UNIVERSITÄT BAYREUTH

Collective Obligations

Ethics

Matthias Brinkmann
University of Bayreuth, Summer 2025

2 Exam

- Date: **5 August at 14:00** (sharp!) in **S5 (GWII)**, come **before** 13:45
- Second date (*Nachschreibtermin*) not yet determined
- Further Information on Exam on ELearning
- Mock Exam will be published in last week
- Q&A during last lecture
- **Office Hours**: make appointments via ELearning

Game 1

If you show A, **one piece of candy** is added to the pile. For every two people showing A, another piece of candy is added.

If you show B, **you gain priority** in picking candy.

At the end of the game, everyone gets to pick two pieces of candy, or less if not enough is available.

Game 2

If more than half of the people show A, **all candy** is added to the pile.

If you show B, **you gain priority** in picking candy.

At the end of the game, everyone gets to pick two pieces of candy, or less if not enough is available.

Game 3

Everyone who has already celebrated their birthday should leave the room (i.e., born on or before 1 July).

Congratulations! You can pick as much candy as you want. The leftovers are for the other people when they come back.

6 Today

1. Collective Action Problems
2. The Problem of Imperceptible Harm
3. The Problem of Causal Inefficacy
4. The Principle of Fairness

7 Collective Obligations

‘Collective Obligations’ can have several meanings

- An obligation **of an individual** to act in a certain way in a problem that involves many other people
  - E.g., an individual duty to reduce one’s GHG emissions

  ⟹ Our topic today
- An obligation **of everyone in a group** to act in a certain way
  - E.g., fathers have a duty to think that their child is the cutest
  - But the group here plays no essential role; merely a shorthand way to talk about individual duties
- An obligation **of a group or institution** to act in a certain way
  - E.g., the University of Bayreuth has a duty to provide you with a good learning environment
  - E.g., Germany has a duty to abide by international law
  - The institution or group itself has a duty here; and this duty might not be reducible to individual duties!

Collective Action Problems
and their Moral Significance

9 Prisoner's Dilemma

	Cooperate	Defect
Cooperate	4,4	-2,6
Defect	6,-2	0,0

In a Prisoner's Dilemma, the incentives are such that

- It would be **collectively better** (i.e., in terms of total welfare) if everyone chose to cooperate
- But it is **individually rational** to make a choice that leads to collectively inferior results

10 Prisoner's Dilemma

	Cooperate	Defect
Cooperate	4,4	-2,6
Defect	6,-2	0,0

Different ways how we might 'solve' the Prisoner's Dilemma can be imagined

- Repeated Interaction (esp. punishment of non-compliers)
- Reputational Mechanisms (esp. in small groups)
- External Incentives (e.g., punishment for noncompliance)
- Internalised Moral Obligations

What if the Prisoner's Dilemma happens on a very large scale between anonymous players?

11 **Collective Action Problem** with Linear Structure

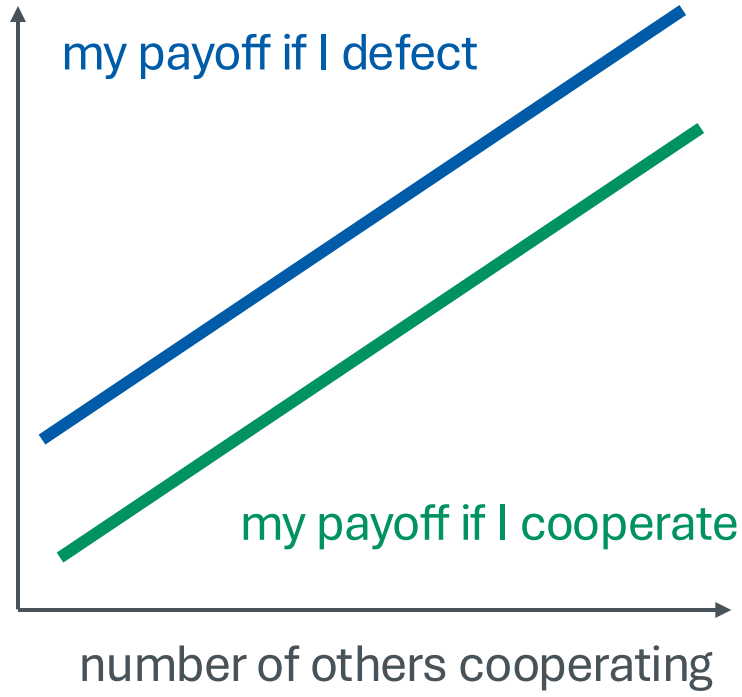

Imagine that for 20000 players, cooperating gives **everyone** a payoff of 1; defecting gives **only me** a payoff of 3

