Decision Theory

1.2 Decisions Under Ignorance

George Mason University, Spring 2018

What should you do in a choice situation where you do not know the probabilities of the outcomes of one or more of your potential actions?

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

Def 1.2.1. $a_1 \in \mathcal{A}$ *weakly dominates* $a_2 \in \mathcal{A}$ if and only if $u(g(a_1, s)) \geq u(g(a_2, s))$ for every state $s \in \Omega$.

Def 1.2.1.1. $a_1 \in \mathcal{A}$ *dominates* $a_2 \in \mathcal{A}$ if and only if $u(g(a_1, s)) \geq u(g(a_2, s))$ for every state $s \in \Omega$ and $u(g(a_1, s)) > u(g(a_2, s))$ for some state $s \in \Omega$.

Thesis 1.2.1 (Dominance Principle). Rationality forbids any dominated action.

The Dominance Principle presupposes an ordinal utility function.

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2

	s_1	s_2	s_3	s_4
a_1	1	-2	3	5
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_1	1	-2	3	5
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

As we will see, the Dominance Principle is one of the few uncontroversial theses in decision theory (though it must be applied with caution when actions and states are causally dependent).

However, the Dominance Principle usually doesn't pin down a single action in $\mathcal{A}$. So it must be supplemented with additional principles.

Let $min(a)$ designate the minimum utility obtainable by performing act $a \in \mathcal{A}$. That is, $min(a) = min_{s \in \Omega}(u(g(a, s)))$.

Thesis 1.2.2 (Maximin Rule). Rationality forbids any action $a_i \in \mathcal{A}$ such that $min(a_i) \neq max_{a \in \mathcal{A}}(min(a))$.

In other words, maximize the minimum utility obtainable.

The Maximin Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

$$min(a_3) = 4 > min(a_2) = -3.$$

	s_1	s_2	s_3	s_4
a_3	10	7	4	5

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_2	5	-2	3	5
a_3	0	1	2	20
a_4	-3	10	-5	0

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_2	5	-2	3	5
a_3	0	1	2	20
a_4	-3	10	-5	0

$$min(a_1) = min(a_3) = max_{a \in \mathcal{A}}(min(a)) = 0.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

To break ties, we can consider second worst outcomes, third worst outcomes, and so forth.

Let $min^1(a)$ designate the lowest utility obtainable with $a \in \mathcal{A}$.
Let $min^2(a)$ designate the second lowest utility obtainable with $a \in \mathcal{A}$.
Let $min^n(a)$ designate the nth lowest utility obtainable with $a \in \mathcal{A}$.

Thesis 1.2.3 (Lexical Maximin Rule). Rationality forbids any action $a_1 \in \mathcal{A}$ such that there is some $n > 0$ and alternative action $a_2 \in \mathcal{A}$ where $min^n(a_2) > min^n(a_1)$ and $min^m(a_2) = max_{a \in \mathcal{A}}(min^m(a))$ for all $m < n$.

In other words, maximize the worst utility obtainable. But in the case of a tie, maximize the second worst utility obtainable. And so forth.

The Lexical Maximin Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^1(a_1) = min^1(a_3) = 0.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^2(a_1) = min^2(a_3) = 1.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^3(a_1) = 9 > min^3(a_3) = 2.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0

Problem. The maximin rules are overly conservative and pessimistic. It seems eminently rational to accept a slightly lower minimum utility for the chance of massive gain.

	s_1	s_2
Gamble A	\$101	\$101
Gamble B	\$100	\$10000

This kind of decision matrix suggests that we should also consider the maximum utility obtainable.

Let $max(a)$ designate the maximum utility obtainable by performing act $a \in \mathcal{A}$. That is, $max(a) = max_{s \in \Omega}(u(g(a, s)))$.

Thesis 1.2.4 (Maximax Rule). Rationality forbids any action $a_1 \in \mathcal{A}$ such that $max(a_1) \neq max_{a \in \mathcal{A}}(max(a))$.

In other words, maximize the maximum utility obtainable.

