Title

A puzzle about conditionals and conditionalizationfrom the perspective of dynamic epistemic logic

Author

Holliday, Wesley Halcrow

Publication Date

2023-03-02

A puzzle about conditionals and conditionalization from the perspective of dynamic epistemic logic

Wesley H. Holliday

March 2, 2023

Abstract

Philosophers of language and formal epistemologists have discussed apparent violations of the rule of Bayesian conditionalization involving indicative conditionals. We attempt to clarify this issue using an analogy from dynamic epistemic logic.

Introduction

Recently there has been much discussion in the philosophy of language and formal epistemology of puzzles like the following, drawn from Ciardelli and Ommundsen Forthcoming (also see, e.g., Goldstein and Santorio 2021; Khoo 2022; Fusco Forthcoming; McNamara and Zhang 2023). A fair six-sided die was rolled. After learning that the die landed on an even number (Even), intuitively you should assign probability 1 to the following indicative conditional: *if* the die landed on a low number, i.e., 1, 2, or 3, *then* it landed on 2 ($\mathsf{Low} > \mathsf{2}$). Thus, we should have

$$Pr_{\mathsf{Even}}(\mathsf{Low} > \mathsf{2}) = 1,$$

where Pr_{Even} is your probability function on sentences after learning Even.[1] On the other hand, where Pr is your prior probability function on sentences, we should also have

$$\frac{Pr((\mathsf{Low} > \mathsf{2}) \wedge \mathsf{Even})}{Pr(\mathsf{Even})} \leq \frac{Pr(\mathsf{Low} > \mathsf{2})}{Pr(\mathsf{Even})} < 1,$$

assuming that $Pr(\mathsf{Low} > \mathsf{2}) = \frac{1}{3}$ in line with Stalnaker's Thesis (Stalnaker 1970) for simple indicative conditionals. Thus, we have an apparent violation of the rule of Bayesian conditionalization, as usually formulated using the ratio definition of conditional probability.[2] A similar puzzle arises for the epistemic modal 'must' ($\Box$) if it is possible to have $Pr(\mathsf{Even}) = .5$, $Pr(\Box\mathsf{Even}) < .5$, and then $Pr_{\mathsf{Even}}(\Box\mathsf{Even}) = 1$.

To help resolve this puzzle, we will consider an analogy from *dynamic epistemic logic* (for introductions to the field, see Pacuit 2013 and Baltag and Renne 2016). To keep things simple, let us restrict attention to cases of learning *non-conditional* and *non-modal* sentence like Even.

[1] Below we will discuss probabilities of *propositions* instead of sentences.

[2] Of course, one response is to say that it does not make sense to assign probabilities to indicative conditional sentences, but for the purposes of this note, let us assume the position of those who think that it does make sense.

Dynamic epistemic logic

We need only consider the most basic dynamic epistemic system (known as Public Announcement Logic). Consider a formal language $\mathcal{L}$ generated by the grammar

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid [\varphi]\varphi$$

where p belongs to a set Prop of propositional variables. The material conditional $\to$ and material biconditional $\leftrightarrow$ are defined using $\neg$ and $\wedge$ as usual. We interpret $\Box$ as epistemic necessity and read $[\varphi]\psi$ as "after information update with φ, ψ."

We define when a formula is true relative to a pair $\mathcal{M} = \langle W, V \rangle$ of a nonempty set W and a valuation $V : \mathsf{Prop} \to \wp(W)$, a world $w \in W$, and a set $i \subseteq W$ of epistemic possibilities[3] with $w \in i$, as follows:

1. $\mathcal{M}, w, i \vDash p$ iff $w \in V(p)$;
2. $\mathcal{M}, w, i \vDash \neg\varphi$ iff $\mathcal{M}, w, i \nvDash \varphi$;
3. $\mathcal{M}, w, i \vDash \varphi \wedge \psi$ iff $\mathcal{M}, w, i \vDash \varphi$ and $\mathcal{M}, w, i \vDash \psi$;
4. $\mathcal{M}, w, i \vDash \Box\varphi$ iff for all $v \in i$, $\mathcal{M}, v, i \vDash \varphi$;
5. $\mathcal{M}, w, i \vDash [\varphi]\psi$ iff $\mathcal{M}, w, i \vDash \varphi$ implies $\mathcal{M}, w, i \cap [\![\varphi]\!]^{\mathcal{M},i} \vDash \psi$,

where $[\![\varphi]\!]^{\mathcal{M},i} = \{u \in W \mid \mathcal{M}, u, i \vDash \varphi\}$. Say a formula is *valid* if it is true relative to all such triples.

