

Rationality, Morality and Economics

Topic 1, Lecture 2

Decisions Under Ignorance

Rob Trueman
rob.trueman@york.ac.uk

University of York

Decisions Under Ignorance

Re-Cap

Milnor's Axioms

The Veil of Ignorance

Milnor's Axioms (Ethical Remix)

Acts, States and Outcomes

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 1 | 10 |
| a_2 | 10 | 1 |

Acts, States and Outcomes

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 1 | 10 |
| a_2 | 10 | 1 |

- **Acts:** the things you can do in the decision problem

Acts, States and Outcomes

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 1 | 10 |
| a_2 | 10 | 1 |

- **Acts:** the things you can do in the decision problem
- **States:** ways the world might be, which you don't control

Acts, States and Outcomes

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 1 | 10 |
| a_2 | 10 | 1 |

- **Acts:** the things you can do in the decision problem
- **States:** ways the world might be, which you don't control
- **Outcomes:** what you get when you combine an act and a state; the numbers represent an agent's preferences

Risk versus Ignorance

- **Decisions under risk**
 - You know the possible outcomes of each act, and you can assign subjective probabilities to those outcomes

- **Decisions under ignorance**
 - You know the possible outcomes of each act, but you cannot assign subjective probabilities to those outcomes

The Principle of Maximising Expected Utility

- The standard principle for decision making is the **Principle of Maximising Expected Utility**
 - You should choose the act with the greatest expected utility
 - $EU(a) = \sum_{i=1}^n [P(s_i) \times U(a \wedge s_i)]$
- But you cannot apply this principle to decisions under ignorance, because we don't have the probabilities to calculate expected utilities!

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- Action a_i is rationally preferable to action a_j iff the worst-case outcome obtainable by performing a_i is better than the worst-case outcome obtainable by performing a_j

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$

- **The Leximin Rule**

- Action a_i is rationally preferable to action a_j iff the worst-case outcome obtainable by performing a_i is better than the worst-case outcome obtainable by performing a_j , or if a_i and a_j are tied at the worst-case outcome, then the second-worst outcome obtainable by performing a_i is better than the second-worst outcome obtainable by performing a_j , or if a_i and a_j are tied at the worst case, then the third-worst...

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$

- **The Leximin Rule**

- $a_i \succ a_j$ if and only if there is some positive integer n such that $\min^n(a_i) > \min^n(a_j)$, and $\min^m(a_i) = \min^m(a_j)$ for all $m < n$

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$

- **The Leximin Rule**

- $a_i \succ a_j$ if and only if there is some positive integer n such that $\min^n(a_i) > \min^n(a_j)$, and $\min^m(a_i) = \min^m(a_j)$ for all $m < n$

- **The Laplace Rule**

- If you have no reason to think that one state is any more or less probable than any other state, then assign every state equal probabilities, and then apply the Principle of Maximising Expected Utility

Three Ways of Coping with Ignorance

- **The Maximin Rule**

- $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$

- **The Leximin Rule**

- $a_i \succ a_j$ if and only if there is some positive integer n such that $\min^n(a_i) > \min^n(a_j)$, and $\min^m(a_i) = \min^m(a_j)$ for all $m < n$

- **The Laplace Rule**

- $a_i \succ a_j$ if and only if $\sum_{x=1}^n \frac{1}{n} u(a_i, s_x) > \sum_{x=1}^n \frac{1}{n} u(a_j, s_x)$

Objections

- Maximin is inconsistent with the **Strong Dominance Principle**:
 - $a_i \succ a_j$ if: $u(a_i, s) \geq u(a_j, s)$ for all states s , and there is some state s' such that $u(a_i, s') > u(a_j, s')$
- Both Leximin and the Laplace Rule are very sensitive to how we *represent* decision problems
 - These rules recommend different acts when we change the way that we are dividing up the possible states of nature

Decisions Under Ignorance

Re-Cap

Milnor's Axioms

The Veil of Ignorance

Milnor's Axioms (Ethical Remix)

