

Newcomb University: A Play in One Act

Adam Elga

August 19, 2019. Please cite published version: *Analysis* 80(2), 2020, [212–221](#).

Scene 1

Professor Causey (in full tweed) and Ms. Neutra (a potential donor) walk down the hall of a gleaming research building.

CAUSEY: . . . and so you see that our experimental causal decision-making lab has state-of-the-art facilities which – pardon me for a moment.

Causey reaches down to what looks like an electrical outlet along the wall and sticks his finger in the socket. A buzzing sound is heard.

CAUSEY: Ouch! Anyway, as I was saying, the facilities are excellent—

NEUTRA: Sorry to interrupt, but are you alright?

CAUSEY: Oh, that. I'm fine. Just a modest electrical shock from putting my finger in the outlet. No permanent damage. Ah, I see we are approaching another outlet. Excuse me.

Causey puts his finger in another outlet and another buzzing sound is heard.

CAUSEY: Yowza!

NEUTRA: Why do you keep putting your finger in those outlets?

CAUSEY: To reduce the amount of electrical shocking I experience, of course.

Neutra pauses.

NEUTRA: I'm confused. It looks like every time you put your finger in an outlet, you get painfully shocked.

CAUSEY: Yes.

NEUTRA: But I and the other passers-by who don't touch the outlets don't get shocked.

CAUSEY: Yes.

NEUTRA: And you expect that pattern to continue?

CAUSEY: Yes.

NEUTRA: So why stick your finger in?!

CAUSEY: Didn't I tell you a moment ago? To reduce getting shocked. Or perhaps you don't know that we have special outlets in this building? They are called Libet-Ahmed-Newcomb¹ outlets, LANs for short. Here's how they operate. Each outlet has a brain scanner. Whenever someone approaches the outlet, the outlet scans nearby brains to determine whether the person approaching the outlet is likely to stick her finger in. If so, the outlet activates program A. Otherwise it activates program B.

NEUTRA: What do those programs do?

CAUSEY: Program A has the outlet deliver a shock upon being touched, but if it is not touched then it remotely delivers a slightly stronger shock to the person who approached. Program B has the outlet never deliver any shocks, but has it play a pleasing melody if someone touches it. Here's a decision table for you:

	Program A (LAN predicted touch)	Program B (LAN predicted no touch)
<i>no touch</i>	worse shock	nothing
<i>touch</i>	regular shock	nice melody

As you can see, if the outlet is running Program A, one is better off touching it (to get a regular shock rather than a worse shock). If the outlet is running Program B, one is also better off touching it (to hear the nice melody). And by the time one makes one's decision, the outlet has irrevocably chosen its program. So one is guaranteed to do better by touching the outlet.

NEUTRA: Fascinating.

CAUSEY: Yes, the original funder of the building, Dr. Newcomb, had some eccentric requirements. Ah, I see we are approaching an outlet. *Reaches down and is zapped by it.* Augghh! I suppose that now that you know how these outlets operate, you'll be sticking your finger in?

NEUTRA: No thank you. *Walks past the outlet without incident.*

CAUSEY: Oh. Quite irrational of you. A shame – you missed out on hearing a wonderful melody.

NEUTRA: You've heard it?

CAUSEY: Actually, no. The brain-scanners in these outlets are infernally accurate. To my knowledge, no one has yet heard the melody. And no one has yet been shocked by the stronger shock.

¹See Libet (2005), Ahmed (2014, Chapter 4), Nozick (1969).

NEUTRA: But you expect one of those things to happen soon?

CAUSEY: No.

NEUTRA: I see. *Pauses.* Is it not a workplace safety issue, to have outlets around with the potential to remotely shock people? I notice that other students, not to mention children and some pet cats, wander through the building.

CAUSEY: We were worried about that, but it turned out to not be a problem. The cats rarely touch the outlets, and when one does, it gets shocked and takes care to not touch outlets in the future.² And incoming freshmen are informed about how the outlets are designed. But they are by and large irrational – they don't put their fingers in. (Haven't had the benefit of my *Advanced Causal Decision Theory* course yet, ha ha!)

NEUTRA: So children, animals, and freshman are safe and don't get shocked. It is just shock-avoidance experts such as yourself who stick their fingers in the outlets and receive shocks.³

CAUSEY: Yes.

