

Dynamic Causal Decision Theory

Johan E. Gustafsson*

Draft
June 27, 2024

ABSTRACT. Causal decision theorists are vulnerable to a money pump if they update by conditioning when they learn what they have chosen. Nevertheless, causal decision theorists are immune to money pumps if they instead update by imaging on their choices and by conditioning on other things (and, in addition, evaluate plans rather choices). I also show that David Lewis's Dutch-book argument for conditioning does not work when you update on your choices. Even so, a collective of causal decision theorists are still exploitable even if they start off with the same preferences and the same credences and will all see the same evidence. Evidential decision theorists who consistently update by conditioning are not exploitable in this way.

On your way through downtown, passing by the typical pack of hustlers and con men, something unusual catches your eye. Rather than the old three-card monte, there is a new game on offer:

The One-Box Monte A man presents a closed box containing either \$5 or nothing. The contents have been determined by a prediction machine. You happen to be familiar with this kind of machine, and you are very confident in its predictions. The machine has put \$5 in the box if and only if has predicted that you won't pay the man \$1 to go away. If you don't pay the man \$1 to go away, he'll offer to sell you the contents of the box (unseen) for \$4. You're fairly confident that the box is empty.

Should you pay the man to go away? Of course not: You can walk away from all offers for free. Nonetheless, if you were a causal decision theorist who updated by Bayesian conditioning, you would pay him.¹

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at johan.eric.gustafsson@gmail.com.

¹ This is a simplified version of Ahmed's (2014, pp. 226–30) Psycho-Insurance case. A difference, however, is that the One-Box Mote is BI-terminating (see note 7). Note

Let us explore why causal decision theorists who update by conditioning would fall for this scheme. Causal Decision Theory can be stated as follows:²

Causal Decision Theory Choose an option x such that there is no option y such that $V_{CDT}(y) > V_{CDT}(x)$, where

$$V_{CDT}(x) = \sum_{o \in O} P(x \square \rightarrow o)V(o).$$

Here, O is the set of possible outcomes, $V(o)$ is your utility for outcome o , $P(x \square \rightarrow o)$ is your credence in the subjunctive conditional that, if x were the case, then o would be the case.

According to Bayes's Rule of Conditioning, your credence in a proposition A after learning that an event e has occurred should be equal to your conditional credence in A given e . Let this conditional credence be

$$P(A | e) =_{df} \frac{P(A \& e)}{P(e)}.$$

And let your credence in a proposition A after learning that an event e has occurred be $P_e(A)$. Then we can state the rule as follows:³

The Rule of Conditioning $P_e(A) = P(A | e)$.

That is, your updated credence of A upon learning e should be equal to your conditional probability of A given e .


Let S_1 be the state of nature in which the machine predicted that you would pay the man \$1 to go away. And let S_2 be the complementary state of nature. Throughout, we'll assume that the states of nature in our decision problems are causally independent of the agents choices. We can diagram the One-Box Monte as follows:

that, unlike a three-card monte (Nash 1976, pp. 22–4), this scheme does not require any deception or sleight of hand. Though, it does require a convincing prediction machine, which may be hard to come by.

² Gibbard and Harper 1978, p. 128. See also Cartwright (1979, p. 431), Lewis (1981, pp. 11–12), Joyce (1999, p. 4), and Pearl (2000, p. 108) who state slightly different versions of the theory. These differences, however, won't matter for the arguments of this paper.

³ Bayes 1763, p. 381; 1958, pp. 300–1.

The One-Box Monte



$$P(S_1) > .9, \text{ and } P(S_1 \mid \text{down at node 1}) < .1.$$

Here, the boxes represent choice nodes, and the numbers on the right represent the utility levels for each outcome in each state of nature. (We let utility levels correspond to the monetary pay-offs.) At node 1, you start off being fairly confident that you are in S_1 . And you are very confident that you go up at node 1 if and only if you are in S_1 . If you were to go down at node 1, you would reach node 2. If you were to reach node 2, you would update by the Rule of Conditioning and realize that, since you went down at node 1, you're unlikely to be in S_1 — that is, you are likely to be in S_2 . At node 2, going up gives you a pay-off of -4 if you are in S_1 and 1 if you are in S_2 , whereas going down gives you a 0 pay-off. Accordingly, since you would be fairly confident at node 2 that you are in S_2 , you would go up at that node.

Next, we assume that you rely on *sophisticated choice* — that is, you make choices taking into account what you would choose at future choice nodes with backward induction.⁴ (For myopic and naive choice, see Appendix A.) Accordingly, at node 1, you take into account your prediction that you would go up at node 2 using backward induction. So you find that, if you were to go down at node 1, you would bring about a pay-off of -4 if you are in S_1 and 1 if you are in S_2 . Going up at node 1, gives you a pay-off of -1 in any case. And, since you are fairly confident that you are in S_1 , you go up at node 1.

Hence you end up with -1 for sure, which seems irrational when you could have walked away with 0 for sure. That is, you end up paying the exploiter \$1 to go away, even though you knew that you could have walked away from all offers for free. Hence you are vulnerable to a money pump.⁵

⁴ von Neumann and Morgenstern 1944, pp. 116–17, Pollak 1968, p. 203, and Hammond 1976, p. 162.