An Axiomatic Approach

- We can propose some basic axioms about how rational preferences should behave
- The different decision principles are then characterised by different sets of axioms
- This approach was taken by John Willard Milnor (1954), who presented the characteristic axioms for a number of different decision rules
 - See Box 3.2 in Peterson's textbook for a complete summary of Milnor's results
- To keep things simple, we will focus only on Maximin and the Laplace Rule

Axioms for Maximin and Laplace

| | Maximin | Laplace |
|-------------------------|---------|---------|
| Ordering | ✓ | ✓ |
| Symmetry | ✓ | ✓ |
| Strict Dominance | ✓ | ✓ |
| Continuity | ✓ | |
| Irrelevant Alternatives | ✓ | ✓ |
| Column Linearity | ✗ | ✓ |
| Column Duplication | ✓ | ✗ |
| Randomization | ✓ | |

The axioms marked ✓ are individually necessary and jointly sufficient for the relevant decision rule

The axioms marked ✗ are incompatible with the relevant decision rule

The Uncontroversial Axioms

- **Ordering:** \succeq is transitive and complete
 - *Transitive:* for any acts a , b and c , if $a \succeq b$ and $b \succeq c$, then $a \succeq c$
 - *Complete:* for any acts a and b , either $a \succeq b$ or $b \succeq a$ (or both)
- **Symmetry:** Preferences over acts do not depend on how the acts and states of nature have been labelled
- **Strict Dominance:** $a_i \succ a_j$ if: $u(a_i, s) > u(a_j, s)$, for all states s

The Varieties of Dominance

- The first three axioms are uncontroversial because they are all implied by both Maximin and Laplace (and any other decent decision rule too!)
- *Hold on! You said that Maximin is incompatible with Dominance!*
- Maximin is incompatible with *Strong* Dominance, but it implies *Strict* Dominance
 - **Strict:** $a_i \succ a_j$ if: $u(a_i, s) > u(a_j, s)$, for all states s
 - **Strong:** $a_i \succ a_j$ if: $u(a_i, s) \geq u(a_j, s)$ for all states s , and there is some state s' such that $u(a_i, s') > u(a_j, s')$
 - **Weak:** $a_i \succeq a_j$ if: $u(a_i, s) \geq u(a_j, s)$, for all states s

The Axioms of the Laplace Rule

- We get the Laplace Rule if we add the following take the uncontroversial axioms, and then add two more
- **Irrelevant Alternatives:** If $a \succeq b$, then this will not change if we add new acts to choose from
- **Column Linearity:** If $a \succeq b$, then this will not change if the utilities for all outcomes in one state are increased by the same amount

Maximin and Column Linearity

- Maximin is incompatible with Column Linearity

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 0 | 7 |
| a_2 | 5 | 6 |

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 10 | 7 |
| a_2 | 15 | 6 |

- In the left decision problem Maximin implies $a_2 \succeq a_1$, but in the right problem Maximin implies $a_2 \not\succeq a_1$

Laplace and Column Duplication

- However, the Laplace Rule is also incompatible with one of the characteristic axioms of Maximin
- **Column Duplication:** If $a \succeq b$, then this will not change if we duplicate one of the states of nature

| | s_1 | s_2 |
|-------|-------|-------|
| a_1 | 12 | 0 |
| a_2 | 5 | 5 |

| | s_1 | s_2 | s_3 |
|-------|-------|-------|-------|
| a_1 | 12 | 0 | 0 |
| a_2 | 5 | 5 | 5 |

- In the left problem Laplace implies $a_1 \succeq a_2$, but in the right problem Laplace implies $a_1 \not\succeq a_2$

Choosing between Maximin and Laplace

- Choosing between Maximin and the Laplace Rule is, then, partly a matter of choosing which of these two axioms you prefer
 - **Column Linearity:** If $a \succeq b$, then this will not change if the utilities for all outcomes in one state are increased by the same amount
 - **Column Duplication:** If $a \succeq b$, then this will not change if we duplicate one of the states of nature
- For now, I will leave it to you to think over which you find most compelling, and we can talk about it during the seminar!