A basic exercise

One of the first facts that students of dynamic epistemic logic learn is that there is a valid law that allows us to calculate what is known *after* information update in terms of what is true in the *prior* state:

$$[\varphi]\Box\psi \leftrightarrow \underline{\hspace{3cm}}.$$

Students are asked to fill in the blank. The instructor expects that the first answer will be

Naive Conditionalization: $[\varphi]\Box\psi \leftrightarrow (\varphi \to \Box(\varphi \to \psi))$.

However, this is not a valid principle for all ψ. The reason is that our language contains the modal operator $\Box$ that is sensitive to the set i that is modified by information update. For example, take φ to be p and ψ to be $\Box p$. Then the relevant instance of Naive Conditionalization is

$$[p]\Box\Box p \leftrightarrow (p \to \Box(p \to \Box p).$$

But suppose that $W = \{w, v\}$, $V(p) = \{w\}$, and $i = W$. Then it is easy to see that at $\mathcal{M}, w, i$, the left-hand side is true while the right-hand side is false, because $\mathcal{M}, w, i \vDash p$ but $\mathcal{M}, w, i \nvDash \Box p$.

After appreciating the mistake of Naive Conditionalization, students learn the answer of

Correct Conditionalization: $[\varphi]\Box\psi \leftrightarrow (\varphi \to \Box(\varphi \to [\varphi]\psi))$.

[3] A more standard presentation in dynamic epistemic logic uses an accessibility relation R, so $R(w) = \{v \in W \mid wRv\}$ appears in place of i in the truth definition, and information update changes R; but the presentation with i is just the special case where $R(w) = R(v)$ for all $w, v \in W$, and this presentation is more familiar in the philosophy of language (Yalcin 2007).

It is not hard to see that this principle is valid for all ψ. In the example above, we have the instance

$$[p]\Box\Box p \leftrightarrow (p \to \Box(p \to [p]\Box p),$$

and while above we had $\mathcal{M}, w, i \nvDash \Box p$, now we have $\mathcal{M}, w, i \vDash [p]\Box p$, so $\mathcal{M}, w, i \vDash \Box(p \to [p]\Box p)$.

The same point again with probability

To make essentially the same point with probability, we add to our model $\mathcal{M}$ a probability measure μ on *propositions*, i.e., $\mu : \wp(W) \to [0,1]$. This can be lifted to a probability function $\mu^{\mathcal{M},i} : \mathcal{L} \to [0,1]$ on *sentences* by $\mu^{\mathcal{M},i}(\psi) = \mu(\llbracket\psi\rrbracket^{\mathcal{M},i})$. Now suppose that when learning a proposition $E \subseteq W$ with $\mu(E) \neq 0$, we update μ by standard Bayesian conditionalization: $\mu_E(A) = \mu(A \cap E)/\mu(E)$. Then in the spirit of Naive Conditionalization, one might think that one's probability function for sentences should update similarly:

$$(\mu^{\mathcal{M},i})_\varphi(\psi) = \frac{\mu^{\mathcal{M},i}(\psi \wedge \varphi)}{\mu^{\mathcal{M},i}(\varphi)}.$$