- It is collectively optimal if everyone cooperates (resulting welfare for everyone: 20000)
- But it is individually rational for me to defect (resulting welfare for everyone: 3)

# others cooperating	0	1	n	19998	19999
my payoff for cooperate	1	2	n + 1	19999	20000
my payoff for defect	3	4	n + 3	20001	20002

12 Collective Action Problem

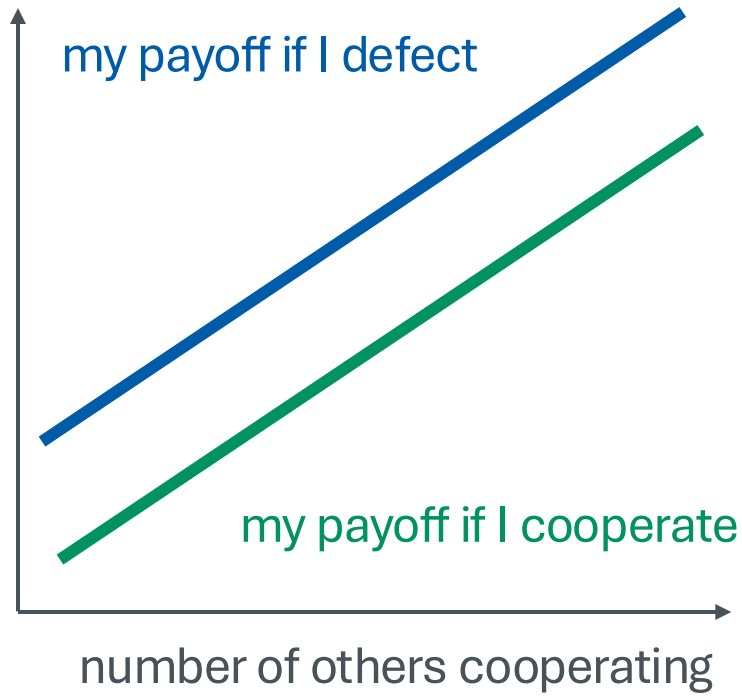

In a collective action problem, a **public good** is produced (e.g., peace, fireworks, clean air, herd immunity, social trust, dressing well)

Subtype: collective harm problem

- A public good is **non-rivalrous in consumption**: if I consume the good, I do not diminish the good for others
- A public good is **non-excludable**: one cannot exclude others from enjoying the good

Collective action problems incentivise **free-riding**: people who enjoy the public good but do not contribute to its production

