

# On Algebra Relativisation

Chloé de Canson

**Abstract.** Katie Steele and H. Orri Stefánsson argue that, to reflect an agent’s limited awareness, the algebra of propositions on which that agent’s credences are defined should be relativised to their awareness state. I argue that this produces insurmountable difficulties. But the project of relativising the agent’s algebra to reflect their partial perspective need not be abandoned: the algebra can be relativised, not to the agent’s awareness state, but to what we might call their subjective modality.

An agent’s awareness is always what Katie Steele and H. Orri Stefánsson call “limited”: the agent does not have an exhaustive grasp of all the possible ways the world might be (2021a: 1223). In order to accommodate this facet of the agent’s epistemic state, Steele and Stefánsson propose to *relativise* the algebra of propositions on which the agent’s credences are defined, such that this algebra no longer represents a comprehensive landscape of possibilities, but instead, corresponds to the agent’s personal, partial perspective—what they call the agent’s “awareness context” (1214). They suggest building this algebra from the bottom up; thus they start from a set of “basic propositions of which [the agent] is aware” (1223), and from it, “generate a Boolean algebra ... in the usual way” (1224). This allows them to represent awareness growth by “expansion” (1212): as agents become aware of propositions beyond the ones they were entertaining, the set of basic propositions expands, and so does the resulting algebra. They illustrate their proposal with the following example:

Suppose you are deciding whether to see a movie at your local cinema. ... You know that the movie’s ... genre will affect your viewing experience. ... The genres you consider are thrillers and comedies. But then ... you realize that the genre could be drama. ... In [this] shift ... you have ... *expanded* or *extended* the set of possibilities you entertain to include the further genre of drama. (1212, emphasis original)

The idea of relativising algebras to the agent’s awareness has a history. For instance, Hacking (1967), Shimony (1970), Wenmackers and Romeijn (2016), and Bradley (2017) all mention a

similar kind of relativisation.<sup>1</sup> What Steele and Stefánsson do is develop the general idea into a precise proposal. In what follows, I argue that the proposal produces insurmountable difficulties: awareness-relativisation should be abandoned. Another form of relativisation however, to what we might call the agent's *subjective modality*, is possible.

Assessing the awareness-relativisation of algebras requires a clear sense of what, exactly, the algebra is being relativised to, that is, of what is meant by *awareness*. Steele and Stefánsson do not discuss this in detail in their paper, but their proposal is consistent with two accounts. On the first account, the agent is aware of a proposition if they are actively thinking about it at the time of the decision: this is what we might call the *occurrence* account of awareness. Richard Bradley (2017), whom Steele and Stefánsson cite favourably, adopts this account: on his view, an agent counts as unaware of a proposition if, for instance, they have temporarily forgotten it, or if they are deliberately excluding it from their decision-making considerations. In their book, Steele and Stefánsson also allow for this account; they write that unawareness includes situations where agents “fail to account for relevant possibilities in [their] decision-making due to a momentary lack of perspective—a failure to consider some otherwise familiar contingency that bears on the question at hand” (2021b: 1).

What propositions an agent happens to think of or find relevant is highly changeable, so on the occurrence account of awareness, the resulting algebra will be rather fickle. This is not a problem in itself; in fact, it can be very useful to model decision-making as it appears in the mind of the agent at the time of the decision. But I take it that it isn't what most philosophers, Steele and Stefánsson included, intend to do with probabilistic modelling. Instead, many of us are interested in more substantive facts about agents; in the lineage of Frank Ramsey (1926), we use credences to capture something robust about the agent, that has normative weight, in the sense that it informs or depicts the decisions they're disposed to take across a range of situations. We are less interested in the decisions agents happen to make based on temporary cognitive whims than in the decisions they would make based on a thorough sweep of what's available to them; that is, we care less about how agents do act than how they should act.

And when agents make decisions, they should entertain all the relevant propositions that are available to them. This suggests that we need a more demanding account of awareness, one that does not allow agents to count as unaware of propositions they have merely temporarily forgotten or excluded. Steele and Stefánsson give the outline of such an account when they describe an unawareness situation as one in which “the agent has no idea how to specify the propositions' content” (2021a: 1226), where they “lack the conceptual resources to entertain the possibilities in question” (2021b: 1) or “cannot (yet) articulate” them (2021b: 34). On this second account, awareness is understood as *intelligibility*: an agent counts as aware of all propositions they are able to understand or articulate. Thus the agent in Steele and Stefánsson's example who is

---

<sup>1</sup> The idea that an agent's algebra should be relativised to their awareness context is taken up in Mahtani (2021), Pettigrew (forthcoming), and Roussos (forthcoming).

unaware that the film could be a drama has not merely forgotten to include drama in their deliberation; drama is genuinely unintelligible to them, they do not have the conceptual apparatus necessary to make sense of this possibility. As the discussion in Canson (2024) shows, this account of awareness is the one that has traditionally been adopted in the literature.