⁵ Gustafsson 2022, pp. 1–2.

Could this money pump be blocked by foresight?⁶ No. In fact, what made the scheme work is that you used foresight and relied on backward induction at node 1 given your prediction of what you would do at node 2.⁷ (As we shall see later on, a variation of the case works if you don't rely on backward induction.)

Could causal decision theorists avoid this money pump by adding a ratificationist requirement — that is, a requirement that one won't make a choice that one would instantly regret after you have made it?⁸ They cannot. Note that, according to Causal Decision Theory, choosing to go up is ratifiable at each choice node in this case.

Next, could the money pump be blocked by adopting the Tickle Defence? According to the Tickle Defence, a rational agent knows their own mind. Any information about yourself that you learn from choosing something, you should already possess if you are rational.⁹ So you should already possess whatever information about you which the predictor machine based its prediction on. Accordingly, when you learn that you won't pay the exploiter to go away, you don't learn anything new about the prediction. And, then you could remain confident that you are in S_1 and walk away from all offers.

In addition to other problems with the Tickle Defence, causal decision theorists can't rely on this defence without robbing themselves of one of their main arguments against Evidential Decision Theory — namely, its prescriptions in Newcomb cases. Evidential Decision Theory can be stated as follows:¹⁰

Evidential Decision Theory Choose an option x such that there is no option y such that $V_{EDT}(y) > V_{EDT}(x)$, where

$$V_{EDT}(x) = \sum_{o \in O} P(o | x)V(o).$$

A standard objection to Evidential Decision Theory is that it seemingly

⁶ Schick 1986, pp. 117–18.

⁷ In fact, since this decision problem is BI-terminating, the recommendation of backward induction can be defended with a very minimal form of backward induction. A decision problem is *BI-terminating* if backward induction only prescribes choices that aren't followed by any further potential choices. See Rabinowicz 1998, pp. 97–8, 108–9.

⁸ Egan 2007, pp. 107–8. For ratificationism in general, see Jeffrey 1981, pp. 487–8.

⁹ Skyrms 1980, pp. 130–1 and Eells 1982, pp. 170–4.

¹⁰ Gibbard and Harper 1978, p. 129. See also Jeffrey 1965, pp. 1–6 and Ahmed 2014, pp. 43–6 who define the view in terms of a partition of states rather than outcomes.

prescribes the wrong option in Newcomb cases, such as the following:¹¹

The Smoking Lesion Suppose you believe that smoking is strongly correlated with cancer, but only because there is a common cause — a condition that tends to cause both smoking and cancer. You prefer smoking without cancer to not smoking without cancer. You prefer smoking with cancer to not smoking with cancer. And you strongly prefer not to get cancer.

Evidential Decision Theory seems to prescribe non-smoking since learning that you chose to abstain from smoking is better news than learning that you chose to smoke. But, if the Tickle Defence is available, it may be that rational agents already know if they have an urge to smoke. So they would already possess the information they would get about their condition when they learn what they chose. And then Evidential Decision Theory prescribes smoking.

But, apart from this dialectical disadvantage of the Tickle Defence, it's implausible that a rational agent would need to know everything about themselves which may correlate with both their choices and some state of nature. Since those states of nature may be reliable predictions about what the agent will choose, the agent would need to know in advance what they are going to choose. This leaves very little work for the decision theory to guide choice.¹²

Could you block this money pump by adopting resolute choice? If you rely on *resolute choice*, you stick to a plan you prefer most at the outset.¹³ So, in the One-Box Monte, you stick to the plan to walk away from all offers. But there are well known problems with resolute choice.¹⁴ Moreover, resolutely following the plan to walk away from all offers shouldn't look attractive to causal decision theorists at node 2. At node 2, you would regard it as very likely that walking away (that is, going down) would bring about a less preferred outcome than going up.¹⁵

¹¹ Skyrms 1980, pp. 128–9 and Egan 2007, p. 94. For the original Newcomb case, see Nozick 1963, p. 223; 1969, pp. 114–15.

¹² See Lewis 1981, pp. 10–11, Skyrms 1984, p. 74, and Horwich 1985, pp. 439–41.

¹³ McClennen 1990, pp. 12–13.

¹⁴ See Gustafsson 2022, pp. 66–74 for an overview.

¹⁵ Note that this money pump also works against revised versions of Causal Decision Theory which tell you to maximize the expected causal difference of an option conditional on that it is chosen. (Proposed in Gallow 2020, p. 131 and Barnett 2022, p. 67.) At node 2, you are confident that you are in S_2 no matter what you do. So going up is

Finally, a potential loophole that might save Causal Decision Theory with conditioning is to maintain that, in the One-Box Monte, a high credence in S_1 is irrational.¹⁶ So it's neither Causal Decision Theory nor conditioning that is to blame: your credences are to blame. If vulnerability to money-pumps is irrational, then you cannot believe that you will go up at node 1 unless you believe that you are irrational. And, if you are rational, you do not believe that you are irrational.