13 **Collective Action Problem** with Tipping Point Structure

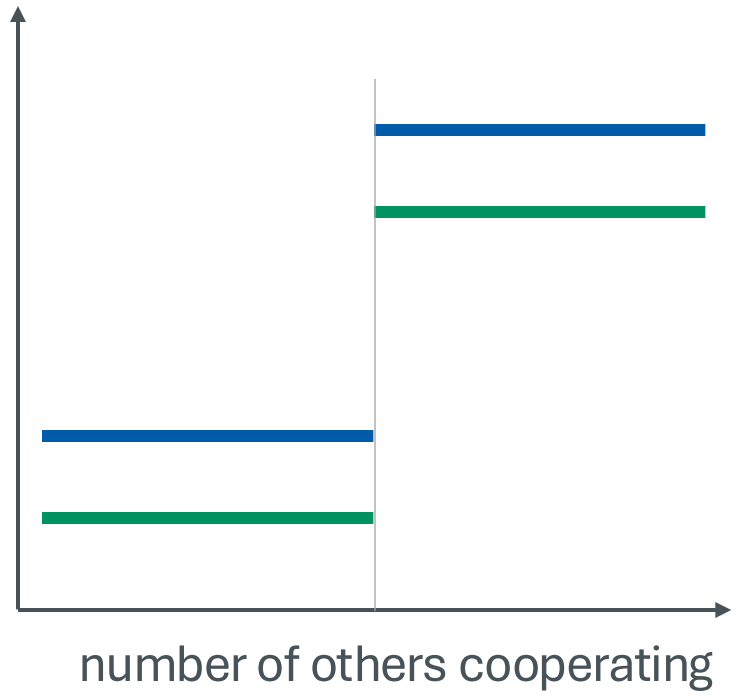

my payoff if I defect
my payoff if I cooperate

Imagine that there are 20000 players. If 10001 players cooperate, **everyone** receives a benefit of 5000; defecting always gives **me** a benefit of 100

- It would be **collectively optimal** if 10001 persons cooperated
- It is **individually rational** for me to defect, **except** if I can 'tip' the result

# others cooperating	…	9999	10000	10001	…
my payoff for cooperate	0	0	5000	5000	5000
my payoff for defect	100	100	100	5100	5100

14 **Collective Action Problem** with Step Ladder Structure

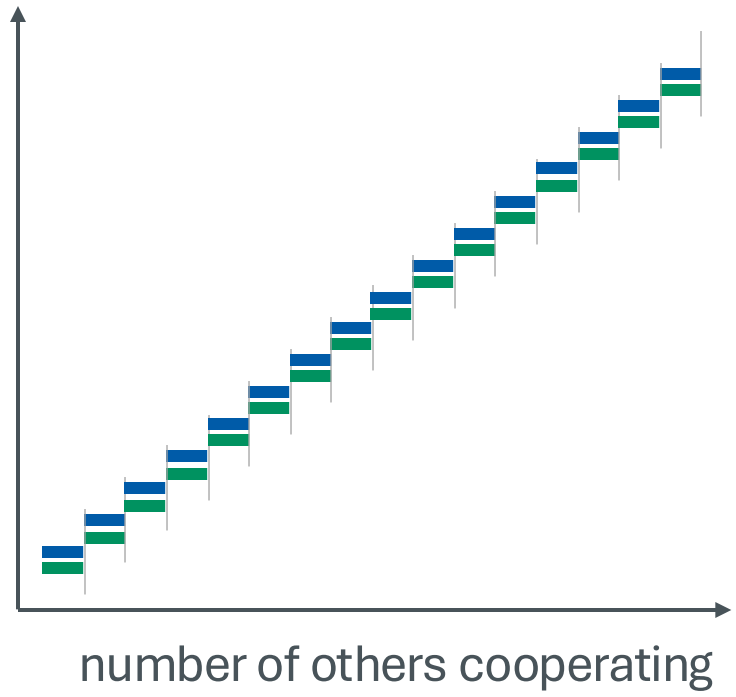

my payoff if I defect
my payoff if I cooperate

Imagine that there are 20000 players. For each 1000 players who cooperate, **everyone** receives a benefit of 100; defecting always gives **me** a benefit of 3

- It would be collectively optimal if everyone cooperated
- It is individually rational for me to defect, **except** if I can bring about the next 'step'

# others cooperating	…	4998	4999	5000	…
my payoff for cooperate	400	400	500	500	500
my payoff for defect	403	403	403	503	503

Question

Which real-world collective action problems have which of these structures (linear, tipping-point, stepladder)?