NEUTRA: *Blinks.*

CAUSEY: The strange thing is, so far as their outward appearance and observed behaviour goes, these LANs look just like ordinary electrical outlets. In fact, only last week I learned that the outlets one floor up from here *are* ordinary outlets. Turns out I had been sticking my fingers in them completely unnecessarily! Now I know not to do it.

NEUTRA: Thanks very much for the tour.⁴

²Compare to the claim concerning Newcomb's problem: "As surely as water will wet us, as surely as fire will burn," so too will one-boxing enrich us, and two-boxing will certainly not.' (Ahmed 2018, 72).

³Neutra's comment here echoes the so-called "Why ain't cha rich?" objection to causal decision theory (Lewis (1981b), Joyce (1999, 151-4), Bales (2017, §3.1), Ahmed (2018)).

⁴In response to the seemingly self-punishing nature of Causey's actions, causal decision theorists have various lines of reply available, many familiar from the debate about Newcomb's problem. These include claiming: (1) In the above situation, *all* of the outcomes typically available to agents like Causey involve being shocked, and so it is no embarrassment to choose a regular shock over a worse one. (Opponents might put forward a competing notion of "available outcome" that blocks this argument – see the definition of "E-opportunity" in Ahmed (2018, §4.1).) (2) Everyone should agree that if one expects to be near many LANs, it would be beneficial to transform oneself into the sort of person who does not stick their finger in LANs. Readers might be count Causey as irrational because he fails to effect such a transformation – or because he continues to work at a department chock full of LANs. These criticisms of Causey may be fair, but even if true they do not count against causal decision theory. (3) The misleading appearance and fanciful construction of LANs invites readers to insufficiently distinguish them in imagination from ordinary outlets, and so treat Causey's action as if it were sticking one's finger in an ordinary outlet. (Opponents might reply that the reason it is silly to stick one's finger in an ordinary outlet is that ordinary outlets are such that usually those who stick their fingers in them get shocked and those who refrain do not, and that LANs have that same property.)

Scene 2

It is one week later. Neutra enters a building marked 'Evidential altruism lab' and knocks on an office door.

EVIDENTIA: Come in. Ah, Ms. Neutra. Nice to meet you, I am Professor Evidentia. Welcome! I take it that last week you met Professor Causey and heard about his... theories?

NEUTRA: Yes. He kept shocking himself. *Shudders*

EVIDENTIA: So sad. Anyway, you'll find none of that today. Here at the evidential altruism lab we always choose so as to *minimize* the expected number of shocks people receive.

NEUTRA: That is a relief.

EVIDENTIA: Speaking of altruism, I will now subject you to a painful shock. *Takes out an electric stun gun and prepares to fire it.*

NEUTRA: Wait! Stop! What are you doing?

EVIDENTIA: Shocking you. For your own sake, of course. To reduce the expected number of shocks you receive.

NEUTRA: *Pause.* OK, I get it. You mean that shocking me now will prevent me from getting some even worse shock in the future? Then I'm sorry to hear about the whole situation, but I guess I'll just grit my teeth and bear this one shock for the sake of its future good effects.

EVIDENTIA: *Adjusting the stun gun.* No, you've misunderstood. The shock I'm about to administer won't cause any good effects at all, to you or anyone else. Its only relevant causal consequence will be the pain you feel.

NEUTRA: No good effects? Pure pain for no reason? How can this help me?

EVIDENTIA: Perhaps Professor Causey didn't explain all the features of our building. Do you see that device? *Points up to what looks like a smoke detector on the ceiling.* That is one of our Decision Prediction Units (DPUs). Last week, I noted in my calendar that you would be visiting me today. In response, the DPU scanned our brains to predict whether I would shock you. If it predicted that I would shock you, it left you alone last week. But if it predicted that I would refrain, then it ensured that last week you received *twenty* shocks, each as painful as the one delivered by this stun gun. Here's a decision table for you, below. I've crossed out those outcomes I'm almost certain don't obtain:

	DPU predicted shock	DPU predicted no shock
<i>shock today</i>	1 shock total	21 shocks total
<i>no shock today</i>	0 shocks total	20 shocks total

NEUTRA: That is horrifying! But there’s nothing to be done *now* about shocks I may or may not have received last week, so that information is not relevant to your current decision.

