Knowledge-Based Decision Theories Are Vulnerable to Money Pumps

Petra Kosonen

kosonenpetra@gmail.com

Maximizing expected utility

- In this talk, I'll show that if you conditionalize on knowledge before maximizing expected utility, you are vulnerable to a money pump: you would pay for something that you could keep for free.
- On the standard decision theory, rational agents always maximize expected utility, where expected utility is a function of utility and (subjective) probability.
- On this view, knowledge does not play a role in instrumental rationality.

Knowledge, credences and action

- Having knowledge that P is independent of whether one should act on one's belief that P:
- Knowledge that P is not sufficient for action as one might know that P even though one lacks a sufficiently high credence to justify acting on the proposition that P.
- And knowledge that P is not necessary for action as a high credence can license an act even when the agent lacks knowledge.[1]
- So, it is subjective degrees of belief—not knowledge—that matter for rational action.

[1] Hawthorne and Stanley (2008, p. 571).

Knowledge-Based Decision Theories

- Some have proposed that the standard view should be revised to account for the agent's knowledge.
- On the proposed revision, agents should somehow take into account what they know when maximizing expected utility—let's call these views *Knowledge-Based Decision Theories*.[2]
- On these views, what one rationally ought to do depends on what one knows.
- According to the simplest version of this view, agents should conditionalize on their knowledge before maximizing expected utility in such a way that known propositions receive probability 1.[3]

[2] Goldschmidt (2024, p. 1). Some have proposed Knowledge-Based Decision Theories as a solution to fanaticism about tiny probabilities of vast utilities. See for example Hong (2024).

[3] This view is tentatively suggested by Hawthorne and Stanley (2008).

Clarification

- Knowledge$\neq$credence 1 (epistemically)
- You might know that P and still have a credence less than 1 in P.
- The agent conditionalizes on what they know *only for the purposes of decision-making*.
- 'Action-credences' (used in expected utility calculations) vs. 'epistemic-credences' (what the agent believes)

Losing knowledge (Epistemic defeat)

- Next, I'll show that Knowledge-Based Decision Theories are vulnerable to a money pump if it is possible to lose knowledge.
- Consider the following lotteries:

  Ticket A Gives a great payoff if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

  Ticket B Gives a modest positive payoff if the lottery machine returns numbers 1, 2, 3, 4 and 5 (and otherwise it gives nothing).

Values of A and B

- Ticket A gives a great payoff if you guess all seven lottery numbers correctly, while ticket B gives a modest positive payoff if you guess at least five lottery numbers correctly.
- Suppose you start knowing that ticket A wins nothing, but you do not know that ticket B wins nothing.
- If it is possible to have knowledge in lottery cases, then there must be some (possibly vague and context-dependent) threshold for when a probability is high enough to count as knowledge.
- We may suppose that the probability of not winning with A is above this threshold, but the probability of not winning with B is below this threshold.
- Consequently, B is worth some amount to the agent, while A is worthless (or at most marginally better than nothing).

Money Pump Setup

- The setup of the money pump is as follows:
- You currently have B, which is worth something to you.
- If the lottery machine return numbers 1–5, then you will be offered A in exchange for B.
- If you learn that the lottery machine returns numbers 1–5, you no longer know that the lottery machine will not return numbers six and seven as well.
- In that case, you would only need to get two more numbers right, and for all you know, you might do so.
- So, you would then prefer A to B, and happily accept the trade.

Money Pumped

- This is unfortunate: Right now you know that you will not win anything with A. So, it would be better to keep B.
- However, you also know that if you win anything with B, you will accept the trade and end up with A.
- Luckily, you are offered a chance to avoid this situation: If you pay some amount $\$\epsilon$, you will not be offered A in exchange for B in case the lottery machine returns number 1–5.
- And, given that B is worth some positive amount while A is worth nothing, you accept this offer.
- But then you have paid for something you could have kept for free: you've been money pumped.

Money Pump

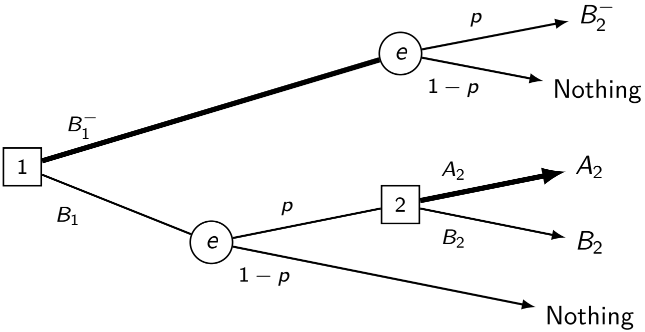

$$A_2 \succ B_2, \text{ and } B_1 \succ B_1^- \succ A_1.$$

Ticket A Gives a great payoff if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

Ticket B Gives a modest positive payoff if the lottery machine returns numbers 1, 2, 3, 4 and 5 (and otherwise it gives nothing).

Knowledge in lottery cases

- One might object that it is not possible to have knowledge in lottery cases.
- However, the case can be changed such that it does not involve lottery tickets.
- The important features in this case are that the agent starts knowing that some event E_1 cannot happen, the agent does *not* know that some event E_2 cannot happen, and if event E_2 happens, then the agent no longer knows that event E_1 cannot happen.
- E.g. perhaps the agent knows they don't have some very rare condition, but if they developed symptoms associated with that condition, they no longer know they don't have it. And they don't know they won't develop such symptoms.