The Maximax Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_1	5	10	3	0
a_2	20	-2	-50	-10
a_3	0	1	4	7
a_4	-30	20	-5	0

	s_1	s_2	s_3	s_4
a_1	5	10	3	0
a_2	20	-2	-50	-10
a_3	0	1	4	7
a_4	-30	20	-5	0

$$max(a_2) = max(a_4) = max_{a \in \mathcal{A}}(max(a)) = 20.$$

	s_1	s_2	s_3	s_4
a_2	20	-2	-50	-10
a_4	-30	20	-5	0

To break ties, we could also introduce a Lexical Maximax Rule and consider second best outcomes, third best outcomes, and so forth.

The Maximax Rule is quite optimistic. But it has few adherents.

A more flexible rule allows for compromise between the maximum and minimum utility obtainable.

Def 1.2.2. Given an *optimism index* $\alpha \in \mathbb{R}[0,1]$ that represents the degree of optimism of the decision maker, the α-*index* of $a \in \mathcal{A}$ is $\alpha \times max(a) + (1-\alpha) \times min(a)$.

Thesis 1.2.5 (Optimism-Pessimism Rule). Rationality forbids any act $a_1 \in \mathcal{A}$ whose α-index is lower than that of some alternative act $a_2 \in \mathcal{A}$.

In other words, maximize the α-index.

When $\alpha = 1$, the OP Rule collapses to the Maximax Rule.

When $\alpha = 0$, the OP Rule collapses to the Minimax Rule.

The Optimism-Pessimism Rule presupposes an interval utility function.

	s_1	s_2	s_3
a_1	10	4	0
a_2	2	5	6

If $\alpha = 0.5$, then the α-index of a_1 is $0.5 \times 10 + 0.5 \times 0 = 5$ and the α-index of a_2 is $0.5 \times 6 + 0.5 \times 2 = 4$. So the OP Rule requires a_1.

If $\alpha = 0.2$, then the α-index of a_1 is $0.2 \times 10 + 0.8 \times 0 = 2$ and the α-index of a_2 is $0.2 \times 6 + 0.8 \times 2 = 2.8$. So the OP Rule requires a_2.

	s_1	s_2	s_3
a_1	8	4	0
a_2	3	5	7

If $\alpha = 0.5$, then the α-index of a_1 is $0.5 \times 8 + 0.5 \times 0 = 4$ and the α-index of a_2 is $0.5 \times 7 + 0.5 \times 3 = 5$. So the OP Rule requires a_2.

The fact that a rule presupposes an interval utility function is sometimes considered a disadvantage.

Problem. It seems rational to consider non-extreme utilities when making a decision.

	s_1	s_2	s_3
Gamble A	\$100	\$1	\$0
Gamble B	\$20	\$79	\$80

A natural response is to assign α-values to non-extreme outcomes.

Problem. Since α is subjective, the Optimism-Pessimism Rule is unstable.

An agent's optimism index can change over time, so the OP Rule imposes no consistency on an agent's decisions over time.

Moreover, we could rationalize impulsive behavior by declaring an appropriate degree of optimism.

The next rule focuses on *missed opportunities*.

Let $max(s)$ designate the maximum utility obtainable in state $s \in \Omega$. That is, $max(s) = max_{a \in \mathcal{A}}(u(g(a,s)))$.

Def 1.2.3. The *regret value* of $a \in \mathcal{A}$ in $s \in \Omega$ is $r(a,s) = u(g(a,s) - max(s))$.

Let $max\text{-}regret(a)$ designate the maximum regret value obtainable by performing act $a \in \mathcal{A}$. That is, $max\text{-}regret(a) = max_{s \in \Omega}(r(a,s))$.

Thesis 1.2.6 (Minimax Regret Rule). Rationality forbids any act $a_1 \in \mathcal{A}$ such that $max\text{-}regret(a_1) \neq min_{a \in \mathcal{A}}(max\text{-}regret(a))$.

In other words, minimize the maximum possible regret.

The Minimax Regret Rule presupposes an interval utility function.

	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

Step 1. Identify max(s) for each $s \in \Omega$.

	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

Step 1. Identify max(s) for each $s \in \Omega$.

	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

Step 2. Calculate $r(a, s)$ for each $a \in \mathcal{A}$ and $s \in \Omega$.

	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

The resulting matrix is called a *regret matrix*.

	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

Step 3. Identify $\textit{max-regret}(a)$ for each $a \in \mathcal{A}$.

	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

Step 3. Identify *max-regret*(a) for each $a \in \mathcal{A}$.