EVIDENTIA: Actually it is highly relevant. For I know that the DPUs are excellent predictors. So (1) if I shock you now, you will almost certainly end up experiencing one shock in total. And (2) if I refrain now, you will almost certainly end up experiencing twenty shocks in total (delivered last week).⁵ So shocking you reduces the total number of shocks I expect that you receive.⁶ Rationality, altruism, and simple kindness all demand that I shock you.

NEUTRA: Well, (1) and (2) are *not* both true.⁷ And I assure you: *shocking me now cannot help me!*

EVIDENTIA: Well, from *your* point of view I can see why you are saying that. But you must understand that as an evidential altruist, I must act in the light of *my* evidence.

NEUTRA: OK, I see where this is going – I’m pretty sure what you are going to decide. But let me try to reason with you anyway.⁸

⁵These are what might be called “Horgan conditionals”, following Horgan (1981) (though note that Horgan (2015, 7) prefers an argument for one-boxing that does not employ such conditionals).

⁶If it is thought that “thank goodness that’s over” (Prior (1959), Parfit (1984, 165-7)) considerations enter here, we may suppose that receiving 20 shocks has severe long term after-effects, so that Ms. Neutra would prefer (and Professor Evidentia would altruistically prefer on her behalf) to be about to receive 1 shock over having received 20 shocks already. If concerns about harming someone without permission are thought to be relevant, we may suppose that receiving 20 shocks is such a harmful outcome that it is morally permissible to shock someone once without their consent in order to save them from being shocked 20 times.

⁷Here Ms. Neutra assumes that she may in conversation correctly deny a conditional that is false when evaluated with respect to her own evidence. For influential examples illustrating how one’s evidence interacts with the conditional claims one may assert, see Gibbard (1981).

⁸Ms. Neutra believes what she has been told about the setup and also recalls whether she was shocked 20 times last week (spoiler: she wasn’t). As a result, she is already confident that Evidentia will decide to shock her. Given this, why would she bother trying to persuade Evidentia? That of course depends on further details of the case, but under at least some assumptions Neutra can gain a (small) advantage by trying to persuade – even according to evidential decision theory. Here is a toy model that illustrates this: Suppose that Neutra’s current credence function P' is the result of conditionalizing her prior P on K , the claim that she wasn’t shocked last week. P counts it as certain that Evidentia’s description of the overall setup is correct. Let R be the claim that Neutra argues as she does in the text and S be that Evidentia shocks Neutra with the stun gun. Assume that P rules out that Evidentia would refrain from using the stun gun without Neutra’s attempted persuasion, but does not rule out that Neutra successfully persuades Evidentia to refrain from using the stun gun even though Neutra was not shocked last week. (So $P(\bar{S}\bar{R}) = 0$ but $P(\bar{S}RK) > 0$.) Assume that Neutra’s disutility is proportional to the number of shocks she receives, that the only alternative to trying to persuade is to do nothing, and that Neutra is currently uncertain whether she will try to persuade (so that $P'(R)$ and $P'(\bar{R})$ are both greater than 0). Then the evidential expected value of attempting persuasion ($V(R)$) is (at least a tiny bit) higher than that of doing nothing ($V(\bar{R})$) since under these conditions: $-V(R) = 1 \times P'(S|R) + 0 \times P'(\bar{S}|R) = P(S|RK) = 1 - P(\bar{S}|RK) = 1 - P(\bar{S}RK)/P(RK) < 1$, while $-V(\bar{R}) = 1 \times P'(S|\bar{R}) + 0 \times P'(\bar{S}|\bar{R}) = P(S|\bar{R}K) = 1 - P(\bar{S}|\bar{R}K) = 1 - P(\bar{S}\bar{R}K)/P(\bar{R}K) = 1$.

EVIDENTIA: Feel free.

NEUTRA: You admit that I have relevant evidence that you lack (memory of whether I was shocked last week)? That I prefer that you not shock me? And that I would continue to do so even if you shared with me all of your relevant evidence?

EVIDENTIA: Of course.

NEUTRA: You agree that we share the goal of minimizing the number of shocks I experience?

EVIDENTIA: Yes.

NEUTRA: So if you shared all your evidence with me then you'd know, of a reasonable person with *your* goals and *strictly stronger* evidence, that they prefer that you refrain from shocking. Doesn't that count for something in this case?⁹

EVIDENTIA: No. *Starts to pull the trigger.*¹⁰

NEUTRA: Wait! This is a long-shot, but let me give you *my* evidence. Then you might not think shocking me will help me. Last week what happened was—

EVIDENTIA: *Puts fingers in own ears.* LA LA LA LA! I'm not listening to you! LA LA LA—

NEUTRA: What are you doing?