This loophole, however, gets the dialectic wrong. The irrationality of vulnerability to money pumps is not something that follows from Causal Decision Theory. Starting out with a high credence in S_1 in the One-Box Monte is fully consistent with following Causal Decision Theory and the Rule of Conditioning. So we can consistently object to Causal Decision Theory combined with the Rule of Conditioning by showing that it conflicts with the irrationality of vulnerability to money pumps. If a decision theory is a correct account of rational decision making and vulnerability to money pumps is irrational, the theory should be able to account for the irrationality of this vulnerability.

So far, we have been concerned with causal decision theorists who update by the Rule of Conditioning.¹⁷ This is crucial, because causal decision theorists avoid exploitation in this case if they, rather than conditioning, update by imaging when they learn what they have chosen.¹⁸ To *image* on e , you transfer your credence in each world W where e is false to the world closest to W where e is true.¹⁹ We then have the following alternative to the Rule of Conditioning:

The Rule of Imaging $P_e(A)$ = the probability of A after imaging on e .

the only option that will make a positive causal difference. Taking this into account at node 1, you find that going down at node 1 will make you roughly 2 units better off in expectation conditional on you choosing to go down at node 1. But you also find that going up at node 1 will make you roughly 3 units better off in expectation conditional on going up at node 1. So you go up at node 1. Unlike standard Causal Decision Theory, this revised version can't be saved by adopting the Rule of Imaging for choices. See note 30.

¹⁶ See Maher 1990, p. 491 for a similar suggestion.

¹⁷ As Lewis (1981, p. 6) suggests.

¹⁸ Cantwell (2010, pp. 142–5) suggests that causal decision theorists should update on their choices by an indicative variant of imaging and on other news by conditioning. Pearl (2000, p. 23; 2021, pp. 427–8) suggests using a similar variant of imaging for updating on choices and using conditioning for updating on observations.

¹⁹ Lewis 1976, pp. 310–11.

Crucially, imaging on your choices lets you retain your credences in each state of nature after you have learned what you've chosen. (We partition states of nature so that they are causally independent of your choices.) So, in the One-Box Monte, you then retain your high credence in S_1 even if you were to reach node 2. So then you would go down at node 2 if you were to reach that node. Taking this prediction into account at node 1, you go down and end up walking away from all offers.

It may be objected that, if you update by anything other than the Rule of Conditioning, you are vulnerable to a Dutch book, which is also kind of money pump.²⁰ So adopting the Rule of Imaging for learning from one's choices would still leave causal decision theorists open to exploitation. Or so it may seem.

In fact, the Dutch-book argument for the Rule of Conditioning only works if what you learn is something other than your choice. To see this, we will first consider the Dutch-book argument for conditioning when you learn by observing a chance event. This argument is often put in terms of series of bets such that the agent accepts them all but will then faces a certain loss. Since we assume that the agent relies on backward induction, the argument can be simplified considerably.²¹ All bets except one can be replaced by an initial offer to pay the exploiter to go away.²²

Suppose that $P(e) > 0$ and that $P_e(A) < P(A | e)$. (The case of $P_e(A) > P(A | e)$ can be handled in a similar manner; see Appendix B.) Let ϵ be an amount such that

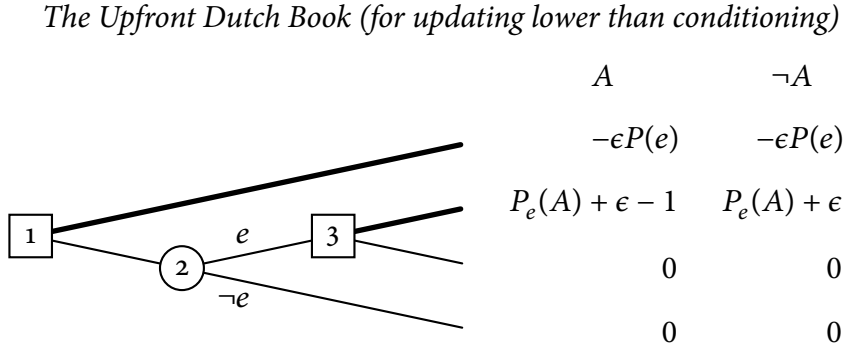
$$0 < \epsilon < \frac{P(A | e) - P_e(A)}{2}.$$

²⁰ Putnam 1967, p. 113, Teller 1973, pp. 222–5, Lewis 1999, pp. 405–6, and Skyrms 1993, pp. 321–5. This is one of Cantwell's (2010, p. 145) main worries about using the Rule of Imaging for updating on your choices. One advantage is, as Cantwell (2010, p. 144) points out, that it can explain why it makes sense, in Egan's (2007, p. 97) psychopath-button case, to push the button rati-fiedly without regret. Similarly, we get rati-fiable choices in Gibbard and Harper's (1978, pp. 157–9) death-in-Damascus case.

²¹ If the agent relies on backward induction, Teller (1973, pp. 222–5) and Lewis's (1999, pp. 405–6) Dutch book doesn't work, as Maher (1992, pp. 124–5) shows.

²² It may be objected that the term 'Dutch book' should be reserved to schemes that only rely on combinations of bets that bring about a sure loss. We could conceive of the upfront payment to the exploiter as a bet against the tautology or paying for a bet that has a zero pay-off in any event. But I see little reason to do so. Once we allow that the agent relies on backward induction, the agent needn't assess the bets individually anyway. Besides, if you don't think that the term 'Dutch book' applies here, consider the Upfront Dutch Book as a money pump. The upshot remains the same.

