in the exceptions we ordinarily deem acceptable or even required.

Utilitarianism tells us how to determine if an act would be an exception to any general rule or guideline.

Universal prescriptions do not tell us their limits.

Instead of a lot of detailed rules and confusion about when they apply, utilitarianism gives you one flexible, general always-applicable guideline.

Considering U2, utilitarianism is an Epicurean philosophy: happiness, or pleasure, is the goal of life.

The ultimate end, with reference to and for the sake of which all other things are desirable (whether we are considering our own good or that of other people), is an existence exempt as far as possible from pain, and as rich as possible in enjoyments, both in point of quantity and quality; the test of quality, and the rule for measuring it against quantity, being the preference felt by those who in their opportunities of experience, to which must be added their habits of self-consciousness and self-observation, are best furnished with the means of comparison (Mill 144a).

Mill’s godfather, Jeremy Bentham, provides seven ways to measure happiness.

B1. Intensity

B2. Duration

B3. Certainty

B4. Propinquity (proximity) or remoteness

B5. Fecundity (capability to produce more, followed by similar feelings)

B6. Purity (the chance an experience has of not being followed by opposite sensations)

B7. Extent

B1 - B4 measure the pain itself.

B5 and B6 measure the tendency of an action related to other actions.

B7 measures the way in which an act affects other people.

To calculate the total effects of an act, you just add up all the effects on each individual.

To add effects, we can imagine some basic units of happiness which have been called utils.

We must be careful to distinguish utils from money and other fungible goods.

For example, money abides by a law of diminishing returns: more money is always more money, but more money does not always lead to more happiness.

The first cold drink on a hot day is great but the twelfth is not so good.

Utils do not suffer diminishing returns.

They are just ways we can imagine of measuring happiness itself.

There may be problems quantifying happiness.

Some people argue that there are no valid interpersonal comparisons of happiness.

For example, my daughter's happiness when getting a piece of chocolate seems much greater than my own.

But, this is just an argument that she should have more chocolate.

Economists make interpersonal comparisons all the time.

Though economists often work with preferences, rather than happiness, they do so with the implicit assumption that we are happy when our preferences are fulfilled.

Mill's principal argument for hedonism is that any thing we value we do so because of the happiness it brings us or others.

If human nature is so constituted as to desire nothing which is not either a part of happiness or a means of happiness, we can have no other proof, and we require no other, that these are the only things desirable. If so, happiness is the sole end of human action, and the promotion of it the test by which to judge of all human conduct... (Mill 145b).

Every one wants to be happy, and if we all followed utilitarianism, then happiness would increase.

The egalitarian clause of utilitarianism, U3, is not very controversial.

It is true that we value some people more than others: the President of the United States, Albert Pujols, Lady Gaga.

The utilitarian can account for valuing some people over others without abandoning egalitarianism.

On such an account, Lady Gaga is not intrinsically more valuable than other people.

But actions that affect her and which she performs have a greater range of effects.

So, there is no need to count each one as more than one.

There is a good question about how widely to extend the egalitarianism.

Should replicants be counted? Aliens? Dolphins or chimps?

The answers to the questions we studied earlier in the course about personal identity and consciousness are essential to determining the range of our moral theory.

There are different versions of utilitarianism.

Many abandon U2, exchanging it for other ways to evaluate consequences.

Others abandon U3, though giving up egalitarianism is not a standard move.

But, any version of utilitarianism will be consequentialist.

III. Utilitarianism, Egoism, and Prudence

It is tempting to misinterpret utilitarianism as egoism, so it will be worthwhile to distinguish the two theories.

Utilitarianism says that the right act is the one which produces the greatest happiness for the greatest number.

Ethical egoism says that the right act is the one which produces the greatest happiness for me.

Consider an ethical egoist named Al.

Al believes that every one else should act to serve Al.

The utilitarian, in contrast, says that we should all act for the benefit of everyone.

Make sure to distinguish ethical egoism from the more plausible claim that every one should act to serve him or her self.

The latter claim is just self-interest, or prudence, and is really a form of subjectivism.

The egoist says that everyone should serve me.

The person who urges selfishness or prudence for everyone says that we should all act to serve ourselves.

One might be inclined to follow only one's self-interest because of psychological egoism.

The psychological egoist says no one ever does anything that he or she does not want to do.

Faced with someone who sacrifices for others, a Mother Teresa, say, the psychological egoist says that sacrifice must be rewarding to that person.

Psychological egoism may well be right.

But it is irrelevant to ethics.

It leaves us with no way to distinguish among good, not so good, and truly bad people or actions.

It is philosophically vacuous, even if psychologically compelling.