On the intelligibility account of awareness, the relativisation of the algebra to the agent's awareness amounts to the inclusion in the algebra of all and only the propositions that the agent can articulate. Using Steele and Stefánsson's example, and assuming that the propositions *comedy* and *thriller* are the only propositions initially intelligible to the agent, the agent's algebra would contain these two propositions as what they call "basic propositions" (2021a: 1223), from which a Boolean algebra would be generated. It follows that the proposition *comedy*  $\vee$  *thriller* would amount to the universal proposition, traditionally denoted  $\Omega$ . More generally, on the account of awareness as intelligibility, the tautology is identified with the disjunction of all the intelligible propositions.

Provided that we accept probabilism—and, in the context of modelling limited awareness and awareness growth, Steele and Stefánsson call it "non-negotiable" (2021a: 1213)—this entails that agents ought to have credence 1 in the disjunction of propositions they are able to articulate. And provided that we accept, as Steele and Stefánsson also do, an orthodox kind of decision theory, this entails that agents ought to be willing to bet everything—to bet their life!—on this disjunction. This is patently absurd. The fact that the agent in Steele and Stefánsson's example cannot produce the name of a cinematic genre besides comedies and thrillers should not condemn them to accepting such a bet, because the mere fact that they cannot name such a genre does not entail that they think it cannot exist. This reveals the core of the issue with relativising algebras to the agent's state of awareness: what is intelligible to the agent does not coincide with what the agent takes to be possible, and it is the latter which is normatively relevant.

This speaks in favour of the view that Steele and Stefánsson built theirs in contrast to: what I elsewhere call "orthodoxy" (Canson 2024, 2). On orthodoxy, the algebra's sample space, and by extension its tautology, bear no trace of the agent. Instead, they are identified either with the logically or metaphysically necessary, or with the determinable a priori. This view entails that any agent with a limited awareness should be modelled with what's known as a *catch-all proposition*: the complement of the intelligible propositions against necessity. Steele and Stefánsson had rejected orthodoxy on these grounds: since catch-all propositions "cannot be articulated", it wouldn't be "cogent" for the agent to be modelled as having a credence in them (2021a: 1226). But in my discussion of catch-all propositions, I point out that these unintelligible propositions are perfectly mundane objects of credence: cinematic novices regularly think the genre of a film might be a genre they have never heard of; many scientists believe that the correct theory of the phenomenon they study hasn't yet been articulated. Agents routinely take such unintelligible propositions to be options for what might be the case, and as such, have a credence in them. Steele and Stefánsson's argument against catch-all propositions fails for the same reason their proposal to relativise algebras to the agent's awareness cannot be right: because there is a difference between

what is intelligible and what is possible, and it is the latter that is epistemically and practically relevant, and as such, that should feature in a model of normative decision-making.

However, I do raise the possibility that an agent may be mistaken about what is necessary. We might imagine for instance an agent who not only cannot name a film genre besides thriller and comedy, but who positively takes the existence of such a genre to be impossible. It seems that if we wanted to describe the decision problem from this agent's perspective, if we wanted to figure out what the agent's reasoning should be in this case, we should model them as entertaining *comedy* and *thriller* as disjoining to the universe—it is, after all, what they take to be necessary. Such an agent, unlike the one who merely fails to conceive of an alternative to comedies and thrillers, would be rational, from their own perspective, in betting their life on this disjunction; that is simply what it means to consider a proposition to be necessary.

This amounts to a relativisation of the algebra, not to the agent's state of awareness, but to what we might call the agent's *subjective modality*. Where on orthodoxy, the universe of discourse corresponds to an objective notion of necessity, here, it corresponds to what the agent considers necessary. There is precedent for such a view. Richard Pettigrew proposes that the sample space be composed of all the “personally possible worlds”, where “a world is personally possible for a particular individual at a particular time if by this time this individual hasn't ruled it out by”, among other things, “their logical [and] conceptual thinking” (2021: 9995). Thus the algebra contains all the propositions that the agent thinks could obtain, and the universal proposition  $\Omega$  corresponds to what the agent takes to be necessary. Importantly, the use of catch-all propositions remains necessary whenever an agent's awareness is limited, that is, whenever the propositions that are intelligible to them do not fill out their subjective modal space.

In their book, Steele and Stefánsson do endorse the occasional use of what they call a “subjective catch-all”, which is “meant to represent the possibilities that the agent cannot specify but thinks she might have left out of her reasoning”, to accommodate the fact that some agents are “aware of their unawareness” (2021b: 36). This stands in tension with the claim they make in their paper and which I discussed above, that catch-all propositions mustn't ever be used because they “cannot be articulated” (2021a: 1226). One might interpret this evolution as reflecting a change of view between their paper and their book, from the position that algebras must be relativised to the agent's awareness context on the intelligibility account of awareness, to the same position but on the occurrence account of consciousness. In their book, they would then be arguing that propositions, including catch-all propositions, ought to be included in the agent's algebra to the extent that they are actively entertained by the agent at the time of the decision. It would follow that an agent who is explicitly considering the possibility that an unspecified option may be correct would have to be modelled with a catch-all proposition even if, by definition, the agent cannot articulate this unspecified option.