And suppose that, before you learn whether e , you can pay an exploiter the amount $\epsilon P(e)$ to go away. If you don't do so and then learn that e , the exploiter will offer to pay you $P_e(A) + \epsilon$ for the following bet: -1 if A ; otherwise 0 . We can diagram the decision problem as follows:²³



$$P(e) > 0, P_e(A) < P(A | e), \text{ and } 0 < \epsilon < \frac{P(A | e) - P_e(A)}{2}.$$

Here, the circle represents a chance node where you learn whether e . Suppose that the choices at nodes 1 and 3 are causally and evidentially independent of A .

If you were to reach node 3, you would have updated your credences after learning e so that your credence in A would be $P_e(A)$. Given your updated credence, you find that accepting the bet offered at node 3 has an expected pay-off of ϵ . So you would accept that bet if you were to reach node 3. Taking this into account at node 1, you find that not paying the exploiter upfront to go away (that is, going up) has an expected pay-off of

$$P(e)P(A | e)(P_e(A) + \epsilon - 1) + P(e)(1 - P(A | e))(P_e(A) + \epsilon) = -P(e)(P(A | e) - P_e(A) - \epsilon).$$


And, given our constraints on ϵ , this pay-off is lower than $-\epsilon P(e)$. So you go up at node 1. That is, you pay the exploiter $\epsilon P(e)$, even though you

²³ This upfront Dutch book simplifies Skyrms's (1993, pp. 321–5) Dutch book in the same way Gustafsson and Rabinowicz's (2020, p. 583) upfront money pump simplifies Rabinowicz's (2000, pp. 140–1) earlier money pump with foresight. In both cases, the upfront variants are BI-terminating whereas the non-upfront ones are not (see note 7). As noted by Gustafsson and Rabinowicz (2020, p. 583n9), however, the disadvantage of not being BI-terminating is less clear for Skyrms's Dutch book than for Rabinowicz's money pump with foresight.

could have walked away from all offers for free.²⁴

This argument for conditioning, however, supposes that what you learn isn't your choice. Suppose we replace the chance node in the Upfront Dutch Book with a choice node:

*The Upfront Dutch Book (for updating lower than conditioning),
attempted for choices*



$$P_e(A) < P(A | e), \text{ and } 0 < \epsilon < \frac{P(A | e) - P_e(A)}{2}.$$

Like before, you predict that you would go up at node 3 if you were to reach that node. If you were to reach node 2 (now, a choice node), you would assess whether to make e true or false. But you would do this differently depending on whether you follow Evidential or Causal Decision Theory.

On the one hand, if you follow Evidential Decision Theory, you find that the value of going up at node 2 will be the same as the value of going down at node 1 in the original, chance-node version of the Upfront Dutch Book. And, as we saw earlier, that value is negative.

On the other hand, if you follow Causal Decision Theory, you disregard the news you would get from learning that you have made the choice. Hence you evaluate the prospect of going up at node 2 using your unconditional credence in A . Given that $P(A)$ is sufficiently high, you find that the expectation of going up at node 2 is negative.

So, on both Evidential and Causal Decision Theory, the expectation of going up at node 2 is negative. Hence, if you follow Evidential or Causal

²⁴ Maher (1993, pp. 110–13) objects (in response to Skyrms's diachronic Dutch book) by allowing that it need not be irrational to accept a sure loss, since the alternative may be to risk a greater loss. But, in this case (and in Skyrms's case), one alternative to accepting a sure loss is to walk away from all offers — a prospect whose potential outcomes are each preferred to every potential outcome of accepting a sure loss.

Decision Theory, you would go down at node 2 if you were to reach that node. Taking this prediction into account at node 1, you refuse to pay the exploiter and walk away from all offers. (Note that if you were to become certain at node 1 that you would go down at node 2, the price of going up at node 1 would be zero. So then you could just as well walk away by going up as going down at node 1.) Hence the Upfront Dutch Book does not support the Rule of Conditioning when you learn what you have chosen.²⁵

Nevertheless, even though the Upfront Dutch Book doesn't work as a general argument for the Rule of Conditioning, it may still work against those who use the Rule of Imaging when they learn what they have chosen. If you update by imaging when you learn what you have chosen at node 2, then your credence in A remains unchanged. That is, $P(A) = P_e(A)$. And then, taking into account the prediction that you would go up at node 3, you would go up at node 2 (since the utility of going up at node 2 is then the same as the utility of going up at node 3). Then, at node 1, you take these predictions into account — that is, that you would go up at nodes 2 and 3. Next, you take the evidence you get from your prediction that e will occur into account and conclude that the probability of A , conditional on going down at node 1, is equal to $P(A | e)$. Hence you take the expected pay-off of going down at node 1 to be

$$P(A | e)(P_e(A) + \epsilon - 1) + (1 - P(A | e))(P_e(A) + \epsilon) = \epsilon + P_e(A) - P(A | e),$$

which, given our constraints on ϵ , is less than $-\epsilon$, that is, less than $-\epsilon P(e)$ with $P(e) = 1$ from your prediction that you would go up at node 2. Hence you go up at node 1 and pay ϵ , even could have walked away for free.

