


If you know you are soft-spoken, sit here so I can give you a mic so you don't have to scream!



THE UNIVERSITY
OF BRITISH COLUMBIA

MCLD 2018



CPSC 430

Computers & Society

Dr. Firas Moosvi | 2024_S2

Slides courtesy of Dr. Kevin Leyton-Brown

Class Outline

1. Announcements (10 mins)
2. Kantianism (20 mins)
3. Activity: Companion Bots, Mobile Surveillance (20 mins)
4. Break (10 mins)
5. Utilitarianism (30 mins)
6. Break (10 mins)
7. Rule vs. Act Utilitarianism Activity (30 mins)
8. Reminders before next class (5 mins)

Announcements

Kantianism

2.6 Kantianism

- Key goal: derive morality from more basic principles
- Is anything good regardless of its consequences?
- Immanuel Kant: Only thing in the world that is good without qualification is a good will (desire to do the right thing)
 - other things we might call good (e.g., giving to charity) really depend on consequences
- Reason should cultivate desire to do right thing. Make this precise?



Categorical Imperative (1st Formulation)

Act only from moral rules that you can at the same time will to be universal moral laws.

Illustration of 1st Formulation

- Question: Can a person in dire straits make a promise with the intention of breaking it later?
- Proposed rule: “I may make promises with the intention of later breaking them.”
- The person in trouble wants his promise to be believed so he can get what he needs.
- Universalize rule: Everyone may make & then break promises
- Everyone breaking promises would make promises unbelievable, contradicting desire to have promise believed
- The rule is flawed. The answer is “No.”

Categorical Imperative (2nd Formulation)

Act so that you treat both yourself
and other people as ends in themselves
and never only as a means to an end.

“This is usually an easier formulation to work
with than the first formulation of the
Categorical Imperative.”

...but it depends critically on the “only”.
IMO, it’s “unworkable”.

Second formulation follows from the first

From Wikipedia (“Categorical Imperative”):

*The free will is the source of all rational action. But to treat it as a [means to an end] is to deny the possibility of freedom in general. Because the autonomous will is the one and only source of moral action, **it would contradict the first formulation to claim that a person is merely a means to some other end, rather than always an end in themselves.***

Case for Kantianism

- Rational
- Produces universal moral guidelines
- Treats all people as moral equals

Perfect and Imperfect Duties

- Perfect duty: duty obliged to fulfill without exception
 - Examples: Telling the truth
- Imperfect duties are still duties that can be inferred by the application of “pure reason”: i.e., the first or second formulations of the categorical imperative. But they’re:
 - Activities you couldn’t keep doing forever; never “done”
 - Cause for praise if you do it; not cause for blame if you don’t.
- So what *are* imperfect duties?
 - Examples: helping others; developing your talents.

Perfect and Imperfect Duties

- More generally:
 - Furthering the ends of ourselves and others.
 - Not following maxims that lead to undesirable states of affairs (as distinct from logical contradictions) when universalized
- Not everything we think of as “good” is even an imperfect duty (e.g., doing my chores)

Case Against Kantianism

- Sometimes no rule adequately characterizes an action
- Sometimes there is no way to resolve a conflict between rules
 - In a conflict between a perfect duty and an imperfect duty, perfect duty prevails
 - In a conflict between two perfect duties, no solution
 - Doesn't allow for tradeoffs between moral imperatives
- Kantianism allows no exceptions to perfect duties
- Second formulation of the categorical imperative is really easy to misuse (as, indeed, is Kantianism in general)

Kantianism Activity

Breakout Discussion: Companion bots

Companion robots are being developed to provide care and therapy to elderly patient



Use Kantianism to determine whether or not the use of these robot companions is ethical

Source: [Robots in aged care: a dystopian future?](https://www.ge.com/news/reports/a-robotic-companion-for-the-elderly)

Breakout Discussion: Mobile Surveillance

- “a special investigative team covertly infiltrates the mobile devices of Canadians. The tools, which have been used on at least 10 investigations between 2018 and 2020, give the police access to text messages, email, photos, videos, audio files, calendar entries and financial records. The software can also remotely turn on the camera and microphone of a suspect’s phone or laptop.”
- **Is it ethical for the Canadian government to use this system?**

Consider from the perspective of Kantianism.