Decisions Under Ignorance

Re-Cap


Milnor's Axioms

The Veil of Ignorance

Milnor's Axioms (Ethical Remix)

From Decision Theory to Ethics

- John Harsanyi and John Rawls have both argued that we can use decision theory to tell us something important about ethics
- They both think that we can use principles governing decisions under ignorance to figure out what kind of society would be *just*
- However, they have *very* different ideas about what a just society would actually look like!



vs.



Choosing between Societies

- Decision problems are usually conceived of as problems about what act to perform
- But we can re-interpret our decision problems without changing their structure in any way
- For example, we could think of a decision problem as a problem about what kind of society we would rather live in

Choosing between Societies

| | Huey | Dewey | Louie |
|----------|------|-------|-------|
| <i>A</i> | 4 | 5 | 3 |
| <i>B</i> | 2 | 2 | 2 |
| <i>C</i> | 1 | 5 | 9 |

Choosing between Societies

| | Huey | Dewey | Louie |
|----------|------|-------|-------|
| <i>A</i> | 4 | 5 | 3 |
| <i>B</i> | 2 | 2 | 2 |
| <i>C</i> | 1 | 5 | 9 |

- **Kinds of society:** the alternatives that we choose from

Choosing between Societies

| | Huey | Dewey | Louie |
|----------|------|-------|-------|
| <i>A</i> | 4 | 5 | 3 |
| <i>B</i> | 2 | 2 | 2 |
| <i>C</i> | 1 | 5 | 9 |

- **Kinds of society:** the alternatives that we choose from
- **People:** the members of the society

Choosing between Societies

| | Huey | Dewey | Louie |
|----------|------|-------|-------|
| <i>A</i> | 4 | 5 | 3 |
| <i>B</i> | 2 | 2 | 2 |
| <i>C</i> | 1 | 5 | 9 |

- **Kinds of society:** the alternatives that we choose from
- **People:** the members of the society
- **Utilities:** these numbers measure the well-being of each person in each kind of society

Which Society?

| | Huey | Dewey | Louie |
|---|------|-------|-------|
| A | 4 | 5 | 3 |
| B | 2 | 2 | 2 |
| C | 1 | 5 | 9 |

- Imagine we ask Huey which society he would rather live in
- This is not an inherently ethical question: Huey might choose A, just because his well-being is at its highest in A
- However, Harsanyi (1953) and Rawls (1971, 1974) both think that we can make it an ethical question by restricting the type of information we have access to

The Veil of Ignorance

- Imagine Huey *forgot* who he was: he knew he was one of the people in the society being discussed, but he didn't know *which* person he was
 - Rawls described this as putting Huey behind a “veil of ignorance”
- If we now asked Huey which kind of society he would prefer, he couldn't just pick the one that works out best for him, because he doesn't know which society that is!
- Harsanyi and Rawls both think that the society Huey would choose to live in from behind the veil of ignorance is the society which he thinks is most just

Two Assumptions of this Approach

- This approach to social justice makes two big assumptions
 - **Welfarism:** whether one kind of society is more just than another is entirely determined by the utilities of the people in that society
 - **Comparability:** it makes sense to compare the utilities of two different people
- Neither of these assumptions are uncontroversial, but we will accept them for now, just for the sake of argument

Harsanyi and the Laplace Rule

- The question now is what decision rule we should use from behind the veil of ignorance
- Harsanyi (1953) thinks we should use a version of the Laplace Rule
- You might be anyone in the society, and you should assign each possibility exactly the same probability; then you should just maximise your expected utility

Applying the Laplace Rule

| | Huey | Dewey | Louie |
|---|------|-------|-------|
| A | 4 | 5 | 3 |
| B | 2 | 2 | 2 |
| C | 1 | 5 | 9 |

- $EU(A) = \frac{4}{3} + \frac{5}{3} + \frac{3}{3} = 4$
- $EU(B) = \frac{2}{3} + \frac{2}{3} + \frac{2}{3} = 2$
- $EU(C) = \frac{1}{3} + \frac{5}{3} + \frac{9}{3} = 5$

MOST JUST

From Laplace to Utilitarianism

- As Harsanyi points out, using the Laplace Rule immediately yields a version of **utilitarianism**:
 - The most just society is the one which maximises the average utility of its members
- And if we assume that the population of the society is held fixed, we get an even simpler version of utilitarianism:
 - The most just society is the one which maximises the total utility of its members
- So if you're trying to pick the most just society, just add up the utilities of all its members!