So causal decision theorists who update by imaging when they learn what they have chosen but by conditioning when they learn other things are open to the Upfront Dutch Book. The culprit, however, needn't be the Rule of Imaging nor Causal Decision Theory, but rather your taking the evidence from your predicted choice at node 2 into account at node 1 even though you would ignore this evidence at node 2 (since you are a causal decision theorist) and you would also ignore it at node 3 (since

²⁵ This objection also rebuts Skyrms's (1993, pp. 321–5) more complicated Dutch book if it's applied to updates on the agent's choices. The exact same reasoning applies in his set-up at the node determining e in case you were not to accept the initial bets (in case your probability in A at that node is sufficiently high).

you update with imaging when you learn what you have chosen). So the theory needs one further modification to avoid this Dutch book: causal decision theorists need to consistently ignore evidential news from their choices before, during, and after their choices. In other words, causal decision theorists need to ignore the evidential evidence from their choices throughout dynamic decision problems. They must ignore evidential evidence not just from their choice at the current node but also their choices at future choice nodes. They need not only adopt the Rule of Imaging for learning from choices but also adopt a plan-based form of Causal Decision Theory. Given the standard choice-based form of Causal Decision Theory, one still takes evidential evidence from one's future choices into account. The plan-based form of Causal Decision Theory can be stated as follows:²⁶

Plan-Based Causal Decision Theory Choose an option x such that x is the initial segment of an available plan s such that there is no available plan s' such that $V_{CDT}(s') > V_{CDT}(s)$, where

$$V_{CDT}(s) = \sum_{o \in O} P(s \square \rightarrow o) V(o).$$


Here, an available plan is a specification of choices at the current choice node and the later choice nodes that the agent might reach by following the plan at earlier nodes. So a plan can be available in this sense even if you predict that you will depart from the plan after having made the first choice of the plan.

For a simpler case that shows the need for causal decision theorists to adopt Plan-Based Causal Decision Theory, consider the following variation of the One-Box Monte, where the prediction machine has put \$5 in the box if and only if you will turn down the offer to buy its contents:²⁷

²⁶ Rothfus 2022, p. 267.

²⁷ This is a simplified binary variation of Spencer's (2021, p. 55) Two Rooms case, which in turn is a sequential variation of Spencer and Wells's (2019, p. 34) Frustrater case.

The Binary Frustrater



$P(S_1) > .9$, and $P(S_1 \mid \text{up at node 2}) < .1$.

At node 2, you would (after having imaged on going down at node 1) still be fairly confident that you're in S_1 . So you would go up at node 2.

Taking this prediction into account at node 1, you find that, if you were to go down at node 1, you would probably be in S_2 (since you, following standard choice-based Causal Decision Theory, do take into account conditional probabilities of your predicted choices at future choice nodes). So you regard value the outcome of going down at node 1 as having an expected pay-off of roughly -4 . Hence you go up at node 1.

Plan-Based Causal Decision Theory, however, will regard the plan to go down at node 1 and up at node 2 as the best (evaluating plans holding fixed one's high credence in S_1). So, following Plan-Based Causal Decision Theory, you go down at node 1 and then up at node 2. Hence you avoid exploitation.

Putting this together, we have the following form of Causal Decision Theory:

Dynamic Causal Decision Theory Plan-Based Causal Decision Theory with the Rule of Imaging for updating on having made a choice and the Rule of Conditioning for updating on other news.


Dynamic Causal Decision Theory avoids the Dutch books and money pumps we have discussed, since — following this view — agents consistently ignore evidential news from their choices before, during, and after the choices.²⁸

Similarly, evidential decision theorists who always update by the Rule of Conditioning avoid these Dutch books and money pumps, since they

²⁸ Note that, since Dynamic Causal Decision Theory relies on imaging for updating on having made a choice, it is not open to Rothfus's (2022, pp. 270–271) Two-Drawers objection to Plan-Based Causal Decision Theory.

consistently take into account evidential news from their choices before, during, and after the choices.²⁹ Evidential decision theorists who rely on imaging when they learn their choices will be dynamically inconsistent, however: they take the evidential news from their choices into account before the choices but not after. And, indeed, they are open to a money pump. Consider the following flipped variant of the One-Box Monte, which is the same except that the machine has put \$5 in the box if and only if it predicted that you would pay the man \$1 to go away:

The Flipped One-Box Monte



$$P(S_1) > .9, \text{ and } P(S_1 \mid \text{down at node 1}) < .1.$$

At node 1, you're fairly confident that you're in S_0 . But, if you update by imaging, you would remain confident that you're in S_1 if you were to learn that you went down at node 1. So, if you were to reach node 2, you would go up at that node. At node 1, you're fairly confident that, conditional on going down at node 1, you're in S_2 . At node 1, assuming that you rely on sophisticated choice, you take into account your prediction that you would go up at node 2. (For myopic and naive choice, see Appendix A.) So, at node 1, you find that the evidential value of going down is close to -4 . Hence, following Evidential Decision Theory combined with updating by imaging for choices, you go up at node 1 and pay the exploiter — even though you could have walked away from all offers for free.³⁰

²⁹ Arntzenius (2008, pp. 289–90) raises a money-pump puzzle for Evidential Decision Theory, but see Ahmed and Price 2012, pp. 23–7 for a solution. Evidential Decision Theory is, however, dynamically inconsistent in the sense that it may require you to change preference between two plans at different nodes in a decision problem. See Rothfus 2020, pp. 3928–30.