<https://bit.ly/3BJUvNX>

Break

Utilitarianism

Principle of Utility

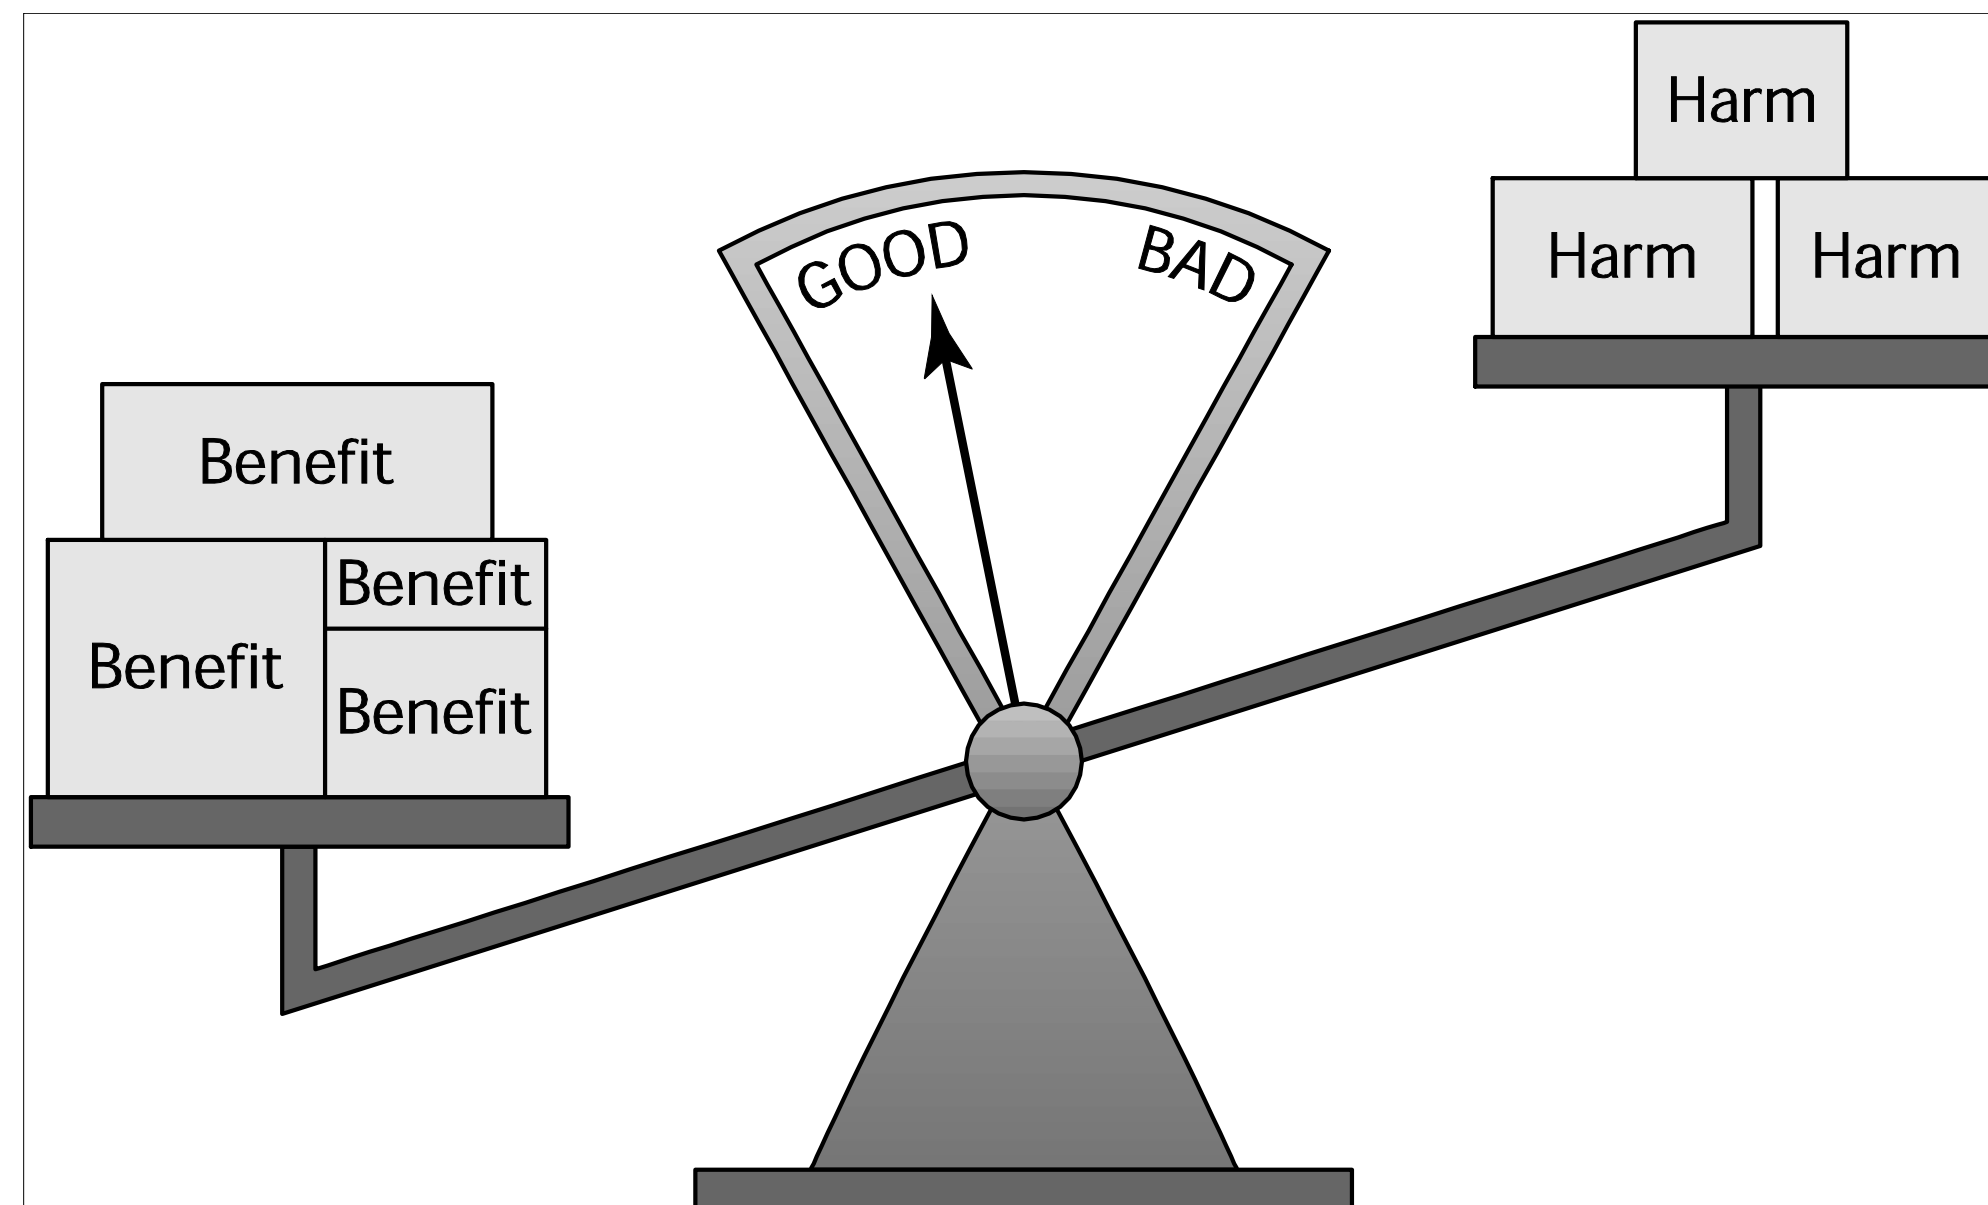
- An action is good if it benefits someone
- An action is bad if it harms someone
- Utility: tendency of an object to produce happiness or prevent unhappiness for an individual or a community
- Happiness = benefit = good = pleasure
- Unhappiness = cost = evil = pain