Summary

- To summarize, Knowledge-Based Decision Theories are vulnerable to money pumps if it is possible to lose knowledge.
- The next section introduces another money pump against these views.

Gaining knowledge

- This section shows that Knowledge-Based Decision Theories are vulnerable to a money pump if it is possible to gain knowledge.
- Consider the following lottery tickets:

  Ticket A Gives a 0.99999 probability of $1,000,001 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

  Ticket B Certainly gives $1,000,000 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

B has higher EU than A

- At the start, the agent does not know if ticket A wins the prize, and similarly, they do not know if ticket B wins the prize. For all the agent knows, tickets A and B might or might not win.
- Consequently, B is better than A, as it has a higher expected utility/monetary value (1,000,000p vs. 999,991p, where p is the probability of the lottery machine returning numbers 1–7).[4]

[4] $p \cdot 1,000,000 = 1,000,000p$ and $p \cdot 0.99999 \cdot 1,000,001 = 999,991p$.

A better than B conditional on knowledge

- However, the agent knows that if the lottery machine returns numbers 1–7, they come to know that both A and B will win, even though there is a tiny probability that A does not do so—knowledge does not require certainty.
- After conditionalizing on each of A and B winning, A is better than B, as it gives a better prize.

Ticket A Gives a 0.99999 probability of $1,000,001 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

Ticket B Certainly gives $1,000,000 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise it gives nothing).

Money Pump

- Suppose the agent starts with B.
- Before the lottery numbers are drawn, the agents knows that they will prefer A if the lottery machine returns numbers 1–7. So they know that, in that case, they would trade B for A.
- But right now the agent prefers B.
- They are therefore willing to pay some amount of money to avoid being offered A if the lottery machine returns numbers 1–7.
- But then they have been money pumped: they have paid for something they could have kept for free.

Money Pump

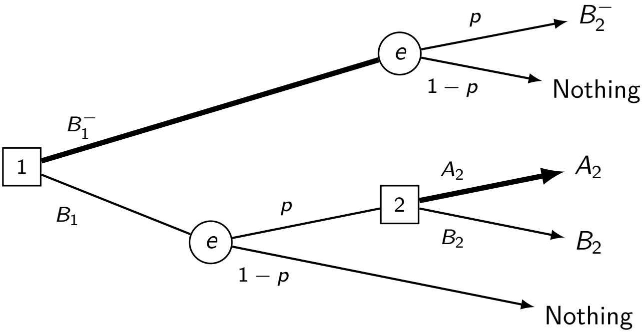

$$A_2 \succ B_2, \text{ and } B_1 \succ B_1^- \succ A_1.$$

Ticket A Gives a 0.99999 probability of \$1,000,001 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise nothing).

Ticket B Certainly gives \$1,000,000 if the lottery machine returns numbers 1, 2, 3, 4, 5, 6 and 7 (and otherwise nothing).

What happened

- Knowledge-Based Decision Theories get one in trouble if there can be cases where, after obtaining some evidence, one comes to know that P and Q, but there is a tiny probability that P is false.
- (There might also be a tiny probability that Q is false. The case will work as long as the probability of P being false is greater than the probability of Q being false.)

What happened

- Ticket A: lower probability of winning, better prize.
- Ticket B: higher probability of winning, worse prize.
- At the start, the agent lacks knowledge in either ticket winning. B has higher expected utility, so B is better.
- Some event happens (lottery machine returns numbers 1–7), and the agent then knows that both tickets will win.
- A gives a better prize, so A is better.
- The money pump exploits this reversal of preference.
- To summarize, Knowledge-Based Decision Theories are vulnerable to money pumps if—as seems undeniable—it is possible to gain knowledge.

Knowledge requires certainty

- One might object that knowledge requires certainty.[5]
- In that case, one cannot know that ticket A wins because there is a tiny probability that A does not win.
- However, if knowledge requires certainty, then it seems that we very rarely have knowledge.
- Consequently, Knowledge-Based Decision Theories very rarely differ from Expected Utility Theory; if one ought to conditionalize on one's knowledge before maximizing expected utility, but one knows almost nothing, then one should almost always simply maximize expected utility.
- Knowledge-Based Decision Theories are irrelevant in practice.

[5] The Knowledge-Based Decision Theory proposed by Hawthorne and Stanley (2008) requires that the probability of known propositions be one.

Conclusion

- Knowledge-Based Decision Theories tell one to take into account what one knows when maximizing expected utility.
- One way to do this is to conditionalize on one's knowledge before maximizing expected utility.
- I've shown that this view is vulnerable to money pumps.

References I

Goldschmidt, Z. (2024), 'Foundations for knowledge-based decision theories', *Australasian Journal of Philosophy* **102**(4), 939–958.

Hawthorne, J. and Stanley, J. (2008), 'Knowledge and action', *The Journal of Philosophy* **105**(10), 571–590.

Hong, F. (2024), 'Know your way out of St. Petersburg: An exploration of "knowledge-first" decision theory', *Erkenntnis* **89**(6), 2473–2492.