EVIDENTIA: *Removes fingers from ears.* Sorry, I had to stop you from telling me about what happened last week. If you had told me (regardless of whether you told me that you were shocked or that you weren't shocked), I would have come to think that shocking you now would serve no purpose, and indeed would be cruel and inhumane.¹¹ So I

⁹Compare Gallow (2019, 4). To sidestep potential complications arising from epistemic permissivism (Schoenfield 2014, White 2005) (according to which there is more than one rationally permissible response to a given batch of total evidence), we may assume that Neutra and Evidentia share a prior probability function, and realize that they do.

¹⁰In response to the seemingly irrational nature of Evidentia's choice here, evidential decision theorists have various lines of reply available, many familiar from the debate about Newcomb's problem. These include claiming: (1) Evidentia would be able to sense in herself whatever factors the DPU used to predict her choice, and once those factors are taken into account, evidential decision theory recommends refraining from shocking. (This is the so-called "tickle defence" (Eells 1981, §4).) (2) The presentation invites readers to see things from Neutra's point of view, and Neutra recalls whether she was shocked last week. Everyone agrees that if Evidentia had *that* information, she should refrain from shocking. But because Neutra's point of view is so salient, it is difficult to keep firmly in mind that Evidentia lacks that information when evaluating her choice.

¹¹This is because: Learning that there were no shocks last week would make Evidentia regard her current choice as settling whether Neutra gets 0 or 1 shocks total. Similarly, learning that there were 20 shocks last week would make Evidentia regard her current choice as settling whether Neutra gets 20 or 21 shocks total. So either way, Evidentia would regard choosing to shock now as pointless, just as evidential decision theorists favor taking two boxes in the variant of Newcomb's problem in which both boxes are transparent (Gibbard and Harper 1978, 153-4).

wouldn't have shocked you. So the DPU would have last week predicted that I wouldn't shock you, and would have ensured that you were last week shocked 20 times. So for your sake, I can't let you tell me what happened last week.¹²

NEUTRA: *Stunned silence*

An administrator bursts into the room.

ADMINISTRATOR: There has been a mistake! Last week the DPU did not scan Professor Evidentia. It accidentally scanned Ms. Neutra instead! If it predicted that Ms. Neutra would today decide to give a shock to Professor Evidentia, it left everyone alone last week. But if it predicted that Ms. Neutra would refrain today, then it ensured that last week Professor Evidentia received twenty shocks.

NEUTRA: So I must now decide whether to shock her? *Points at Evidentia.*

Administrator plucks the stun gun out of Professor Evidentia's hands.

ADMINISTRATOR: Yes. *Hands the stun gun to Ms. Neutra and exits.*

NEUTRA: I don't want to shock anyone.

EVIDENTIA: But if you are to reduce the number of shocks you expect me to receive, you must.

NEUTRA: This is difficult for me, but I see where you are coming from. If I shock you now, you'll almost certainly receive a total of just one shock. But if I refrain, you'll almost certainly receive a total of 20. So if I'm rational and kind, I should shock you. And I shouldn't let you tell me whether you have been shocked last week.

EVIDENTIA: Now you've got the hang of it!

NEUTRA: OK, if you want me shock you— *Points the stun gun at Evidentia.*

EVIDENTIA: Wait! *(Quietly)* I don't want you to shock me.

NEUTRA: *Lowers stun gun.* What? Didn't we just agree that the rational, kind thing for me to do is to shock you, for your own good?

EVIDENTIA: Yes, yes, that is all true. But... I prefer that you don't shock me!

NEUTRA: Why?

¹²Compare Wells (2019, §5) and Gallow (2019, 4), which also make the point (about an analogous situation) that evidential decision theory would enjoin Evidentia to inflict a small additional punishment on Neutra, if that is what it took to keep Neutra from revealing whether she was shocked the previous week.

EVIDENTIA: Shocking me (slightly) increases the expected number of shocks I suffer.

NEUTRA: Auugggh! We already agreed I should think: shocking you *reduces* the expected number of shocks you suffer.