	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

Step 4. Perform act that minimizes *max-regret*.

	s_1	s_2	s_3
a_2	-6	0	0

To break ties, we could introduce a Lexical Minimax Regret Rule.

Why would this work?

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Matrix

	s_1	s_2	s_3
a_1	-5	0	-6
a_2	0	-8	0

Regret Matrix

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Matrix

	s_1	s_2	s_3
a_1	-5	0	-6
a_2	0	-8	0

Regret Matrix

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Matrix

	s_1	s_2	s_3
a_1	-5	0	-6
a_2	0	-8	0

Regret Matrix

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

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

Decision Matrix

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

Regret Matrix

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

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

Decision Matrix

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

Regret Matrix

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

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

Decision Matrix

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

Regret Matrix

The final decision-theoretic principle brings probabilities into the picture.

Guiding idea: If we are in a state of total ignorance, then we should treat all states in Ω as equally likely since there is *insufficient reason* for any other probability measure.

	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$	$[\frac{1}{10}]$
	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
a_1	10	9	8	7	6	5	4	3	2	1
a_2	10	0	0	0	0	0	0	0	0	0

Combining this guiding idea with the principle that we should maximize *expected utility* (more on this later) gives us:

Thesis 1.2.7 (Principle of Insufficient Reason). Rationality forbids any action $a_1 \in \mathcal{A}$ such that there is some alternative action $a_2 \in \mathcal{A}$ where $\sum_{s \in \Omega} \frac{1}{|\Omega|} \times u(g(a_2, s)) > \sum_{s \in \Omega} \frac{1}{|\Omega|} \times u(g(a_1, s))$.

$|\Omega|$ is the number of states in Ω.

The PIR presupposes an interval utility function.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_1 is $\frac{1}{3} \times 6 + \frac{1}{3} \times 3 + \frac{1}{3} \times 0 = 3$.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_2 is $\frac{1}{3} \times 12 + \frac{1}{3} \times 12 + \frac{1}{3} \times 9 = 11$.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_3 is $\frac{1}{3} \times 0 + \frac{1}{3} \times 9 + \frac{1}{3} \times 0 = 3$.

	s_1	s_2	s_3
a_2	12	12	9

Problem. The Principle of Insufficient Reason is extremely sensitive to how states are individuated.

	$[\frac{1}{2}]$ fresh	$[\frac{1}{2}]$ unfresh
eat raw	5	-3
eat grilled	2	2
eat nothing	0	0

The expected utility of eating raw oysters is 1 while the expected utility of eating grilled oysters is 2. Thus, the PIR forbids eating raw oysters.

Problem. The Principle of Insufficient Reason is extremely sensitive to how states are individuated.

	$[\frac{1}{3}]$ extremely fresh	$[\frac{1}{3}]$ fresh	$[\frac{1}{3}]$ unfresh
eat raw	5	5	-3
eat grilled	2	2	2
eat nothing	0	0	0

The expected utility of eating raw oysters is $\frac{7}{3}$ while the expected utility of eating grilled oysters is 2. Thus, the PIR forbids eating grilled oysters.

Problem. There is arguably insufficient reason to adopt *any* probability measure over Ω including the uniform distribution.

Problem. The Principle of Insufficient Reason leads to counterintuitive results.

	s_1	s_2
Gamble A	\$300	-\$100
Gamble B	\$90	\$90

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Maximin Rule allows only a_1.

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Minimax Regret Rule allows only a_2.

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Optimism-Pessimism Rule (with $\alpha = 0.5$) allows only a_3.

How to decide between the 7 principles?

Some decision theorists have tried to provide an *axiomatic* justification of their preferred decision rule. They argue that only their preferred rule satisfies certain reasonable conditions or axioms.

But this just pushes back the debate to the axioms themselves.

Resnik: "The debate could go back and forth over the conditions in this fashion with results no more conclusive than our previous discussions of the rules themselves. The situation is likely to remain this way, I think, until we have amassed a rich backlog of studies of genuine real-life examples of good decision making under ignorance. My hope is that these would help us sharpen our judgments concerning the various rules. I would also conjecture that we will ultimately conclude that no rule is always the rational one to use but rather that different rules are appropriate to different situations. If this is so, it would be more profitable to seek conditions delimiting the applicability of the various rules rather than to seek ones that will declare in favor of a single rule."