Decisions under Ignorance

Johan E. Gustafsson

The Decision Matrix

We will model choice situations in terms of acts, states of nature, and outcomes.

Example:

States of nature ↓

		Rain	*No rain*
Acts →	*Bring umbrella*	4	8
	Leave umbrella	2	10

↑ Values of outcomes

Acts

The acts are the alternative things one can do in a choice situation.

Some standard requirements on the set of acts in a choice situation:

> The acts should be jointly exhaustive in the sense that one has to perform at least one of them in a choice situation.
>
> The acts should be mutually exclusive in the sense that if one performs one of them one does not perform any of the other acts in the situation.

States of Nature

The states of nature are different descriptions of the world over which one has no control in a choice situation.

Some standard requirements on the set of states of nature in a choice situation are:

> The states should be jointly exhaustive in the sense that at least one of the states is true.
>
> The states should be mutually exclusive in the sense that at most one of the states is true.

Outcomes

The outcomes are what will happen given a combination of an act and a state of nature. The outcomes are the things over which the agent has preferences.

A Taxonomy

Decisions under certainty
are decisions made when one knows what the outcome of each act would be.

Decisions under risk
are decisions made when one does not know the outcome of each act yet can assign a subjective probability to the possible outcomes of each act.

Decisions under ignorance
are decisions made when one cannot assign a subjective probability to the possible outcomes of any act.

Decisions under Ignorance

For decisions under ignorance there are no probabilities available to the agent. So the standard principle for decision making—that is, the principle of maximizing expected utility—cannot be used.

The Maximin Rule

The maximin rule says that an act is preferred to another act if and only if its worst possible outcome is preferred to the worst possible outcome of the other act. And two acts are indifferent if and only if their worst possible outcomes are indifferent.

Example:

	s_1	s_2	s_3
a_1	4	3	2
a_2	1	9	7
a_3	5	6	5

Here, the maximin rule says that a_3 is the best act.

The Maximin Rule and Dominance

A standard rationality requirement is following dominance condition:

Dominance
If, for each possible state, the outcome of act x is at least as preferred as the outcome of act y and, for some possible state, the outcome of act x is preferred to the outcome of act y, then x is preferred to y.

Example:

	s_1	s_2	s_3
a_1	5	3	2
a_2	3	3	2

Here, since the worst possible outcome of a_2 is at least as preferred as the worst outcome of a_1, the maximin rule says that a_1 and a_2 are indifferent.

But, according to Dominance, a_1 is preferred to a_2.

The Leximin Rule

The leximin rule is a variation of the maximin rule that avoids this conflict with Dominance.

The leximin rule is like the maximin rule except that ties are broken in favour of the act with the best second worst outcome and, if there is still a tie, it breaks the tie in favour of the act with the best third worst outcome, and so on.

	s_1	s_2	s_3
a_1	5	3	2
a_2	3	3	2

Here, the leximin rule says that a_1 is preferred to a_2.

Hence we avoid conflict with Dominance.

Both the maximin and the leximin rules might seem to yield overly pessimistic recommendations. In the following case, the potential gain from performing a_1 rather than a_2 is very large but the potential loss is very small.

	s_1	s_2	s_3
a_1	100	100	1
a_2	2	2	2

Yet the maximin and the leximin rules each yields that a_2 is preferred to a_1.

Column Linearity

Both the maximin and the leximin rules violate

Column Linearity
Preferences over acts does not change if the utilities for all outcomes in one state of nature are increased by the same amount for all acts.

Example 1:

	s_1	s_2
a_1	3	2
a_2	4	1

Example 2:

	s_1	s_2	
a_1	3	6	$(= 2 + 4)$
a_2	4	5	$(= 1 + 4)$

The maximin and leximin rules prefer a_1 in example 1 but a_2 in example 2.

The Minimax-Regret Rule

Let the regret in an outcome be the how much greater one's utility could have been given that one had performed one of the other acts given that the same state of nature is true.

The minimax-regret rule says then that an act is at least as preferred as another act if and only if its worst possible regret in an outcome is at least as small as the worst possible regret in an outcome for the other act.

Example:

	s_1	s_2	s_3
a_1	5	-2	10
a_2	-1	-1	20
a_3	-3	-1	5
a_4	0	-4	1

Regret table:

	s_1	s_2	s_3
a_1	0	1	10
a_2	6	0	0
a_3	8	0	15
a_4	5	3	19

Since a_2 has the minimal maximal (worst) regret it is the best act according to the minimax-regret rule.