Utilitarianism is thus not the claim that we should act only in our own interests, nor that everyone should act in one's own interest.

It is the claim that we should all act in ways that best further every one's interests.

IV. Utilitarianism in Practice

Imagine a world in which there are two people, John and Harriet, and each has ten units of happiness.

Harriet wonders whether she should gather some flowers.

Having some flowers would increase her happiness by, say, two units.

The new totals would be: John = 10, Harriet = 12.

Since the new total would be 22, which is greater than 20, she should gather the flowers.

That is, we have a moral duty to increase our own happiness.

Another day, John = 10, Harriet = 10

John wonders if he should gather flowers for Harriet.

John does not like flowers, but Harriet does.

John would prefer to go swimming.

The new values would be John = 9, Harriet = 12.

The total would increase from 20 to 21.

So John should forego swimming to get the flowers for Harriet.

We have a moral duty to sacrifice ourselves for others, when the rewards are greater than the sacrifice when every one's happiness is measured.

Utilitarianism encourages working hard to get a raise, to provide better for one's family.

We should go out of our way for a stranger in dire need.

Sacrifice has its limits, though.

We should not give more than is gained.

If John really hates collecting flowers and Harriet only like flowers a little bit, then there is no moral requirement for John to go pick them.

The requirement of self-sacrifice also prevents utilitarianism from supporting immediate pleasure-seeking, in a narrow sense.

We are often required to sacrifice in the short term in the hopes of long term gains, for oneself.

For example, many of us calculate that we should work hard to get a college degree.

This is like moving from $A=B=10$ to $A=9, B=12$, where A represents your current self, and B represents your later self.

The question to ask is whether the happiness I will gain later outweighs the happiness I sacrifice now.

A slightly less abstract, though still artificial, illustration of utilitarian ethics is the famous [trolley problem](#), video [here](#).

According to the utilitarian, all else being equal, it is always right to sacrifice one to save many.

The ends justify the means; what else could?

The utilitarian's demands for sacrifice may be extreme.

The utilitarian must put her own interests aside, and treat herself as one individual affected by her actions. It looks like I should give away most of my belongings in order to try to combat humanity's worst conditions: famine and disease and war.

This impartiality may be impractical.

Mill agrees that it is difficult, but does not see this difficulty as a problem for the theory.

Those among [the objectors to utilitarianism] who entertain anything like a just idea of its disinterested character, sometimes find fault with its standard as being too high for humanity. They say it is exacting too much to require that people shall always act from the inducement of promoting the general interests of society. But this is to mistake the very meaning of a standard of morals, and confound the rule of action with the motive of it. It is the business of ethics to tell us what are our duties, or by what test we may know them... (Mill 144a-b).

When considering large-scale applications of utilitarianism, interesting questions arise about whether to consider average happiness or total happiness.

To increase average happiness in a community, we can (quietly) kill all of those whose happiness is below average.

To increase the total, we can require a population explosion.

Total utilitarianism seems preferable, since the population explosion will have long-term ill-effects.

But, I'll put these questions aside, here.

V. Is Utilitarianism the Right Moral Theory?

Utilitarianism differs from common-sense morality in some important ways.

It does not prescribe universal rules.

It also does not allow room for supererogation, doing more good than is morally required.

Ordinarily, some acts are morally required, like refraining from murder.

But other acts are even better: giving twenty percent of one's income to worthy charities, say.

Or, volunteering on Friday nights at a soup kitchen.

The utilitarian always demands that we maximize the consequences.

So, there is no way to do more than is morally required.

The case of supererogation may just be one in which our moral intuitions are wrong.

If utilitarianism were the right moral theory, then we could easily understand how some people would take an accepting attitude to those of use who failed to be fully moral, and a more laudatory attitude to those who hit the moral jackpot.

The common-sense notion of supererogation is really just this attitude.

Mill defends utilitarianism with a simple argument:
People desire happiness.
Utilitarianism says that we should increase happiness.
So, utilitarianism is the right moral theory.

Hume famously warned that one can not derive an ought from an is.
He argued that the world of values is subjectively imposed on an objective world of facts.
To find moral value in the world itself derives an ought (what we should do) from an is (the world).
Mill's argument seems to commit just that fallacy,
People in fact desire happiness.
Whether they should desire happiness remains an open question.
Whether utilitarianism is the correct guide to our actions depends on the answer to that question.
Still, we should look at the theory and its use, and compare it to other theories.

I will proceed to considering some criticisms of utilitarianism
Along the way, try to notice forms of responses:
One can dismiss the objection.
One can adjust the theory in response to the objection.
One can reject the theory, abandon it, in the face of insuperable criticism.