16 Real-World Collective Action Problems

- Linear Structure: Climate Change
  - I can reduce my greenhouse gas emissions (cooperate). This reduces overall greenhouse gas emissions by a tiny amount (linear change)
  - But I have an incentive to not reduce my emissions (defect)
- Tipping Point Structure: Voting (in a winner-takes-all system)
  - I can vote for the candidate of my party (cooperate). If the candidate receives one more vote than the next-best candidate, s/he wins (tipping point)
  - I have an incentive to not vote and stay home (defect), unless I know that I make the difference
- Step Ladder Structure: Meat Consumption
  - I can reduce my meat consumption (cooperate). The meat industry adjusts its meat production, but will only reduce production if there is a discernible reduction in demand (step ladder)
  - I have an incentive to continue eating meat (defect), unless I know that I reduce meat consumption enough for the meat industry to reduce its production

The Problem of **Imperceptible Harm**

18 Imperceptible Harm

TORTURERS. A thousand victims are tied to electric chairs. A thousand torturers each can press one button. Each button press increases the voltage of each electric chair by one volt. If a thousand buttons are pressed, the victims are in immense pain. But the victims cannot tell the difference if one additional button is pressed (after Parfit 1984, 80)

- Structure (let '$=_F$' stand for 'feels the same to the victim')
  - 0 volt $=_F$ 1 volt; 1 volt $=_F$ 2 volt; 2 volt $=_F$ 3 volt; …; but 0 volt $\neq_F$ 1000 volt;
  - In technical language: the '$=_F$' relation is not transitive
- This is a linear case: each contribution makes a causal difference
- However, the difference is **not perceptible** (more broadly: not morally significant)

19 The Analogy with Climate Change

- GHG emission per capita in Germany were ca. 9 tonnes of CO2 equivalent (2022)
  - Global GHG emissions are ca. 53 billion tonnes of CO2 equivalent (2022)
  - Thus: the average German contributes ca. 0.000000017% of global GHG emissions
- Does it make a difference if there are more GHG emissions to this tiny amount? Broome: yes. $\Longrightarrow$ That is, we are probably in a linear case!
- Thus: our individual GHG emission **do make** a difference, just a difference that nobody can **perceive**

20 The Argument Against Imperceptible Harms

1. Whether an action is right or wrong depends only on its consequences considered in isolation, holding what everyone else does fixed.
2. For an action to be wrong, it needs to harm someone.
3. For an action to harm someone, it needs to be possible for that person to experience the harm.
4. In TORTURERS, the consequence of pressing the button, considered in isolation, does not lead to a perceptible harm to anyone.
   Similarly: If I emit GHG gases, this does not lead to a perceptible harm to anyone.
5. Thus, it is not wrong to press the button in TORTURERS.
   Similarly: It is not wrong for me to emit GHG gases/not reduce my GHG emissions.

Questions

Should we accept the conclusion of the argument?

If not, which premise should we reject?

22 Escaping the Argument Against Imperceptible Harms

1. Whether an action is right or wrong depends only on its consequences considered in isolation, holding what everyone else does fixed.
2. For an action to be wrong, it needs to harm someone.
3. For an action to harm someone, it needs to be possible for that person to experience the harm.
4. In TORTURERS, the consequence of pressing the button, considered in isolation, does not lead to a perceptible harm to anyone.
5. Thus, it is not wrong to press the button in TORTURERS.

- Reject premise 1: an action can be wrong beyond the consequences it has in isolation.
- Reject premise 2: an action can be wrong even if does not harm anyone ($\Longrightarrow$ Principle of Fairness, discussed later)
- Reject premise 3: it is possible to harm someone even if that person does not experience the harm.

23 Rejecting the Isolation Premise

"Even if an act harms no one, this act may be wrong because it is one of a set of acts that together harm other people." (Parfit, 1984, 70)

- One motivation: cases of overdetermination.
  - Imagine that both A and B shoot at C at exactly the same time. If A would not have shot, then C would have died any way; the same for B. So it seems that neither A or B make a difference!
- A non-trivial problem: finding the relevant 'set of acts'
  - Take the set {me lecturing; X killing Y in cold blood}
  - Or {me shouting for people to quiet down; others being loud in an inconsiderate way}
  - Or {poor country increasing from very low to low GHG emissions; rich country increasing from high to very high GHG emissions}

24 In Favour of Imperceptible Harms

1. If you do not reduce your climate emissions, then your yearly climate emissions lead to a loss of half a year of life, spread across billions of people (estimate).
2. There is no moral difference between making one person lose half a year of their life, and making billions of people lose the aggregate total of half a year of life.
3. It is wrong to make one person lose half a year of their life.
4. Thus, it is wrong for you not to reduce your climate emissions.