EVIDENTIA: Well, from *your* point of view I can see why you are saying that. But you must understand that as an evidential altruist, I must prefer and hope in the light of *my* evidence. I hope that you act irrationally and/or unkindly and refrain from shocking me!¹³

NEUTRA: *Places gun on the table.* Thanks very much for the tour.

Scene 3

Several days later. Neutra, Causey, and Evidentia sip coffees at a restaurant.

NEUTRA: I don't mean to be rude, but you are both totally irrational! Professor Causey, *please* stop sticking your fingers in LAN outlets. And Professor Evidentia, *please* stop shocking people for the sake of reducing shocks they may or may not have received by the DPU a full week prior.

CAUSEY: Wait, you think that *both of us* are irrational?

NEUTRA: Bonkers, in fact.

CAUSEY: How odd.

EVIDENTIA: Remarkable!

NEUTRA: But surely almost everyone in your profession thinks that you two act irrationally. That is just common sense.

CAUSEY: Actually, a substantial fraction of decision theorists hold views that say I have acted rationally—¹⁴

EVIDENTIA: —And a substantial fraction of decision theorists hold views that say *I* have acted rationally.¹⁵

¹³For reasons similar to those described in footnote 8, Evidentia is already confident of what Neutra decides, but might count trying to persuade Neutra as doing a tiny bit of good. Note that if side-betting is an option, Evidentia will by her own lights do better to spend her breath arranging side-bets with Neutra that take advantage of Evidentia's extra evidence about whether Neutra will decide to administer a shock.

¹⁴See for example Gibbard and Harper (1978), Joyce (1999), Lewis (1981a), Skyrms (1980), Stalnaker (1981), Wedgwood (2013).

¹⁵See for example Ahmed (2014), Horgan (2015), Jeffrey (1965). For a small-sample ongoing survey on how decision theorists divide between one-boxing and two-boxing in Newcomb's problem, see https://philpapers.org/surveys/results.pl?affil=Target+faculty&areas0=1399&areas_max=1&grain=coarse, a continuation of the survey described in Bourget and Chalmers (2014).

CAUSEY: And almost *no* decision theorists hold views that say *neither* of us has acted rationally. Indeed, aside from inessential window dressing, the two choice situations are the same: If you think that one should avoid sticking one's fingers in LAN outlets, you should also think that you should shock in the DPU situation—

EVIDENTIA: —And if you think that one should refrain from shocking in the DPU situation, you should also think that one should stick one's fingers in LAN outlets. So we ask: what decision rule do you endorse that says that both of us have acted irrationally?¹⁶

*Curtain.*¹⁷

References

Arif Ahmed. *Evidence, Decision and Causality*. Cambridge University Press, 2014.

Arif Ahmed. The “why ain’cha rich?” argument. In Arif Ahmed, editor, *Newcomb’s Problem*, pages 55–72. Cambridge ; New York, NY: Cambridge University Press, 2018.

Adam Bales. Re: A bit of fun re Newcomb’s problem, 2019. Personal communication.

Adam Thomas Bales. *Decision and Dependence: A Defence of Causal Decision Theory*. PhD thesis, University of Cambridge, 2017.

Maya Bar-Hillel and Avishai Margalit. Newcomb’s paradox revisited. *The British Journal for the Philosophy of Science*, 23(4):295–304, 1972.

¹⁶ Author’s note: I agree with Causey and Evidentia that aside from inessential window dressing the two choice situations are the same, and agree with Neutra that there is some initial intuitive pull toward thinking that Causey and Evidentia both act irrationally. I disagree with Neutra that on reflection we should count both Causey and Evidentia as irrational, or even seek a decision theory that does so. Instead I conjecture that our intuitions are sensitive to the window-dressing, and are thereby pulled astray in at least one of the cases. For example, it is an implicit assumption of the play that Causey and Evidentia’s descriptions of the setup are to be trusted. But in imagining the case we may not be respecting that assumption. We may naturally suspect on Neutra’s behalf that the LANs might be regular outlets, and that the DPU might be a regular smoke detector. If Causey and Evidentia take such suspicions seriously enough then according to both causal and evidential decision theory, each acts irrationally (Bales 2019; see also Bales 2017, Bar-Hillel and Margalit 1972, Ninan 2006).

In my view the moral of the play is not that both causal and evidential decision theories are incorrect. Rather it is that to fully defend either theory’s treatment of Newcomb’s problem it is not enough to just put forward an attractive positive analysis. One should also defend an error theory explaining the source of tempting intuitions that conflict with that analysis.