³⁰ This money pump also works against the revised versions of Causal Decision Theory which tell you to maximize the expected causal difference of an option conditional on that it is chosen (see note 15). At node 2, having updated with imaging rather than conditioning, you remain confident that you're in S_1 . So going up is the only option that will make a positive causal difference. Taking this into account at node 1, you find that

As we noted earlier, agents who follow Dynamic Causal Decision Theory are not, individually, open to exploitation. But what about two or more such agents who have the same values and preferences and start off having the same credences and are then exposed to the same evidence? Consider, once more, the One-Box Monte, except this time the choices at nodes 1 and 2 are taken by different agents. These agents have the same values and preferences and they start off with the same credences. If the first agent were to go down at node 1, the second agent would update by conditioning and become fairly confident that they are in S_2 . Note that, if node 2 were reached, the agents would no longer have the same credences or the same preferences: the first agent would prefer going down at node 2, whereas the second agent would prefer going up. So the second agent would go up at node 2 if that node were reached. Taking this prediction into account at node 1, the first agent finds that going down at node 1 has an expected pay-off of roughly -4 . Hence the first agent goes up at node 1 and pays the exploiter, even though both agents prefer that they had walked away from all offers for free.^{31,32}

Evidential decision theorists who consistently follow the Rule of Conditioning are not open to this kind of collective exploitation.³³ If such

going down at node 1 will make you roughly 2 units better off in expectation conditional on going up at node 1. But you also find that going up at node 1 will make you roughly 3 units better off in expectation conditional on going down at node 1. So you go up at node 1.

³¹ Note that this is a worse result than the familiar problem that rational agents who have different preferences may be exploited in a prisoner's dilemma, which relies on the agents having different preferences. (See Luce and Raiffa 1957, pp. 94–5 and Tucker 1980, p. 101.) Here, the problem arises even though the agents start off agreeing about everything and then get the same evidence.

³² It may be objected that causal decision theorists could avoid this problem if they updated by imaging on not only their choices but also on their collaborators' choices. But this raises the problem of where to draw the line. It would be absurd not to learn by conditioning on some peoples choices. Although, a problem with a sharp distinction between one's own choices and those of others is how to deal with cases where the relation that matters in survival seem to depart from identity and come in degrees. (See Parfit 1971, pp. 10–11.) A potential solution is update with a weighted average of imaging and conditioning, where the weight for imaging is your degree to which you are related to the future agent and the weight for conditioning is one minus the weight for imaging.

³³ It may be objected that they will choose dominated options in Newcomb cases like the Smoking Lesion. Isn't that as bad? It's not. Note that in a Newcomb case the dominated option is only dominated in some partitions of the states of nature — not in all. So evidential decision theorist can respond that the relevant dominance principle should be concerned with statewise dominance relative to some other partition of the


agents start off having the same credences and the same preferences, they will stay in agreement if they get the same evidence. Since they update the same way on other agents choices as on their own choices, they will handle collective cases the same as if they individually made all the choices.

Appendices

A. Myopic and Naive Choice

We have assumed that agents rely on sophisticated choice. But what if the agents rely on myopic choices or naive choice? *Myopic choice* is to consider offers in isolation under the assumption that one would reject all future offers (that is, to go down at all future choice nodes in our diagrams).³⁴ *Naive choice* is to (i) consider the outcomes of all available plans and assess which of these outcomes are choice-worthy in a choice between all of them and (ii) choose in accordance with a plan to end up with a choice-worthy outcome, without taking into consideration whether one would later depart from that plan.³⁵

Given either myopic or naive choice, a causal decision theorist would go down at node 1 of the One-Box Monte. And then the money pump is blocked. But consider the following revised case:



$$P(S_1) > .9, \text{ and } P(S_1 \mid \text{up at node 1}) < .1.$$

Here, as a causal decision theorist who relies on either myopic or naive choice, you would go up at node 1, since you are fairly confident that you

states. (See Nozick 1969, pp. 120–1.) In the money-pump cases we have discussed, causal decision theorists can't make this move, since paying the exploiter is certainly worse than walking away. Hence paying the exploiter is statewise dominated relative to all partitions of states of nature.


³⁴ See Dow 1984, p. 96 and McClennen 1990, pp. 11–12.

³⁵ See Pollak 1968, pp. 202–3 and Hammond 1976, p. 162.

are in S_1 and the most preferred outcome given S_1 is achieved by going up at node 1 and then down at node 2. But then, at node 2 after updating on having gone up at node 1 with the Rule of Conditioning, you become fairly confident that you are in S_2 . So you then go up at node 2. Hence you're still open to a money pump.