- Premise 1 is difficult to establish and estimate.
- Should we accept premise 2? E.g., my terrible fashion choices might cause lots of people a small amount of anguish, which in total is quite a lot suffering.

The Problem of Causal Inefficacy

26 The Problem of Causal Inefficacy

Consider now cases with a tipping-point or step-ladder structure. In such cases collective action problems with large numbers of participants, your individual choice does often not make **any difference**:

- If I stop eating meat, the meat industry will almost certainly not notice and the same amount of animals will be slaughtered (step-ladder)
- Whether I vote in election or not, the same candidate or party will win (tipping-point)

27 The Problem of Causal Inefficacy

In argument form:

1. Eating this chicken makes no difference: if I eat it, I cause no additional suffering, and if I do not eat it, I do not prevent any suffering
2. For an action to be morally wrong, it must make a (relevant) difference

Therefore,

3. Eating this chicken is not morally wrong

Questions

Is anything wrong with this argument?

Should either of the premises be rejected?

29 Cheering on my Sports Club

24000 people in the Yellow Wall already are cheering for Dortmund. Why should I?

- **Direct Instrumental Reason**: cheering motivates my team, making victory more likely
- **Indirect Instrumental Reason**: by cheering, others recognise me as a fan
- **Moral Integrity**: if others do their part and cheer, then I should do so too
- **Symbolism**: I'm a Dortmund fan, these are my players, I wish to show support

30 Responses to the Problem of Causal Inefficacy

1. Eating this chicken makes no difference: if I eat it, I cause no additional suffering, and if I do not eat it, I do not prevent any suffering
2. For an action to be morally wrong, it must make a (relevant) difference

Therefore,

3. Eating this chicken is not morally wrong

A. Accept both premises, and thus accept the conclusion
B. Deny premise 1: we **do** make a difference in the relevant cases
C. Deny premise 2: we do not need to make a difference to act wrongly $\Longrightarrow$ Principle of Fairness, discussed later

Response A: Accepting the Conclusion

- Perhaps it is only the responsibility of governments or other powerful agents to act
- But the problem can be recreated on the larger level
  - E.g., Germany unilaterally reducing its GHG emissions might only have a marginal effect (Germany GHG emissions = ca. 2% of global total)
- Powerful agents must also first be created, and here we face a similar collective action problem
- Powerful agents might fail in their duties
  - E.g., if Germany fails to reduce its GHG emissions, then does the duty not ‘fall back’ on me?

32 Response B: Expected Impact (Kagan, Norcross, Singer)

- Idea: reject premise 1. Our actions do have an impact—an **expectable** impact
- Toy Example
  - Assume: an animal farm increases or decreases the number of animals produced to demand only in terms of hundreds (it raises 47700, or 47800, or 47900, … chicken)
  - If I do not eat a chicken, there is a chance I hit the relevant threshold: assume that there is a 1/100 chance that the farm will reduce its chicken production by 100
  - My **expected impact** of not eating a chicken is (1/100)*100 = 1 chicken not being produced (even if my **actual impact** is 0)
- (The actual production function of factory farming is likely more complex!)

33 The Expected Impact Argument

1. Eating this chicken makes an **expected difference**—i.e., on statistical average, it kills one chicken
2. If an action makes a negative expected difference, then it is morally wrong—specifically, it is wrong to act in a way that on statistical average kills one chicken

Therefore,

3. Eating this chicken is morally wrong

Problem: premise 2 seems to be **excessively risk-averse**

Example: any time you drive a car, there is a small chance you injure or kill (there is an expected but not actual harm you do to others); but driving a car is not (normally) morally wrong

