Analysis of Utilitarianism:
When we put together:

Consequentialism: We are morally required to do what would produce the best outcome.
Aggregation: The outcome is better if the sum of what is good for each person minus what is bad for each person is greater.
Hedonism: What is good for us as an end is pleasure, and what is bad for us as an end is pain.

We get:

Utilitarianism: We are morally required to do what would produce the greatest total of pleasure minus pain.

Problems with Hedonism

Is pleasure the only thing that makes our lives better?
- *Thought experiment*: Your life as an episode of MTV's "Punk'd!"
- In the Punk'd version of your life, you would have the exactly the same sensations of pleasure that your actual achievements and relationships bring you in your actual life.
- Why? Because your sensations of pleasure arise from your *beliefs* about your achievements and relationships, and you would have the same *beliefs* in the Punk'd version. The only difference is that your beliefs would be *false*.
- Don't you hope that your life *isn't* an episode of Punk'd?

  Possible alternative to Hedonism, part 1: What is good for us, what makes our lives go well, is not *only* pleasure. It also depends on:
  - our achieving worthwhile things,
  - having meaningful relationships,
  - leading a life that we freely choose.

Do pleasures always make our lives better?
Compare pleasures from:
- torturing vs. helping people
Reply? We have to take into account how the behavior affects *other people*.
- Should the sadist's pleasure count for *anything* in *our decisions about what to do*?
- Are such pleasures even good *for him*?

  Possible alternative to Hedonism, part 2: Whether pleasures are good for us, and/or whether pleasures should count in our decisions about what to do, depends on the value of what *causes* those pleasures.

Why is hedonism, at least at first, plausible and attractive?
(i)   Pleasure and pain clearly matter.
(ii) When we believe that we have achieved things, or that we have good relationships, this typically produces feelings of pleasure. This may make it seem as if these things are good *only* as means to pleasure. But they might *also* be good as ends.
(iii) Pleasures and pains seem *quantifiable*. Thus, makes possible a determinate decision procedure, a method for resolving conflicts.
(iv) *Avoids many controversies* about which achievements and relationships are objectively worthwhile. Every person is, in effect, his or her own judge.

Problems with Aggregation

Can things be good without being good *for* anyone?
According to:

> *Aggregation*: an outcome is better if the sum of what is good for each of us less what is bad for each of us is greater.

Aggregation says:
- only what is good *for us* matters.
- only the *sum total* of what is good for us matters, not how it is distributed.

Is what is good *for us* the only thing that is *good*, or worth bringing about? What about scientific and cultural achievements, the natural environment?

Does justice—people getting what they deserve—matter?
Is it good as an end

> that *justice* is done
> = that people get what they *deserve*
> = that good people are rewarded and bad people are punished?

In other words: Is it good *period* that what is good *for us* is *distributed* in a certain way: namely, according to desert?

Population growth
- Plausible: Population is at its ideal size when average, or per capita, utility is greatest (or when average utility = marginal utility).
- According to Aggregation, however, population is at its ideal size when *total* utility is greatest (i.e., when marginal utility = 0).

Why should we do this? The population is not itself some being who will experience the increase in total utility. There are only individuals, and each individual's utility will be lower.

Reply:

> *Average Utilitarianism*: An outcome is better if the average of what is good for each person (who exists at the time of our decision?) less what is bad for each person (who exists at the time of our decision?) is greater.

Equal division

Can aggregation explain:

> *Equal Division*: When we have a fixed stock of goods, and a group of people none of whom deserves those goods more than anyone else, we ought to give each member of the group an equal share?

Why should we care how many people derive pleasure from those goods?

A utilitarian reply:

> *Diminishing marginal utility:* The same amount of a good produces more pleasure when given to someone who has *less* of it than when given to someone who has *more* of it.

If diminishing marginal utility is true, and if we have a fixed stock of resources, then we always increase total pleasure by taking from someone who has more and giving it to someone who has less. We maximize total pleasure, therefore, when everyone has an equal amount.

Sacrificing the few and badly off for the many and better off

Many decisions do not involve distributing a fixed stock of goods.

- Example 1: Enslaving an especially productive person.
- Example 2: Persecuting a minority whose practices we, the majority, find displeasing.

Each victim suffers more than each beneficiary benefits. But there are more beneficiaries than victims. So the sum total increases. This can happen even if DMU is true.

The basic problem is this:

> Utilitarianism implies that we are morally required to impose great suffering on a few people who are badly off in order to produce a greater sum of smaller benefits for many other people who are better off.

Problems with consequentialism

When we ask: "Which outcome is best?" rather than "Why outcome is best *for me*?" we ignore which person we are.

Does it demand too much from us?

Construction Case: saving someone from harm by suffering almost as bad a harm

> *Permissive Consequentialism*: We are morally required to produce the best outcome, except when we would have to sacrifice too much.

Does it tell us to harm others?

Transplant Case: harming the one to save the three from *natural* harm.

Does Permissive Consequentialism solve the problem?

- No: It seems that we are not even *permitted* to harm the one to save the three from harm.

Does it solve the problem if *killings* make the outcome worse?
- First, is this plausible? Is it a worse outcome if one person is killed but three are saved from dying from natural causes? *Fire case.*
- Second, a new case:

Railway Case: harming the one to save the three from being harmed *by other people*.
- Now by killing one person, you can prevent three other killings.

This is paradoxical:
- If killing is bad, then why shouldn't we be required, or at least permitted, to kill to prevent more killings?
- The only difference (besides the numbers involved) is who *I* kill, who is harmed *through my agency*.
- Why should it matter whether someone is harmed through my agency or another person's?
- Why should the single victim care whether he is killed through *my* agency? Isn't his complaint simply that he will be killed by someone? And can't each of the three victims make the same complaint: that he will be killed by someone?

Does it tell us to be disloyal to our loved ones?
I do more for my daughter than for any stranger's daughter. But if I did the same for each, then the outcome would be better.
- Am I *permitted* to do more for my daughter?
- Am I *required* to do more for my daughter? Would I be acting wrongly if I did the same for each?

A consequentialist response: "Actually, the outcome will be better if you do more for your daughter. You know her needs better, she'll be hurt if you reject her, etc."
- Must this be true?
- Should I even *think* this way? Do I really *love* my child if I think this way?

The dilemma:
We seem to be left, then, with an agonizing choice between utilitarianism and intuitionism.
- Either we give up our cherished beliefs, in order to have a system.
- Or we cling to our cherished beliefs, and resign ourselves to the mess.