Regardless of whether this interpretation is correct, it must be stressed that Steele and Stefánsson are not arguing, at any point, for the relativisation of algebras to the agent's subjective modality. For they are unambiguous in their claim that even “subjective” catch-all propositions are only to be used when the agent actively suspects that a heretofore unarticulated possibility might

obtain; in other words, in those cases where the agent actively entertains a catch-all proposition. But as we have seen, the relativisation of the algebra to the agent's subjective modality mandates the employment of a catch-all proposition whenever the propositions intelligible to the agent do not disjoin up to subjective necessity, whether the agent actively entertains this catch-all proposition or not. To illustrate the difference, where Steele and Stefánsson would model an agent who, because of a "momentary lack of perspective" (2021b: 1), forgets that there exist cinematic genres besides thrillers and comedies without a catch-all proposition; a proponent of algebra relativisation to the agent's subjective modality would employ a catch-all proposition, since the existence of an option beyond thriller and comedy is presumably not something that the agent considers to be impossible.

It is easy to slip between different views concerning the relativisation of Bayesian algebras to the agent's perspective. In fact, Pettigrew claims to inherit the relativisation of algebras to subjective modality from Hacking, who, it is true, had suggested that sample spaces "should consist of states of affairs each of which is 'possible to the agent'" (1967: 317). But in the same discussion, Hacking had stated that "much scientific learning consists in devising new hypotheses or forming new concepts. ... Since new hypotheses and new concepts typically lead to newly intelligible sentences, they lead to a new personal language" (318). This latter view is of course the view precisified and defended by Steele and Stefánsson, according to which algebras should be relativised to the agent's awareness context, which I have shown in this note must be abandoned.

Thus we must be careful. While the relativisation of algebras to the agent's awareness state, intended to reflect either what the agent happens to think about when they make a decision, or what is intelligible to them, faces insurmountable difficulties, the idea of relativising the algebra to the agent's perspective need not be abandoned. But it is not what the agent can articulate that matters but their *modal perspective*: not what they happen to think about or can express but what they take to be possible. As common slips between these show, they are not easily distinguished, but their differentiation is crucial for a formal theory of normative decision-making, for only the latter has any normative significance.

This has important consequences, unsurprisingly, for the modelling of awareness growth. In particular, it poses a problem for accounts of awareness growth that represent it by expansion, such as those defended by Bradley (2017), Steele and Stefánsson (2021a, 2021b) and Roussos (forthcoming). For these accounts treat agents undergoing awareness growth as moving from one algebra relativised to the agent's old state of awareness, to another algebra relativised to the agent's new state of awareness; they depend, that is, on the awareness-relativisation of algebras I have argued is problematic. This tells in favour of accounts of awareness growth that represent it by *refinement*; on which growths of awareness are accounted for by the splitting of a catch-all proposition, and which therefore function with algebras that are not relativised to the agent's awareness context. Elsewhere, I defend such an account on unrelativised algebras (Canson 2024); we can easily imagine extending it to algebras relativised to an agent's subjective modality.

## References

- Bradley, Richard. 2017. *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- Canson, Chloé de. 2024. “The Nature of Awareness Growth.” *Philosophical Review* 133(1): 1–32.
- Hacking, Ian. 1967. “Slightly More Realistic Personal Probability.” *Philosophy of Science* 34(4): 311–25.
- Mahtani, Anna. 2021. “Awareness Growth and Dispositional Attitudes.” *Synthese* 198(9): 8981–97.
- Pettigrew, Richard. 2021. “Logical Ignorance and Logical Learning.” *Synthese* 198(10): 9991–10020.
- Pettigrew, Richard. Forthcoming. “How Should your Beliefs Change When your Awareness Grows?” *Episteme*.
- Ramsey, Frank. 1926/1931. “Truth and Probability.” In R.B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*. London: Kegan Paul.
- Roussos, Joe. Forthcoming. “Awareness Growth and Belief Revision.” *Australasian Journal of Philosophy*.
- Shimony, Abner. 1970. “Scientific Inference.” In R. Colodny (ed.), *The Nature and Function of Scientific Theories*, 79–172. Pittsburgh: University of Pittsburgh Press.
- Steele, Katie, and H. Orri Stefánsson. 2021a. “Belief Revision for Growing Awareness.” *Mind* 130(520): 1207–32.
- Steele, Katie, and H. Orri Stefánsson. 2021b. *Beyond Uncertainty*. Cambridge: Cambridge University Press.
- Wenmackers, Sylvia, and Jan-Willem Romeijn. 2016. “New Theory About Old Evidence.” *Synthese* 193(4): 1225–50.