34 Overview

Causal Structure of the Collective Action Problem

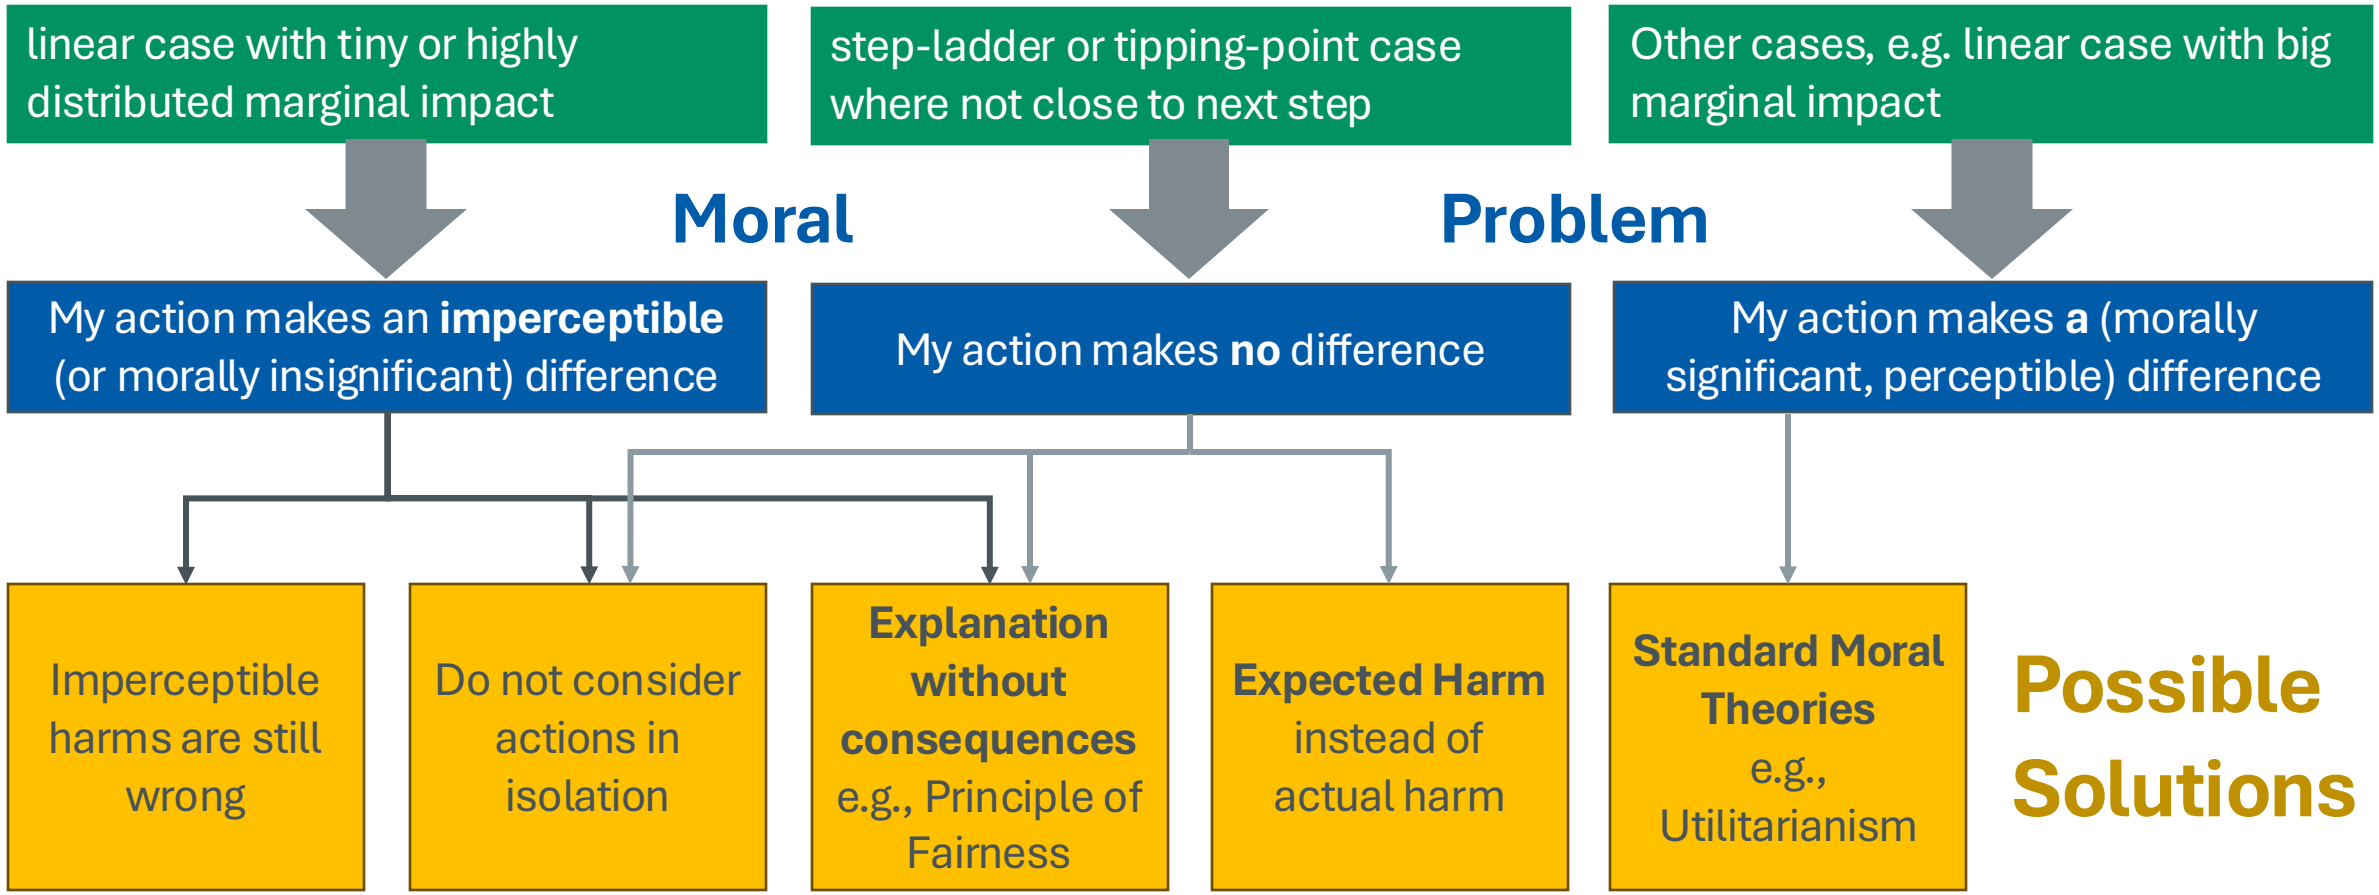

Questions?

The Principle of Fairness

37 Fairness Principle

Hart's Principle. 'when a number of persons conduct any joint enterprise according to rules and thus restrict their liberty, those who have submitted to these restrictions when required have a right to a similar submission from those who have benefited by their submission.' (Hart 1955, 185)

- Put in our language: **Benefit Principle**. Whenever I **benefit** from a public good, then I I have a duty to contribute to its production
- One application: I benefit from the benefits of the state, and thus I have a duty to contribute to its continued existence (e.g., by paying taxes)

38 Objections to the Benefit Principle

Benefit Principle. Whenever I benefit from a public good, then I also have a duty to contribute to its production

- The principle seems **incomplete** in important respects. It does not consider
  - the costs to me, and whether the costs are bigger than the benefits to me
  - whether the costs are fairly distributed
  - whether benefits are equally distributed

  $\Longrightarrow$ we might be able to fix these problems with a more precise formulation
- Nozick's counterexample: my neighbourhood has a system where books are brought every week to everyone's doorstep. People put books on my doorstep which I am happy to read. Do I now have a duty to help finance the book-sharing scheme?

39 Tutorials / Next Week

- **Primary Text** this week: Lomasky, Loren, and Geoffrey Brennan. 2000. ‘Is There a Duty to Vote?’ *Social Philosophy and Policy* 17 (1): 62–86.
- **Secondary Text** this week: Nefsky, Julia. 2019. ‘Collective Harm and the Inefficacy Problem’. *Philosophy Compass* 14 (4).

- **Next Week**: we start with metaethics
  - The topic is Subjectivism: Is morality subjective? What would it mean for morality to be subjective?