Rawls against Utilitarianism

- Rawls (1971: esp. §28) was a staunch critic of utilitarianism, and of the attempt to establish it via the Laplace Rule
- In a utilitarian society, it is acceptable to greatly reduce one person's well-being, so long as doing so increases the overall utility of the society
- Rawls thought that if someone had to pick a society from behind the veil of ignorance, then they would not want to risk being the person whose utility is reduced for the greater good
- So they would not use the Laplace Rule, which opens up that possibility

Rawls and Maximin

- Rawls instead argues that people would use Maximin to choose a society from behind the veil of ignorance
 - The most just society is the society which maximises the utility of its worst off member
 - [This is a bit of a simplification: Rawls (1971: p. 152) claims that his use of Maximin is just a “heuristic”]
- By using Maximin, you ensure that the worst position in society is the best it can be, and so ensure that whatever position you end up occupying, it won't be *too* bad

Applying Maximin

| | Huey | Dewey | Louie |
|----------|------|-------|-------|
| <i>A</i> | 4 | 5 | 3 |
| <i>B</i> | 2 | 2 | 2 |
| <i>C</i> | 1 | 5 | 9 |

- The most just kind of society is *A*, because its worst position is better than the worst positions of *B* and *C*
- The inequalities between Huey, Dewey and Louie in *A* are tolerable because they somehow serve to maximise the utility of the worst position
- Even though *C* has a much higher total utility, it is unjust because the inequalities do not serve to maximise the utility of the worst off

An Added Benefit of Maximin

- Rawls (1974: 143–4) points out that one benefit of using Maximin instead of Laplace is that we need a lot less information
 - To use Laplace/utilitarianism, we need to know if the utility gains of one group of people outweigh the utility losses of another
 - To use Maximin, we only need to identify the worst off in each society, and figure out which of those has the highest utility
- This is because the Laplace Rule requires an *interval* utility scale, whereas Maximin only requires an *ordinal* scale
 - **Ordinal scale:** $u(a) > u(b)$ iff the utility of a is higher than the utility of b
 - **Interval scale:** $u(a) - u(b) > u(c) - u(d)$ iff a is preferred to b more than c is preferred to d

Decisions Under Ignorance

Re-Cap

Milnor's Axioms

The Veil of Ignorance

Milnor's Axioms (Ethical Remix)

An Axiomatic Approach

- Harsanyi and Rawls employ ethical versions of Laplace and Maximin
- Milnor discovered the characteristic axioms of Laplace and Maximin
- We can convert these axioms into ethical axioms, and use them to characterise the ethical versions of Laplace and Maximin

Re-Interpreting \succeq

- In the original axioms, $a \succeq b$ meant that performing act a is at least as rational as performing act b
- Now, $A \succeq B$ will mean that choosing society A from behind the veil of ignorance is at least as rational as choosing society B from behind the veil of ignorance
- If Harsanyi and Rawls are right about the link between justice and the veil of ignorance, then we can also read $A \succeq B$ as saying that society A is at least as just as society B

The Uncontroversial Axioms

- **Ordering:** \succeq is transitive and complete
 - *Transitive:* for any a , b and c , if $a \succeq b$ and $b \succeq c$, then $a \succeq c$
 - *Complete:* for any a and b , either $a \succeq b$ or $b \succeq a$ (or both)
- **Symmetry:** Preferences over acts do not depend on how the acts and states of nature have been labelled
- **Strict Dominance:** $a_i \succ a_j$ if: $u(a_i, s) > u(a_j, s)$, for all states s