Jeremy Bentham &
John Stuart Mill

Principle of Utility (“Greatest Happiness Principle”)

An action is right (or wrong) to the extent that it increases (or decreases) the total happiness of the affected parties.



Act Utilitarianism

- Utilitarianism

- Morality of an action has nothing to do with intent
- Focuses on the consequences
 - A “consequentialist” theory

- Act utilitarianism

- Add up change in happiness of all affected beings following from a given action
 - $\text{Sum} > 0$, action is good
 - $\text{Sum} < 0$, action is bad

Discussion

Discuss with the person beside you and report back:

- How useful is act utilitarianism?
- How do we come up with/argue for utility functions?
- How can utilitarianism be useful even if we don't agree on the utility function?

Case for Act Utilitarianism

- Focuses on happiness, which is intuitive
- Down-to-earth (practical)
 - Straightforward to apply
 - Can therefore be helpful in resolving disputes
 - decomposes big issues into lots of little questions
- Comprehensive
 - Allows an agent to trade off different aspects of a situation
 - Contrast with Kantianism: we needed to find one rule

Case Against Act Utilitarianism

- Unclear whom to include in calculations
 - Utilitarians might say you should never exclude anyone...
- Too much work
 - But it's OK to follow a "rule of thumb" most of the time.
- Ignores our innate sense of duty
 - Suppose I make a promise, but can get \$1 for violating it.
 - Seems to miss the sense that I care about my word.
 - Quinn claims: "Note that it does no good for an act utilitarian to ... say that the hard feelings caused by breaking my word to A will have a negative impact on total happiness of $-N$ units, because then all I have to do is change the scenario so that breaking my promise to A enables me to produce $1 + N$ units of good for B. We've arrived at the same result."
 - But is this a problem?
- Susceptible to the problem of moral luck
 - Whether an action is moral depends on outcome, which can depend on circumstances beyond your control

Rule Utilitarianism

- We ought to adopt **moral rules** which, if followed by everyone, will lead to the greatest increase in total happiness
 - Act utilitarianism applies Principle of Utility to individual actions
 - Rule utilitarianism applies Principle of Utility to moral rules

Case for Rule Utilitarianism

- Not every moral decision requires utilitarian calculus
 - You only have to work out the morality of rules.
- Moral rules survive exceptional situations
 - A rule utilitarian can reason (a bit like a Kantian) that it's better for everyone to keep their promises than to lie, and so reject lying for a \$1 gain
- Avoids the problem of moral luck
 - We look at the overall usefulness of the rule, not individual outcomes

Case Against RU, and Utilitarianism in General

- RU: need to identify a single rule to describe situation
- All consequences must be measured on a single scale
 - All units must be the same in order to do the sum
 - In certain circumstances utilitarians must quantify the value of a human life
 - BUT: good arguments from utility theory
 - We have to figure out what this utility function is!
- Utilitarianism ignores problem of unjust distribution
 - Utilitarianism (as defined here) doesn't mean “the greatest good for the greatest number”
 - That requires a principle of justice
 - We can try to combine these ideas
 - However, what happens when a conflict arises between the Principle of Utility and our principle of justice?

Exercise

- In groups of four, identify two ethical issues at the intersection of computers and society:
 - One that is ethical from at least one Utilitarian perspective
 - One that is not ethical from either perspective
- Be prepared to explain your reasoning.

Activity: Rule vs. Act Utilitarianism

Reminders before next class