But it is easy to see this is a mistake. Add a measure μ to the example above so that $\mu(\{w\}) = \mu(\{v\}) = .5$. Then initially the probability of $\Box p$ is 0, whereas after updating with p, the probability of $\Box p$ should be 1, so the ratio equation above for the probabilities of sentences after update is incorrect. The correct equation with its derivation is

$$\begin{aligned}
(\mu^{\mathcal{M},i})_\varphi(\psi) &= (\mu_{\llbracket\varphi\rrbracket^{\mathcal{M},i}})^{\mathcal{M},i\cap\llbracket\varphi\rrbracket^{\mathcal{M},i}}(\psi) \\
&= \mu_{\llbracket\varphi\rrbracket^{\mathcal{M},i}}(\llbracket\psi\rrbracket^{\mathcal{M},i\cap\llbracket\varphi\rrbracket^{\mathcal{M},i}}) \\
&= \frac{\mu(\llbracket\psi\rrbracket^{\mathcal{M},i\cap\llbracket\varphi\rrbracket^{\mathcal{M},i}} \cap \llbracket\varphi\rrbracket^{\mathcal{M},i})}{\mu(\llbracket\varphi\rrbracket^{\mathcal{M},i})} \\
&= \frac{\mu(\llbracket[\varphi]\psi \wedge \varphi\rrbracket^{\mathcal{M},i})}{\mu(\llbracket\varphi\rrbracket^{\mathcal{M},i})} \\
&= \frac{\mu^{\mathcal{M},i}([\varphi]\psi \wedge \varphi)}{\mu^{\mathcal{M},i}(\varphi)}. \qquad (1)
\end{aligned}$$

As in the previous section, the correct equation replaces ψ in the naive equation with $[\varphi]\psi$.

A non-contextualist presentation

Above we explained dynamic epistemic logic in a "contextualist" style, where the set of worlds at which $\Box p$ is true before update with p (namely, $\varnothing$) is different than the set of worlds at which $\Box p$ is true after update with p (namely, W). But there is a mathematically equivalent non-contextualist presentation, developed in Holliday et al. 2012 and Wang and Cao 2013. In this approach, information update is modeled by transitions within a model rather than by updating an information state i. A model is now a tuple $\mathcal{N} = \langle W, f, R, V\rangle$ where W is a nonempty set, f is a partial function such that for all $w \in W$ and $A \subseteq W$, if f is defined at $\langle w, A\rangle$, then $f(w, A) \in W$,[4] R is an equivalence relation on W, and $V : \mathsf{Prop} \to \wp(W)$ as before. Of course, f is reminiscent of Stalnaker's (1968) selection function semantics for conditionals. Given $w \in W$,

[4]In fact Holliday et al. 2012 and Wang and Cao 2013 use a formula φ in place of the proposition A, but this is immaterial; we can set $f(x, A)$ to be undefined in their models if A is not definable by a formula.

$R(w) = \{v \in W \mid wRv\}$ is the set of epistemically accessible worlds from w. Now the key fact is that one can impose constraints on such models (see Definition 33 of Wang and Cao 2013) so that we get *precisely the same logic for* $\mathcal{L}$ as before but with the following alternative semantic clauses for $\Box$ and $[\varphi]$:

- $\mathcal{N}, w \vDash \Box\varphi$ iff for all $v \in R(w)$, $\mathcal{N}, v \vDash \varphi$;
- $\mathcal{N}, w \vDash [\varphi]\psi$ iff $\mathcal{N}, w \vDash \varphi$ implies $\mathcal{N}, f(w, [\![\varphi]\!]^{\mathcal{N}}) \vDash \psi$.

Since we get precisely the same logic, the correct answer to the fill-in-the-blank question is still Correct Conditionalization rather than Naive Conditionalization. But this approach is non-contextualist in the sense that there is just one set of worlds where $\Box p$ is true, namely $[\![\Box p]\!]^{\mathcal{N}} = \{v \in W \mid \mathcal{N}, v \vDash \Box p\}$, which does not change pre- and post-update by p.[5] As a consequence, when we bring probability into the picture, we can no longer use Bayesian conditionalization on propositions: for if there is just one proposition expressed by $\Box p$, and we initially assign that proposition probability 0, then we can never move to assigning it probability 1 by Bayesian conditionalization. We again lift $\mu : \wp(W) \to [0,1]$ to $\mu^{\mathcal{N}} : \mathcal{L} \to [0,1]$ by $\mu^{\mathcal{N}}(\varphi) = \mu([\![\varphi]\!]^{\mathcal{N}})$. But now suppose that when learning E, we update our probability measure by what we will call *f-conditioning*:[6]