¹⁷ For helpful discussion and feedback thanks to Arif Ahmed (especially with footnotes 8 and 13), Adam Bales, Alan Hájek, Adam Kolber (especially with footnote 8), Alex Meehan, Timothy L. Williamson, participants in a graduate seminar (jointly taught with Boris Kment) at Princeton in Fall 2018 and an undergraduate seminar (jointly taught with Alex Meehan) in Spring 2019, and two anonymous referees. Original cast: Causey: Dylan Mavrides; Neutra: Rohit Dilip and Sonia Murthy, Evidentia: Davey Fitzpatrick. It should be noted that none of the characters are modeled on actual individuals.

- David Bourget and David J. Chalmers. What do philosophers believe? *Philosophical Studies*, 170(3):465–500, 2014.
- Ellery Eells. Causality, utility, and decision. *Synthese*, 48(2):295–329, 1981. ISSN 1573-0964. doi: 10.1007/BF01063891.
- J. Dmitri Gallow. Face death with indifference. Manuscript, 2019.
- Allan Gibbard. Two recent theories of conditionals. In William L. Harper, Robert Stalnaker, and Glenn Pearce, editors, *IFS: Conditionals, Belief, Decision, Chance and Time*, The University of Western Ontario Series in Philosophy of Science, pages 211–247. Springer Netherlands, Dordrecht, 1981. doi: 10.1007/978-94-009-9117-0_7.
- Allan Gibbard and William Harper. Counterfactuals and two kinds of expected utility. In A. Hooker, J. J. Leach, and E. F. McClennen, editors, *Foundations and Applications of Decision Theory*, pages 125–162. D. Reidel, 1978.
- Terence Horgan. Counterfactuals and Newcomb’s problem. *Journal of Philosophy*, 78(6): 331–356, 1981.
- Terry Horgan. Newcomb’s problem revisited. *The Harvard Review of Philosophy*, 22:4–15, 2015.
- Richard Jeffrey. *The Logic of Decision*. University of Chicago Press, 1965.
- James M. Joyce. *The Foundations of Causal Decision Theory*. Cambridge University Press, 1999.
- David Lewis. Causal decision theory. *Australasian Journal of Philosophy*, 59(1):5–30, 1981a.
- David Lewis. ‘Why ain’cha rich?’. *Noûs*, 15(3):377–380, 1981b. ISSN 0029-4624. doi: 10.2307/2215439.
- Benjamin Libet. *Mind Time: The Temporal Factor in Consciousness*. Harvard University Press, Cambridge, Mass., 2005. ISBN 978-0-674-01846-4.
- Dilip Ninan. Illusions of influence in Newcomb’s problem. Manuscript, 2006.
- Robert Nozick. Newcomb’s problem and two principles of choice. In Nicholas Rescher, editor, *Essays in Honor of Carl G. Hempel*, pages 114–146. Reidel, 1969.
- Derek Parfit. *Reasons and Persons*. Oxford University Press, 1984.
- Arthur Prior. Thank goodness that’s over. *Philosophy*, 34(128):12–17, 1959.
- Miriam Schoenfield. Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 48(2):193–218, 2014. ISSN 00294624. doi: 10.1111/nous.12006.

- Brian Skyrms. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. Yale University Press, New Haven, 1980. ISBN 978-0-300-02339-8.
- Robert C. Stalnaker. Letter to David Lewis. In William L. Harper, Robert Stalnaker, and Glenn Pearce, editors, *IFS: Conditionals, Belief, Decision, Chance and Time*, The University of Western Ontario Series in Philosophy of Science, pages 151–152. Springer Netherlands, Dordrecht, 1981. doi: 10.1007/978-94-009-9117-0_7.
- Ralph Wedgwood. Gandalf’s solution to the newcomb problem. *Synthese*, 190(14): 2643–2675, 2013. doi: 10.1007/s11229-011-9900-1.
- Ian Wells. Equal opportunity and Newcomb’s problem. *Mind*, 128(510):429–457, 2019. doi: 10.1093/mind/fzx018.
- Roger White. Epistemic permissiveness. *Philosophical Perspectives*, 19(1):445–459, 2005. ISSN 1520-8583. doi: 10.1111/j.1520-8583.2005.00069.x.