A similar objection applies to the money-pump argument against choice-based causal decision theorists who rely on imaging. If such agents are myopic or naive, they are not vulnerable to the Binary Frustrater. But consider the following myopic variation:

The Myopic Binary Frustrater




$$P(S_1) > .9, \text{ and } P(S_1 \mid \text{down at node 2}) < .1.$$

Here, as a causal decision theorist who are myopic or naive will go up at node 1, since, at that node, the best plan is to go up at node 1 and then down at node 2. This plan has an expected pay-off of roughly 1 (since you take into account the evidential evidence of going down at node 2. But, when you reach node 2, you have imaged on going up at node 1, and are still fairly confident that you are in S_1 . So you go down at node 2.

Another objection of this kind applies to the money pump against evidential decision theorists who update with imaging on their choices. Given either myopic or naive choice, an evidential decision theorist would go down at node 1 of the Flipped One-Box Monte. But we can revise that case as follows:

The Myopic Flipped One-Box Monte



$P(S_1) > .9$, and $P(S_1 \mid \text{up at node 1}) < .1$.

You're fairly confident that, conditional on going up at node 1, you are in S_2 and the most preferred outcome given S_2 is then achieved by then down at node 2 — and this outcome is more preferred than the outcome of going down at node 1. So, as an evidential decision theorist who relies on either myopic or naive choice, you would go up at node 1. Then, at node 2 after updating on having gone up at node 1 with the Rule of Imaging, you remain confident that you are in S_1 . So you go up at node 2. And then you are still open to a money pump.³⁶

B. Updating Higher than Conditioning


The version of the Upfront Dutch Book we discussed earlier only covers the case of $P_e(A) < P(A \mid e)$. But the case of $P_e(A) > P(A \mid e)$ can be handled in a similar manner. Let ϵ be an amount such that

$$0 < \epsilon < \frac{P_e(A) - P(A \mid e)}{2}.$$

And suppose that, before you learn whether e , you can pay an exploiter the amount $\epsilon P(e)$ to go away. If you don't do so and then learn that e , the exploiter will offer you to pay $P_e(A)$ for the following bet: $1 + \epsilon$ if A ; otherwise ϵ . We can diagram the decision problem as follows:

³⁶ For a version of the Dutch-book argument for the Rule of Conditioning which works for agents who rely on myopic or naive choice, see Teller 1973, pp. 222–5 and Lewis 1999, pp. 405–6.

The Upfront Dutch Book (for updating higher than conditioning)



$$P(e) > 0, P_e(A) > P(A | e), \text{ and } 0 < \epsilon < \frac{P_e(A) - P(A | e)}{2}.$$

Suppose that the choices at nodes 1 and 3 are causally and evidentially independent of A .

If you were to reach node 3, you would have updated your credences after learning e so that your credence in A would be $P_e(A)$. Given your updated credence, you find that accepting to pay for the bet has an expected pay-off of ϵ . So you would go up at node 3 if you were to reach that node. Taking this into account at node 1, you find that not paying the exploiter upfront to go way (that is, going up) has an expected pay-off of

$$P(e)P(A | e)(-P_e(A) + 1 + \epsilon) + P(e)(1 - P(A | e))(-P_e(A) + \epsilon) = -P(e)(P_e(A) - P(A | e) - \epsilon).$$

Given our constraints on ϵ , this pay-off is lower than $-\epsilon P(e)$. So you go up at node 1. And you pay the exploiter $\epsilon P(e)$, even though you could have walked away from all offers for free.

I wish to thank Arif Ahmed, Adam Bales, John Cantwell, Kenny Easwaran, Philip A. Ebert, Tomi Francis, Melissa Fusco, Daniel Herrmann, Petra Kosonen, Sebastian Krug, Harvey Lederman, Martin Peterson, Wlodek Rabinowicz, Gerard J. Rothfus, Martin Smith, Dean Spears, and H. Orri Stefánsson for valuable comments.

References

- Ahmed, Arif (2014) *Evidence, Decision and Causality*, Cambridge: Cambridge University Press.
- Ahmed, Arif and Huw Price (2012) ‘Arntzenius on ‘Why Ain’cha Rich?’’, *Erkenntnis* 77 (1): 15–30.
- Arntzenius, Frank (2008) ‘No Regrets, or: Edith Piaf Revamps Decision Theory’, *Erkenntnis* 68 (2): 277–297.
- Barnett, David James (2022) ‘Graded Ratifiability’, *The Journal of Philosophy* 119 (2): 57–88.
- Bayes, Thomas (1763) ‘An Essay towards Solving a Problem in the Doctrine of Chances’, *Philosophical Transactions* 53 (1): 370–418.
- (1958) ‘An Essay towards Solving a Problem in the Doctrine of Chances’, *Biometrika* 45 (3–4): 296–315.
- Cantwell, John (2010) ‘On an Alleged Counter-Example to Causal Decision Theory’, *Synthese* 173 (2): 127–152.
- Cartwright, Nancy (1979) ‘Causal Laws and Effective Strategies’, *Noûs* 13 (4): 419–437.
- Dow, Gregory K. (1984) ‘Myopia, Amnesia, and Consistent Intertemporal Choice’, *Mathematical Social Sciences* 8 (2): 95–109.
- Eells, Ellery (1982) *Rational Decision and Causality*, Cambridge: Cambridge University Press.
- Egan, Andy (2007) ‘Some Counterexamples to Causal Decision Theory’, *The Philosophical Review* 116 (1): 93–114.
- Gallow, J. Dmitri (2020) ‘The Causal Decision Theorist’s Guide to Managing the News’, *The Journal of Philosophy* 117 (3): 117–149.
- Gibbard, Allan and William L. Harper (1978) ‘Counterfactuals and Two Kinds of Expected Utility’, in C. A. Hooker, J. J. Leach, and E. F. McClennen, eds., *Foundations and Applications of Decision Theory*, vol. I, pp. 125–162, Dordrecht: Reidel.
- Gustafsson, Johan E. (2022) *Money-Pump Arguments*, Cambridge: Cambridge University Press.
- Gustafsson, Johan E. and Wlodek Rabinowicz (2020) ‘A Simpler, More Compelling Money Pump with Foresight’, *The Journal of Philosophy* 117 (10): 578–589.
- Hammond, Peter J. (1976) ‘Changing Tastes and Coherent Dynamic Choice’, *The Review of Economic Studies* 43 (1): 159–173.
- Horwich, Paul (1985) ‘Decision Theory in Light of Newcomb’s Problem’, *Philosophy of Science* 52 (3): 431–450.