$$\mu_E(A) = \frac{\mu(\{w \in E \mid f(w, E) \in A\})}{\mu(E)}.$$

Then we obtain precisely the same equation for updating the probabilities of sentences as before:

$$\begin{aligned} (\mu^{\mathcal{N}})_\varphi(\psi) &= (\mu_{[\![\varphi]\!]^{\mathcal{N}}})^{\mathcal{N}}(\psi) \\ &= \mu_{[\![\varphi]\!]^{\mathcal{N}}}([\![\psi]\!]^{\mathcal{N}}) \\ &= \frac{\mu(\{w \in [\![\varphi]\!]^{\mathcal{N}} \mid f(w, [\![\varphi]\!]^{\mathcal{N}}) \in [\![\psi]\!]^{\mathcal{N}}\})}{\mu([\![\varphi]\!]^{\mathcal{N}})} \\ &= \frac{\mu([\![\varphi \wedge [\varphi]\psi]\!]^{\mathcal{N}})}{\mu([\![\varphi]\!]^{\mathcal{N}})} \\ &= \frac{\mu^{\mathcal{N}}([\varphi]\psi \wedge \varphi)}{\mu^{\mathcal{N}}(\varphi)}. \end{aligned} \tag{2}$$

Thus, regardless of whether we give a contextualist treatment of our language with Bayesian conditionalization to update probabilities of propositions or a non-contextualist treatment of our language with f-conditioning to update probabilities of propositions, we get the same result for updating probabilities of sentences, which is not the naive result.

Completing the analogy

Let us return to the initial puzzle. The point of our analogy with dynamic epistemic logic is that thinking that the probability of an indicative conditional after updating with (non-conditional and non-modal) φ should be given by

[5] Yet we can still give this approach a contextualist spin. Given an equivalence class C of the relation R, define $[\![\Box p]\!]^{\mathcal{N},C} = \{w \in C \mid \mathcal{N}, w \vDash \Box p\}$. Then we can think of $\Box p$ as expressing $[\![\Box p]\!]^{\mathcal{N},R(w)}$ pre-update and $[\![\Box p]\!]^{\mathcal{N},R(f(w,[\![p]\!]^{\mathcal{N}}))}$ post-update.

[6] This differs significantly from Lewis's (1976) *imaging*. For us, $w \in E$ does not imply $f(w, E) = w$, as it would with Stalnaker's (1968) selection functions used for Lewisian imaging. In particular, we may have $w \in [\![p]\!]^{\mathcal{N}}$ but $w \notin [\![\Box p]\!]^{\mathcal{N}}$, but we always have $f(w, [\![p]\!]^{\mathcal{N}}) \in [\![\Box p]\!]^{\mathcal{N}}$ in the framework of Wang and Cao 2013, § 4, in which case $w \neq f(w, [\![p]\!]^{\mathcal{N}})$.

Naive Bayesian Conditionalization:

$$Pr_\varphi(\psi > \chi) = \frac{Pr((\psi > \chi) \wedge \varphi)}{Pr(\varphi)}$$

is committing a mistake analogous to that of the student of dynamic epistemic logic whose first guess above is Naive Conditionalization. Like $\Box$, the indicative conditional $>$ is sensitive to information update, so *of course* Naive Bayesian Conditionalization could not be correct.

The analogy also points the way to a solution to the initial puzzle. In the dynamic epistemic case, to work out what is known after information update in terms of what is the case beforehand, we need to think not about the pre-update truth of ψ but rather about the pre-update truth of $[\varphi]\psi$. Similarly, to work out the probability of an indicative conditional after information update in terms of pre-update probabilities, we need to think not about the pre-update probability of $\psi > \chi$ but rather about the pre-update probability of $\varphi > (\psi > \chi)$. In particular, we should have

Correct Bayesian Conditionalization:

$$Pr_\varphi(\psi > \chi) = \frac{Pr((\varphi > (\psi > \chi)) \wedge \varphi)}{Pr(\varphi)},$$

as an instance of the following more general law (again for non-conditional and non-modal φ) that McNamara and Zhang (2023) call Conditional Conditionalization, which is perfectly analogous to (1) and (2):

$$Pr_\varphi(\chi) = \frac{Pr((\varphi > \chi) \wedge \varphi)}{Pr(\varphi)}. \tag{3}$$