VI. Criticisms of Hedonism

Mill says that some critics of utilitarianism say that hedonism provides too low a standard for morality.

To suppose that life has (as they express it) no higher end than pleasure - no better and nobler object of desire and pursuit - they designate as utterly mean and groveling; as a doctrine worthy only of swine... (Mill 141b).

Mill responds to this criticism by distinguishing higher and lower pleasures.
Since the higher pleasures are more valuable, Mill argues, the utilitarian's standard is not too low.
We have to decide between pleasures, so we need a criterion.
We can ask someone who has had both.
Those who have had both higher and lower pleasures will say that some pleasures require work, but are worth it.

It is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied. And if the fool, or the pig, are of a different opinion, it is because they only know their own side of the question (Mill 143a).

A specific list of higher pleasures may be difficult to formulate.
People's preferences vary.
Still the counsel to seek the pleasures which are more valuable is uncontroversial.
Mill's theory need not rely on a pandering form of hedonism.
This objection may thus be easily dismissed.

The real problem with hedonism may be seen by considering the example from Robert Nozick.
Nozick presents a thought experiment: consider an experience machine.
The experience machine could give us any experience we wanted.

We could get as much happiness as we wanted from our experiences while plugged into the machine.
We would not actually experience the simulated experiences.
We would not be what we experienced ourselves as being.
Would we want to plug ourselves in to the machine?
If we decide not to plug in to the machine, then there might be a problem with Mill's utilitarianism.

Nozick's claim is that happiness is not the right aim for a moral theory.
We can be happy and still not satisfied with our lives.
Consider an accident which renders a young pianist unable to use her hands, but whose happiness may be insured by a 'happy pill'.
If the utilitarian says that all that matters is happiness, and the happiness is maintained, it seems that the utilitarian has no means for describing why this is a bad outcome.
Mill says fulfilling your goals is a means to your happiness, but maybe the goal is independent of happiness.
One might try to dismiss the objection based on the pianist case as an implausible scenario.
For this response, the utilitarian would argue that the pianist's happiness would necessarily be decreased by her dissatisfaction.

One utilitarian response to these criticisms of hedonism dismisses these two objections as based on implausible scenarios.
There is no such thing as a happiness pill or an experience machine.
Still, such cases are logically possible.
It would be nice if the utilitarian could account for them.

The utilitarian can also claim that our intuitions in these cases are wrong.
If we really had a happy pill or an experience machine, then even if the pianist lost the use of her hands, as long as she were happy, we should not see this as a problem.
We are possibly misled by our intuitions.

VII. Preference Utilitarianism

A different kind of response to the criticisms of hedonism we have seen is possible.
Instead of dismissing the objection, we might adjust the theory.
If we are convinced that the utilitarian inappropriately emphasizes happiness, we can drop the happiness clause (hedonism).
Without U2, the utilitarian maintains U1 and U3, consequentialism and egalitarianism.
But we are left without a way to evaluate consequences.
We need to replace hedonism with something.
Some versions of utilitarianism replace hedonism with an appeal to personal preferences or desires.
Counting the satisfaction of preferences instead of happiness may capture what is wrong with the scenarios above.
The pianist may be happy, but not in the way she prefers.

This proposal generates what is called preference-satisfaction utilitarianism, or just preference utilitarianism.

Preference Utilitarianism

The right act is the one that creates the greatest fulfillment of personal preferences for the greatest number.

Notice that the utilitarian theory is saved, though amended.

On preference utilitarianism, we work to satisfy people's preferences, like the preference of the pianist to be happy by playing the piano instead of by taking a happy pill, or our preference to have real experiences rather than artificial ones, even if we could be happy with artificial ones.

Unfortunately, there are serious problems with preference utilitarianism.

Some people have unacceptable preferences, crazy and even self-destructive desires.

More severely, we do not want to fulfill certain kinds of preferences, like those of Nazis.

The classical (i.e. hedonistic) utilitarian can oppose genocide, even if the vast majority prefer it, by appeal to the various measures of happiness, quantity, quality, long-term, etc.

The hedonist can argue that the preferences of sadists should not be counted.

We really do not think that racists and pederasts are happy.

There is psychological data to support this opinion.

But, the preference utilitarian can not argue that the sadist's preferences are not his or her preferences.

There are other versions of utilitarianism which avoid hedonism.

Ideal utilitarianism, for example, argues that we should maximize certain goods, like creative expression.

The right act, for an idealist, is the one which fosters most creative expression for the greatest number.

We proceed to other criticisms of utilitarianism.