The minimax-regret rule's evaluation of two acts depend in part on what other acts are available. That is, the rule violates:

> **Row Adjunction**
> Whether x is at least as preferred as y does not depend on which other alternatives are available.

Moreover, if we may only choose acts which are optimal according to the minimax-regret rule, we would violate:

> **Contraction Consistency**
> If an act x is rationally permissible given choice from the set of alternative acts U and x is in the subset V of U, then x is rationally permissible given a choice from V.

Example 1:

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10
a_3	10	5	1

Regret table 1:

	s_1	s_2	s_3
a_1	10	0	6
a_2	5	8	0
a_3	0	5	9

Example 2:

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Regret table 2:

	s_1	s_2	s_3
a_1	5	0	6
a_2	0	8	0

The minimax-regret rule yields that a_2 is rationally permissible in example 1 but not in example 2. Yet the difference between the two examples is merely that a_3 is no longer available in example 2. So the minimax-regret rule violates the principle of contraction consistency.

And, since a_1 is preferred to a_2 in example 2 but not in example 1, the minimax-regret rule violates Row Adjunction.

The Laplace Rule

The Laplace rule is also known as the principle of insufficient reason.

The basic idea behind this rule is that, if for any two states of nature we have no reason to regard one of them as more probable than the other, then we should regard them as equally probable.

The Laplace rule says that one should value acts by their expected value as if each of the possible states of nature were equally credible.

Example:

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10
a_3	10	5	1

Here, the Laplace rule says that the best act is a_2.

A standard objection to the Laplace rule is that it is sensitive to how one individuates states of nature, that is, how one decides what counts as the same state. That is, the Laplace rule violates:

Column Duplication
Whether act x is at least as preferred as act y does not depend whether a state of nature is split into two duplicate states.

Example 1:

	s_1	s_2	s_3
a_1	1	3	3
a_2	4	1	1

In example 1, the Laplace rule says that a_1 is preferred to a_2.

Example 2:

	s_1	s_2 or s_3
a_1	1	3
a_2	4	1

But, in example 2, the Laplace rule says that a_2 is preferred to a_1.

Yet the leximin rule is also sensitive to how one individuates states of nature. In fact, it yields the same rankings as the Laplace rule in the last example. Hence the leximin rule also violates Column Duplication.

Example 1:

	s_1	s_2	s_3
a_1	1	3	3
a_2	4	1	1

In example 1, the leximin rule says that a_1 is preferred to a_2.

Example 2:

	s_1	s_2 or s_3
a_1	1	3
a_2	4	1

In example 2, however, the leximin rule says that a_2 is preferred to a_1.

Ordinal Scales

An ordinal scale is a scale where only the order of the numbers have significance.

For the maximin and leximin rules, an ordinal scale is enough.

Example 1:

	s_1	s_2
a_1	1	3
a_2	2	4

Example 2:

	s_1	s_2
a_1	1	5
a_2	4	9

The order of the outcomes are the same in examples 1 and 2. So, given an ordinal scale, there is no significant difference between these examples.

The maximin and leximin rules both prefer a_2 in both examples.

Note that, unlike the maximin and leximin rules, the minimax-regret rule requires more than an merely ordinal utility scale.

Example 1:

	s_1	s_2
a_1	3	4
a_2	1	5

Regret table 1:

	s_1	s_2
a_1	0	1
a_2	2	0

Example 2:

	s_1	s_2
a_1	3	4
a_2	1	7

Regret table 2:

	s_1	s_2
a_1	0	3
a_2	2	0

The minimax-regret rule prefers a_1 to a_2 in example 1, but it prefers a_2 to a_1 in example 2, even though the ordinal rankings of the outcomes haven't changed.

The Laplace rule also needs more than an ordinal scale.

Example 1:

	s_1	s_2
a_1	1	5
a_2	2	3

Example 2:

	s_1	s_2
a_1	1	5
a_2	3	4

The outcomes are ordered the same in examples 1 and 2; yet the Laplace rule prefers a_1 in example 1 but a_2 in example 2.

Interval Scales

An interval scale is a scale where the relative differences between alternatives are significant.

An interval scale does not gain or loose any significant information given a positive linear transformation.

Mathematically, if $f(x)$ is a function that return the utility of x on an interval scale, the the following function does so to:

$f'(x) = k \times f(x) + m,$

where k and m are positive constants.