- Jeffrey, Richard C. (1965) *The Logic of Decision*, New York: McGraw-Hill.
- (1981) ‘The Logic of Decision Defended’, *Synthese* 48 (3): 473–492.
- Joyce, James M. (1999) *The Foundations of Causal Decision Theory*, Cambridge: Cambridge University Press.
- Lewis, David (1976) ‘Probabilities of Conditionals and Conditional Probabilities’, *The Philosophical Review* 85 (3): 297–315.
- (1981) ‘Causal Decision Theory’, *Australasian Journal of Philosophy* 59 (1): 5–30.
- (1999) ‘Why Conditionalize?’, in *Papers in Metaphysics and Epistemology*, pp. 403–407, Cambridge: Cambridge University Press.
- Luce, R. Duncan and Howard Raiffa (1957) *Games and Decisions: Introduction and Critical Survey*, New York: Wiley.
- Maher, Patrick (1990) ‘Symptomatic Acts and the Value of Evidence in Causal Decision Theory’, *Philosophy of Science* 57 (3): 479–498.
- (1992) ‘Diachronic Rationality’, *Philosophy of Science* 59 (1): 120–141.
- (1993) *Betting on Theories*, Cambridge: Cambridge University Press.
- McClennen, Edward F. (1990) *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.
- Nash, Jay Robert (1976) *Hustlers and Con Men: An Anecdotal History of the Confidence Man and His Games*, New York: M. Evans and Company.
- Nozick, Robert (1963) *The Normative Theory of Individual Choice*, Ph.D. thesis, Princeton University.
- (1969) ‘Newcomb’s Problem and Two Principles of Choice’, in Nicholas Rescher, ed., *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of his Sixty-Fifth Birthday*, pp. 114–146, Dordrecht: Reidel.
- Parfit, Derek (1971) ‘Personal Identity’, *The Philosophical Review* 80 (1): 3–27.
- Pearl, Judea (2000) *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.
- (2021) ‘Causal and Counterfactual Inference’, in Markus Knauff and Wolfgang Spohn, eds., *The Handbook of Rationality*, pp. 427–438, Cambridge, MA: MIT Press.
- Pollak, R. A. (1968) ‘Consistent Planning’, *The Review of Economic Studies* 35 (2): 201–208.
- Putnam, Hilary (1967) ‘Probability and Confirmation’, in Sidney Morgenbesser, ed., *Philosophy of Science Today*, pp. 100–114, New York: Basic Books.
- Rabinowicz, Wlodek (1998) ‘Grappling with the Centipede: Defence of

- Backward Induction for BI-Terminating Games, *Economics and Philosophy* 14 (1): 95–126.
- (2000) ‘Money Pump with Foresight’, in Michael J. Almeida, ed., *Imperceptible Harms and Benefits*, pp. 123–154, Dordrecht: Kluwer.
- Rothfus, Gerard J. (2020) ‘Dynamic Consistency in the Logic of Decision’, *Philosophical Studies* 177 (12): 3923–3934.
- (2022) ‘A Plan-Based Causal Decision Theory’, *Analysis* 82 (2): 264–272.
- Schick, Frederic (1986) ‘Dutch Bookies and Money Pumps’, *The Journal of Philosophy* 83 (2): 112–119.
- Skyrms, Brian (1980) *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*, New Haven, CT: Yale University Press.
- (1984) *Pragmatics and Empiricism*, New Haven, CT: Yale University Press.
- (1993) ‘A Mistake in Dynamic Coherence Arguments?’, *Philosophy of Science* 60 (2): 320–328.
- Spencer, Jack (2021) ‘An Argument against Causal Decision Theory’, *Analysis* 81 (1): 52–61.
- Spencer, Jack and Ian Wells (2019) ‘Why Take Both Boxes?’, *Philosophy and Phenomenological Research* 99 (1): 27–48.
- Teller, Paul (1973) ‘Conditionalization and Observation’, *Synthese* 26 (2): 218–258.
- Tucker, A. W. (1980) ‘A Two-Person Dilemma’, *The UMAP Journal* 1 (1): 101.
- von Neumann, John and Oskar Morgenstern (1944) *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.