VIII. Utilitarianism, Justice, and Rights

Consider a situation in which better consequences arise from performing an act ordinarily considered to be an injustice.

A magistrate or judge is faced with a very real threat from a large and uncontrollable mob of rioters demanding a culprit for a crime. Unless the criminal is produced, promptly tried, and executed, they will take their own bloody revenge on a much smaller and quite vulnerable section of the community (a kind of frenzied pogrom). The judge knows that the real culprit is unknown and that the authorities do not even have a good clue as to who he may be. But he also knows that there is within easy reach a disreputable, thoroughly disliked, and useless man, who, though innocent, could easily be framed so that the mob would be quite convinced that he was guilty and would be pacified if he were promptly executed. Recognizing that he can prevent the occurrence of extensive carnage only by framing some innocent person, the magistrate has him framed, goes through the mockery of a trial, and has him executed (Kai Nielson, "Against Moral Conservatism." In Louis Pojman, *Moral Philosophy: A Reader*, 3rd edition, Hackett 2003, p 150).

Two other cases of utilitarian problems with justice concern a brilliant scientist who murders his wife while developing a cure for cancer and the extreme punishment of parking offenders.

In the scientist's case, the good he will do by furthering his cancer research heavily outweighs the good of incarcerating him.

Publicly executing one or two parking offenders would probably deter most future illegal parking, thus creating great social benefit.

Still, such a punishment seems disproportionate to the crime.

We can summarize these three cases by saying that utilitarianism, which is forward-looking, seems to conflict with justice, which is backward-looking.

The utilitarian account of justice depends on emphasizing the precedent effect.

If you break a promise, you encourage others to break their promises.

Thus, the consequentialist urges us to keep our promises, not because there is something special about making a promise, but because the consequences of breaking that promise are generally worse, in the long run, than the consequences of keeping it.

The precedent of judicial miscreance, in Nielson's case, may have such an overwhelming negative effect that it would not be worth sacrificing the innocent person.

We could probably find other good scientists to take up the murderer's work or arrange a situation in which he could continue his work while incarcerated.

The case of the parking offenders is particularly implausible, on reflection.

Who is really served by such a well-observed law?

Do we really think that Singapore's proscriptions against chewing gum are socially useful?

Utilitarians rely on precedent to deal with other infringements of justice.

Utilitarians generally seek justice, but justice is generally just defended by aggregated utility.

If there are odd cases in which an apparent injustice is licensed, we might have to give up our intuitions about what is the right thing to do.

That is, maybe sometimes sacrificing an innocent is morally acceptable.

Criticisms of utilitarianism based on justice are closely related to ones based on rights.

Consider a beefy spelunker stuck in a cave, blocking the exit for the rest of his group.

The only way out is to use dynamite to blast open the hole.

The dynamite will kill our stuck spelunker and save the rest of the group.

It seems like we should use the dynamite.

But doing so violates the fat man's rights.

We saw that the utilitarian can rely on precedent to accommodate some common-sense moral intuitions about rights and justice.

The utilitarian can argue that she is not ignoring the rights or interests of the fat man, or acting callously.

Utilitarians merely point out that we should also weigh the rights and interests of those who will die unless we kill him.

But, consider the peeping tom, who secretly adds his own happiness to the world's total.

It looks like the utilitarian has to defend the peeper despite the violation of rights.

Justice and rights and utilitarianism are sometimes incompatible.

Utilitarianism requires we look forward, justice that we look backward.

In these cases, either we give up our notions of justice or we give up utilitarianism.

The borderline cases are difficult, and we should not reject a moral theory on the basis of weird cases.

The objections concerning justice and rights are reasonable enough that we might not merely dismiss them.

IX. Utilitarianism Summary

Consider two astronauts on a moonwalk who have an accident which damages their oxygen. They have only enough oxygen for one of them to return to their ship alive. One astronaut sacrifices his life in return for a promise heard by no one except the other astronaut. According to the utilitarian, the surviving astronaut should act as if the promise were never made. The utilitarian defends promise-keeping on the basis of precedents and expectations. You should keep your promises because of the expectations of those to whom you promise, and the precedent set for others who see you break your promises. No one else knows about the astronaut's promise, so breaking it will not create any ill precedents. Any guilt is residual evidence of non-utilitarian presuppositions. But there seems to be something wrong with the assumption that no promise was ever made, even if, in the end, you do break the promise.

We have been trying to show that utilitarianism fails to account for important intuitions: justice, desert, promise-keeping. These problems are all backwards-looking. When presented with such cases, either we give up the theory or the intuitions. But it is not always clear which to cede.

Another way to criticize a theory is to present a preferable alternative. In this vein, we shall examine Kant's ethics.