Applying this to the initial puzzle, intuitively we should have

$$Pr(\mathsf{Even} > (\mathsf{Low} > 2)) = 1, \text{ and}$$

$$Pr((\mathsf{Even} > (\mathsf{Low} > 2)) \wedge \mathsf{Even}) = .5,$$

so

$$Pr_{\mathsf{Even}}(\mathsf{Low} > 2) = \frac{Pr((\mathsf{Even} > (\mathsf{Low} > 2)) \wedge \mathsf{Even})}{Pr(\mathsf{Even})} = \frac{.5}{.5} = 1,$$

as desired. Of course, if $(\varphi > (\psi > \chi)) \wedge \varphi$ were equivalent to $(\psi > \chi) \wedge \varphi$, then there would be no difference between Correct and Naive Bayesian Conditionalization. But the case of the die is precisely the kind of case that motivates the denial of that equivalence (McGee 1985, Santorio 2022, § 4.2): while $Pr(\mathsf{Even} > (\mathsf{Low} > 2)) \wedge \mathsf{Even}) = .5$, we have $Pr(\mathsf{Low} > 2) = \frac{1}{3}$, so $Pr((\mathsf{Low} > 2) \wedge \mathsf{Even}) \leq \frac{1}{3}$. On the other hand, most authors agree that when φ and χ are non-conditional and non-modal, then $(\varphi > \chi) \wedge \varphi$ is equivalent to $\chi \wedge \varphi$, in which case standard Bayesian conditionalization drops out of Conditional Conditionalization as a special case for non-conditional and non-modal sentences.[7]

If a conditional is probabilistically independent of its (again non-conditional and non-modal) antecedent,

[7]But when Conditional Conditionalization deviates from standard Bayesian conditionalization, is an agent who follows Conditional Conditionalization susceptible to Dutch books? McNamara and Zhang (2023) prove that such an agent is in fact immune to Dutch books provided we understand information update in a particular way. They also observe that Conditional Conditionalization leads to violations of *conglomerability*, the requirement that $\min\{Pr_\varphi(\chi), Pr_{\neg\varphi}(\chi)\} \leq Pr(\chi) \leq \max\{Pr_\varphi(\chi), Pr_{\neg\varphi}(\chi)\}$. But from the point of view of dynamic epistemic logic developed above, conglomerability cannot be right for probabilities of *sentences*. A simple counterexample takes φ to be p and χ to be $\Box p \vee \Box\neg p$. Then $\min\{Pr_p(\Box p \vee \Box\neg p), Pr_{\neg p}(\Box p \vee \Box\neg p)\} = \min\{1, 1\} = 1$ even if $Pr(\Box p \vee \Box\neg p) = 0$.

in the sense that $Pr(\varphi > \chi) = Pr(\varphi > \chi \mid \varphi) := Pr((\varphi > \chi) \wedge \varphi)/Pr(\varphi)$, then Conditional Conditionalization is equivalent to what Goldstein and Santorio (2021) call the Update Thesis from McGee 1989, p. 504: $Pr_\varphi(\chi) = Pr(\varphi > \chi)$. However, there are apparent counterexamples to such probabilistic independence and to the Update Thesis (McGee 2000; Kaufman 2004). Here is a simple version of such a counterexample:[8] Urn A contains 10 marbles, 8 of which are red and all of which are spotted; Urn B contains 10 marbles, 2 of which are red and none of which are spotted. You choose an urn by flipping a fair coin, and then you randomly pick a marble from the chosen urn. Thus, $Pr_{\mathsf{red}}(\mathsf{spotted}) = .8$, but one can reasonably have $Pr(\mathsf{red} > \mathsf{spotted}) = .5$, since $\mathsf{red} > \mathsf{spotted}$ is true if and only if Urn A is chosen. If this is correct, then it is a counterexample to the Update Thesis. But it is consistent with Conditional Conditionalization. Since plausibly $(\mathsf{red} > \mathsf{spotted}) \wedge \mathsf{red}$ is equivalent to $\mathsf{spotted} \wedge \mathsf{red}$, we have

$$Pr_{\mathsf{red}}(\mathsf{spotted}) = \frac{Pr((\mathsf{red} > \mathsf{spotted}) \wedge \mathsf{red})}{Pr(\mathsf{red})} = \frac{Pr(\mathsf{spotted} \wedge \mathsf{red})}{Pr(\mathsf{red})} = \frac{.4}{.5} = .8,$$

so there is no contradiction with Conditional Conditionalization.