Example 1:

	s_1	s_2
a_1	0	4
a_2	2	3

Example 2 (multiply by 5 and add 2):

	s_1	s_2
a_1	2	22
a_2	12	17

The two interval scales in examples 1 and 2 represent the same information.

Both the minimax-regret rule and the Laplace rule require no more than utilities represented by interval scales.

Example 1:

	s_1	s_2
a_1	0	4
a_2	2	3

Example 2 (multiply by 5 and add 2):

	s_1	s_2
a_1	2	22
a_2	12	17

The minimax-regret and the Laplace rules' results do not change given a positive linear transformation.

They both prefer a_2 in both examples.

Ordering

Preferences form an ordering if and only if they satisfy completeness and transitivity.

Transitivity
For all alternatives x, y, and z, if x is at least as preferred as y and y is at least as preferred as z, then x is at least as preferred as z.

Completeness
For all alternatives x and y, either x is at least as preferred as y or y is at least as preferred as x.

The Laplace rule holds if and only if following conditions holds:

Ordering
Preferences over acts are complete and transitive.

Symmetry
Preferences over acts does not depend on the labelling of the states of nature and acts.

Strong Domination
If, for each state of nature, the outcome of act x is preferred to the outcome of act y, then x is preferred to y.

Row Adjunction (violated by the minimax-regret rule)
Whether x is at least as preferred as y does not depend on what other alternatives are available.

Column Linearity (violated by the maximin and leximin rules)
Preferences over acts do not change if the utilities for all outcomes in one state of nature are increased by the same amount for all acts.

A connection between ethics and decisions under ignorance

John C. Harsanyi (1953) argues that value judgements are a special kind of preferences, namely, non-egoistic impersonal preferences.

From this, he then argues that one judges that a society X is at least as good as a society Y if X is at least as preferred as Y in a situation where one does not know who one is in the society.

So one's value ordering of two societies should be the same as one's preference ordering between the two societies given that one does not know who one is in the societies.

John Rawls (1974) presents a similar argument for his theory of justice. The main difference is that Harsanyi favours the Laplace rule but Rawls favours the Maximin rule.

Using the Laplace rule, we end up with utilitarianism:

> A society x is at least as good as a society y if and only if the sum total of utility is at least as great in x as in y.

Using the maximin rule, we end up with a Rawlsian theory:

> A society x is at least as good as a society y if and only if the worst off in x is at least as well of as the worst of in y.

Example:

	P_1	P_2	P_3
O_1	100	100	1
O_2	2	3	3

Utilitarianism yields that a_1 is better than a_2, but the Rawlsian theory yields that a_2 is better than a_1.

Similarly, if we combine the veil of ignorance approach with the minimax-regret rule, we get what is know as the complaint model.

On this view a person's complaint against an act in a situation is equal to the difference between that person's well-being in the outcome of the act and their maximum well-being in the outcome of any available act in the situation.

Then, an act x is at least as good as an alternative act y if and only if the maximum complaint for any person given x is at least as small as the maximum complaint for any person given y.

Example:

	P_1	P_2	P_3
O_1	5	2	10
O_2	1	1	20
O_3	7	4	5

Complaint table:

	P_1	P_2	P_3
O_1	2	2	10
O_2	6	3	0
O_3	0	0	15

The complaint model yields that a_2 is the best act.

An Axiomatic Approach to Ethics

Another connection between decisions under ignorance and ethics, is that there are ethical analogues of the conditions for decision under ignorance and conditions in ethics, given that we replace possible states of nature with people.

States of nature ↓

Acts →

	s_1	s_2
a_1	4	8
a_2	2	10

↑ Values of outcomes

People ↓

Outcomes →

	P_1	P_2
O_1	4	8
O_2	2	10

↑ People's well-being

This analogy allows us to use the above results to reason axiomatically in ethics.

The Trolley Problem

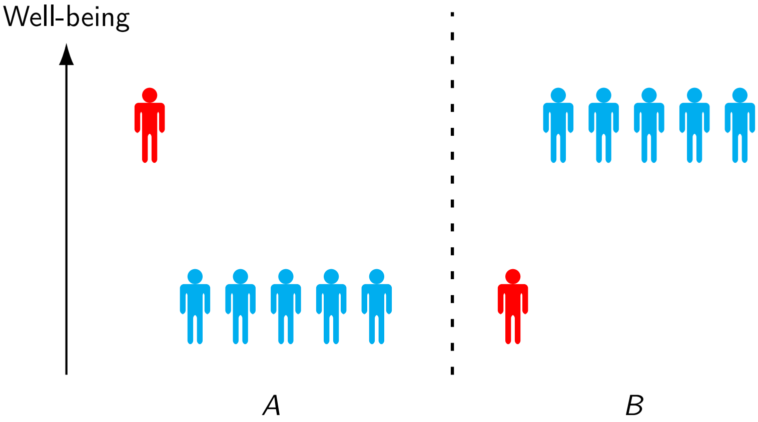

Symmetry and Impartiality

The decision-theoretic condition

> **Symmetry**
> Preferences over acts does not depend on the labelling of the states of nature and acts.

is structurally the same as

> **Impartiality**
> If two outcomes only differ in that some people have switched well-being levels, then the outcomes are equally good.