The Uncontroversial Axioms

- **Ordering:** \succeq is transitive and complete
 - *Transitive:* for any A , B and C , if $A \succeq B$ and $B \succeq C$, then $A \succeq C$
 - *Complete:* for any A and B , either $A \succeq B$ or $B \succeq A$ (or both)
- **Symmetry:** Preferences over acts do not depend on how the acts and states of nature have been labelled
- **Strict Dominance:** $a_i \succ a_j$ if: $u(a_i, s) > u(a_j, s)$, for all states s

The Uncontroversial Axioms

- **Ordering:** \succeq is transitive and complete
 - *Transitive:* for any A , B and C , if $A \succeq B$ and $B \succeq C$, then $A \succeq C$
 - *Complete:* for any A and B , either $A \succeq B$ or $B \succeq A$ (or both)
- **Impartiality:** If two societies differ only in that some people have switched their utilities, then the two societies are equally just
- **Strict Dominance:** $a_i \succ a_j$ if: $u(a_i, s) > u(a_j, s)$, for all states s

The Uncontroversial Axioms

- **Ordering:** \succeq is transitive and complete
 - *Transitive:* for any A , B and C , if $A \succeq B$ and $B \succeq C$, then $A \succeq C$
 - *Complete:* for any A and B , either $A \succeq B$ or $B \succeq A$ (or both)
- **Impartiality:** If two societies differ only in that some people have switched their utilities, then the two societies are equally just
- **Strict Pareto:** $A \succ B$ if: the utility of each person in A is strictly better than their utility in B

The Axioms of Utilitarianism

- We get utilitarianism if we add the following two axioms
- **Ethically Irrelevant Alternatives:** If $A \succeq B$, then this will not change if we add new societies to choose from
- **Ethical Column Linearity:** If $A \succeq B$, then this will not change if the utilities for every person in one society are increased by the same amount

Choosing between Maximin and Laplace

- The ethical version of Maximin is incompatible with Ethical Column Linearity

| | Huey | Dewey |
|---|------|-------|
| A | 0 | 7 |
| B | 5 | 6 |

| | Hewey | Dewey |
|---|-------|-------|
| A | 10 | 7 |
| B | 15 | 6 |

- In the left table Maximin implies $B \succeq A$, but in the right table Maximin implies $B \not\succeq A$

Utilitarianism and Ethical Column Duplication

- However, the utilitarianism is also incompatible with one of the characteristic axioms of the ethical version of Maximin
- **Ethical Column Duplication:** If $A \succeq B$, then this will not change if we duplicate one of the people

| | Hewey | Dewey |
|---|-------|-------|
| A | 12 | 0 |
| B | 5 | 5 |

| | Hewey | Dewey | Louie |
|---|-------|-------|-------|
| A | 12 | 0 | 0 |
| B | 5 | 5 | 5 |

- In the left table utilitarianism implies $A \succeq B$, but in the right table utilitarianism implies $A \not\succeq B$

Choosing between Maximin and Utilitarianism

- Choosing between Rawl's Maximin conception of justice and Harsanyi's utilitarianism is partly a matter of choosing which of these two axioms you prefer
 - **Ethical Column Linearity:** If $A \succeq B$, then this will not change if the utilities for every person in one society are increased by the same amount
 - **Ethical Column Duplication:** If $A \succeq B$, then this will not change if we duplicate one of the people
- For now, I will leave it to you to think over which you find most compelling, and we can talk about it during the seminar!

For the Seminar

- Please read:
 - Peterson, *An Introduction to Decision Theory*, ch. 3
 - Binmore, *Rational Decisions*, §9.1
- Please also complete the exercises at the end of the Peterson chapter
- Please also bring notes on which of Milnor's axioms you find compelling, and which you find doubtful. Think about them in both their original and their ethical flavours

References

- Harsanyi, John C. (1953) 'Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking', *The Journal of Political Economy* 61 (5): 434–435
- Milnor, John (1954) 'Games against Nature', in Thrall et al eds *Decision Processes*, pp. 49–59, New York: Wiley
- Rawls, John (1971) *A Theory of Justice*, 1st ed., Cambridge, MA: HUP
- ——— (1974) 'Some Reasons for the Maximin Criterion', *American Economic Review* 64 (2): 141–146