Though our analogy with dynamic epistemic logic falls far short of a *proof* that Conditional Conditionalization is correct, hopefully it at least suggests that philosophers of language and formal epistemologists may benefit from considering the lessons learned about information update in the context of dynamic epistemic logic.[9]

Acknowledgements

Thanks to Mathias Boehm, Matt Mandelkern, Calum MacNamara, Eric Pacuit, and Snow Zhang for helpful discussion.

References

Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.

Ivano Ciardelli and Adrian Ommundsen. Probabilities of conditionals: updating Adams. *Nous*, Forthcoming. URL https://doi.org/10.1111/nous.12437.

Melissa Fusco. Dutch-booking indicative conditionals. *Philosophy and Phenomenological Research*, Forthcoming. URL https://doi.org/10.1111/phpr.12910.

Simon Goldstein and Paolo Santorio. Probability for epistemic modalities. *Philosophers' Imprint*, 2021(33): 1–37, 2021. URL http://hdl.handle.net/2027/spo.3521354.0021.033.

Wesley H. Holliday and Thomas F. Icard. Indicative conditionals and dynamic epistemic logic. In J. Lang, editor, *Proceedings of Theoretical Aspects of Rationality and Knowledge 2017*. Electronic Proceedings in Theoretical Computer Science, 2017. URL https://arxiv.org/abs/1707.08752.

Wesley H. Holliday, Tomohiro Hoshi, and Thomas F. Icard. A uniform logic of information dynamics. In Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence Moss, editors, *Advances in Modal Logic*, volume 9, pages 348–367. College Publications, London, 2012.

[8] Due to Matt Mandelkern (personal communication).

[9] For another call for interaction between dynamic epistemic logic and philosophy of language, see Holliday and Icard 2017.

Stefan Kaufman. Conditioning against the grain. *Journal of Philosophical Logic*, 33:583–606, 2004. URL https://doi.org/10.1023/B:LOGI.0000046142.51136.bf.

Justin Khoo. *The meaning of "if"*. Oxford University Press, 2022.

David Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3): 297–315, 1976.

Vann McGee. A counterexample to modus ponens. *The Journal of Philosophy*, 82(9):462–471, 1985. doi: 10.2307/2026276.

Vann McGee. Conditional probabilities and compounds of conditionals. *The Philosophical Review*, 98(4): 485–541, 1989. URL https://doi.org/10.2307/2185116.

Vann McGee. To tell the truth about conditionals. *Analysis*, 60(1):107–111, 2000. URL https://doi.org/10.1111/1467-8284.00210.

Calum McNamara and Snow Zhang. Why (not) conditionalize? 2023. Manuscript.

Eric Pacuit. Dynamic epistemic logic II: Logics of information change. *Philosophy Compass*, 8(9):815–833, 2013. URL https://doi.org/10.1111/phc3.12060.

Paolo Santorio. Trivializing informational consequence. *Philosophy and Phenomenological Research*, 104(2): 297–320, 2022. doi: 10.1111/phpr.12745.

Robert C. Stalnaker. A theory of conditionals. In Nicholas Rescher, editor, *Studies in Logical Theory*, pages 98–112. Blackwell, 1968.

Robert C. Stalnaker. Probability and conditionals. *Philosophy of Science*, 37(1):64–80, 1970.

Yanjing Wang and Qinxiang Cao. On axiomatizations of public announcement logic. *Synthese*, 190(1): 103–134, 2013. URL https://doi.org/10.1007/s11229-012-0233-5.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007. URL https://doi.org/10.1093/mind/fzm983.