Impartiality

If two outcomes only differ in that some people have switched well-being levels, then the outcomes are equally good.

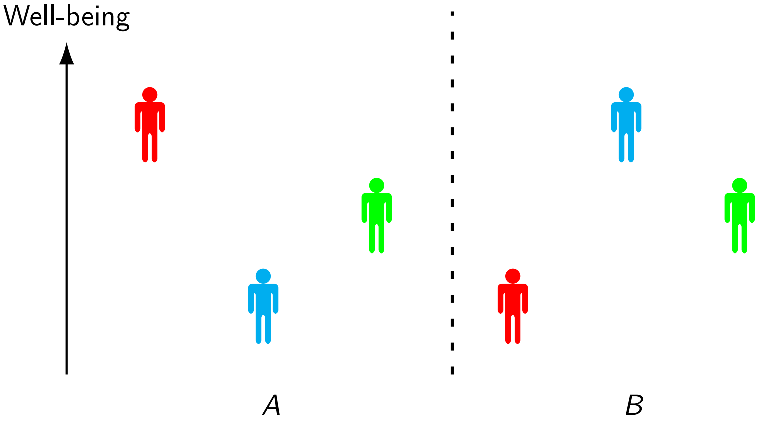

Since *A* and *B* only differ in that Red and Blue have switched well-being levels, *A* is equally good as *B*.

Dominance and Pareto

Likewise

Dominance
If, for each possible state, the outcome of act x is at least as preferred as the outcome of act y and, for some possible state, the outcome of act x is preferred to the outcome of act y, then x is preferred to y.

is structurally the same as

Pareto
If everyone has at least as much well-being in outcome X as in outcome Y and someone has more well-being in X than in Y, then X is better than Y.

Pareto

If everyone has at least as much well-being in outcome A as in outcome B and someone has more well-being in A than in B, then A is better than B.

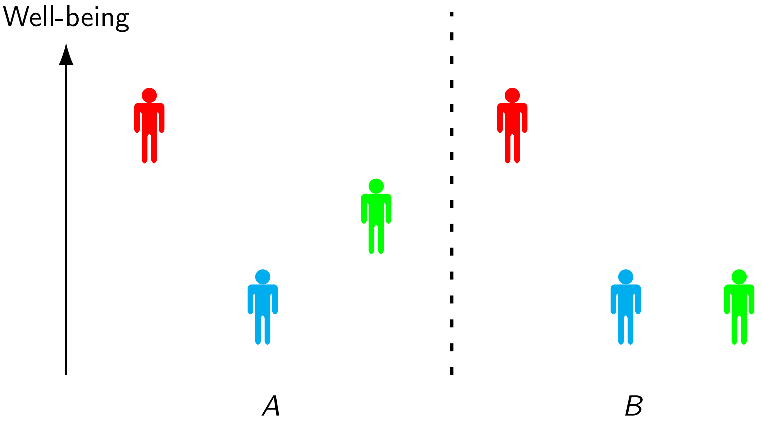

Since everyone has at least as much well-being in A as in B and Green has more well-being in A than in B, A is better than B.

Value comparisons form an ordering if and only if they satisfy completeness and transitivity.

Transitivity
For all alternatives X, Y, and Z, if X is at least as good as Y and Y is at least as good as Z, then X is at least as good as Z.

Completeness
For all alternatives X and Y, either X is at least as good as Y or Y is at least as good as X.

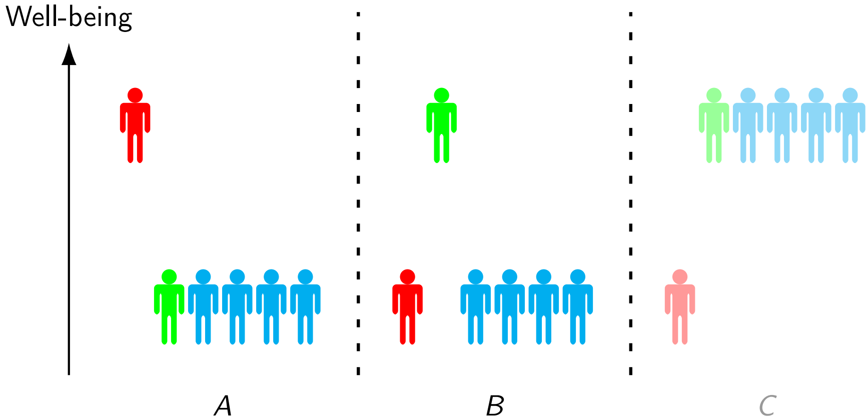

A is equally good as B.

Impartiality entails that A is equally good as B, since these outcomes are the same except that Red and Green has swapped well-being levels.

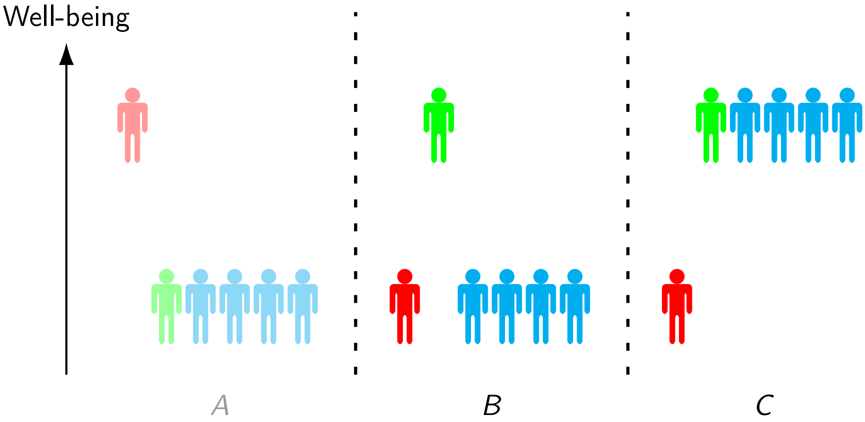

A is equally good as B. B is worse than C.

Dominance entails that C is better than B, since everyone has at least as much well-being in C as in B and each blue person has more well-being in C than in B.

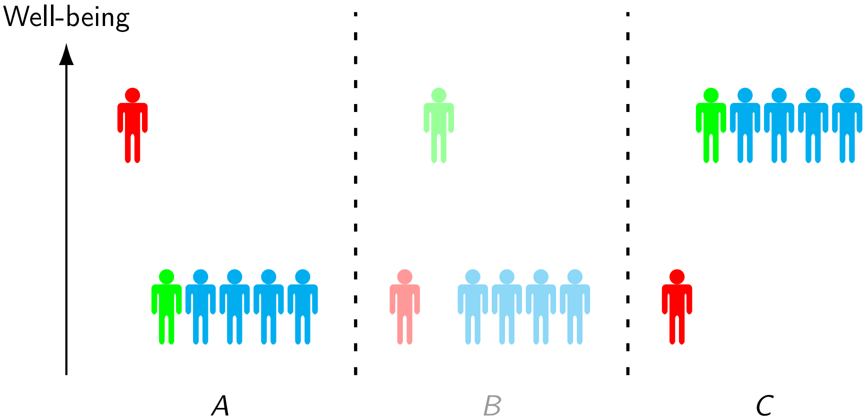

A is equally good as B. B is worse than C.

Hence, by Ordering, A is worse than C.

Hence saving five is better than saving one.

Utilitarianism holds (given a fixed population) if and only if following conditions holds:

Ordering
Value comparisons over outcomes are complete and transitive.

Impartiality
If two outcomes only differ in that some people have switched well-being levels, then the outcomes are equally good.

Strong Pareto
If everyone is at least as well off in outcome X as in outcome Y and someone is better off in X than in Y, then X is better than Y.

Ethical Row Adjunction (violated by the complaint model)
Whether X is at least as good as Y does not depend on what other outcomes are available.

Ethical Column Linearity (violated by Rawls)
Value comparisons for the available outcomes do not change if someone's well-being is increased by the same amount in all of them.

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 Robert M. Thrall, Clyde H. Coombs, and Robert L. Davis, eds., *Decision Processes*, pp. 49–59, New York: Wiley.

Rawls, John (1974) 'Some Reasons for the Maximin Criterion', *American Economic Review* 64 (2): 141–146.