

# Elga & Rayo's "Fragmentation and Logical Omniscience"

DANIEL.HOEK@VT.EDU, EASTERN APA, JANUARY 9TH 2021

## 1. Introduction

- This paper exemplifies a shift away from thinking about the problem of Logical Omniscience as a technical challenge in doxastic logic and for the semantics of belief reports, and towards reconceiving of it as a problem for decision theory.
- Now, I am completely convinced this is the way to do things, and that this is the only way we can make real and lasting progress on the issue.
- In fact I have had that conviction for a few years already. And I think Adam and Agustín have had it for many years more. Perhaps a few other people in the fragmentation literature are leaning this way too. And I think Stalnaker already saw it back in the seventies, though he is not as explicit about it as you might wish. His work is probably the common cause here. There are also traces of this approach in Christopher Cherniak's *Minimal Rationality*.
- But at the end of the day, I think it's fair to say that in all those years we haven't really made many converts, and most of the literature on the topic has not been interested in the decision theoretic aspect of the story.
- So I thought a useful way for me to use these comments is to distinguish, a little more sharply and explicitly, these two approaches. And I'll try and sketch out one or two reasons to think that Adam and Agustín's approach is in fact superior.
- It will be interesting to hear if you guys agree with those motivations.
- I will also have some more critical comments about the particular details of the proposal. But I'll be saving those until the end, if there's time.

## 2. The Traditional Problem of Logical Omniscience

- The problem of logical omniscience is standardly understood as a problem for the logic of belief, or about the semantics of belief reports.
- Here, the question is: *How is it possible for deductive reasoning produce any new beliefs?*
  - This question is raised by what Ramsey and more recently Yalcin have dubbed the "map" picture of belief, which assumes that beliefs are consistent and deductively closed (or that your credences are probabilistically coherent — I'm going to gloss over that distinction). If that assumption is correct, then valid deductive reasoning can only lead to conclusions we already believed in the first place.
  - This is an especially urgent problem if you are under the spell of a functionalist or pragmatist picture of mind and world, or Dennett's *intentional stance*, or anything like that.

That's because by far the best ways we have of spelling out those pictures of belief imply, through the representation theorems, the map picture of belief.

- The semantic problem of logical omniscience is all about finding an alternative, less idealised model of belief or credence states.
  - In the case of belief states, this boils down to finding doxastic closure conditions that are weak enough to avoid collapse into logical omniscience, but strong enough to sustain some of our ordinary reasoning about belief.
  - Cherniak: "A collection of mynah bird utterances or snippets from the *New York Times* are chaos, and so (at most) just a sentence set, not a belief set." (*Minimal Rationality*, p. 6)

### 3. The Decision-Theoretic Problem of Logical Omniscience

- What is distinctive about Adam and Agustín's decision-theoretic approach is that it really tries to tackle an older, more ambitious question: *How can deductive inquiry ever be useful?*
  - What is the use of deductive inquiry, in view of the fact that by its nature it produces no new information?
  - You can find this problem in Plato's *Meno*, and it is also a prominent theme in Kant's work. Dummett calls it the Problem of Deduction.
- This problem is very closely related to the semantic problem.
  - Deduction would not be useful if it didn't make a difference to our doxastic state.
  - And if fruitful deductions do affect our doxastic state, then at least we have the beginnings of an account of how it can be useful.
- Nevertheless, I think it's helpful distinguish the two problems. The decision-theoretic problem is strictly harder.
  - If you have some fancy fine-grained account of doxastic states, that does not by itself tell you which of two informationally equivalent states is "more useful" than another.
  - **But Elga and Rayo's decision-theoretic solution *does* give you that!**
    - If you are comparing the access table of a really fragmented agent, to the access table of an agent whose beliefs are informationally equivalent, but less fragmented (in that they have fewer different rows in their table), then you will be able to prove that the less fragmented agent is going to be less susceptible to Dutch books.
    - Likewise, if one agent's belief fragments have a greater common core of information, and less variation between the fragments, then there are going to be fewer Dutch books you can make against them.

- In particular, that means that the closer agents are to being *globally Boolean* in Adam and Agustín's sense, the less Dutch Bookable they are.
- To sharpen the formulation of the central question a little bit, the decision-theoretic problem of logical omniscience asks how it is deductive accomplishments lead to *better choices*.
  - First you give an account of the (rational) behavioural dispositions of agents with doxastic states that are not deductively closed.
  - Then you give some systematic dynamic story about how deductive reasoning affects an agent's doxastic state
  - Then we are in a position to do a before-and-after comparison, and to say what is better about the posterior state.
  - In the literature on logical omniscience, this paper is almost unique in that it has this sort of form. In my opinion that makes it a blueprint that future work on logical omniscience should aim to follow.

#### 4. Comparison Between the Two Approaches

- Why? I said that the decision theoretic problem is harder than the semantic problem. So doesn't it make more sense to start with the easier problem, and do things step by step?
- Well, I think not. And the reason I think that is that the semantic problem is in a way too easy, and too amorphous to be fruitfully pursued on its own.
  - At the very least, if our ultimate goal is to solve the decision theoretic problem, we had better keep that goal in mind as we are pursuing the semantic problem: otherwise we are bound to go astray.
- To make that a bit more precise, let me note that **every** potential solution to the semantic problem involves a hyperintensional individuation of belief states, in the sense that states are sometimes distinguished in spite of containing the same information.
  - This is true of Elga and Rayo and even true of Stalnaker, the supposed enemy of hyperintensionality. He, too, distinguishes an unfragmented agent from an agent who has the same information, but distributed over two fragments.
- So questions about when to distinguish, and when to identify, intensionally equivalent belief states, and intensionally equivalent individual beliefs, are inevitable.
- Now there is a wealth of choices here, and few appreciable semantic constraints. So there are going to be lots of different ways to "solve" the semantic problem.
  - You can make hyperintensional distinctions using fragmentation, structured propositions, truthmakers, impossible worlds, language of thought, subject matters, etc...


- Direct intuitions about belief reports could provide constraints in theory, and whittle down the option space; but our intuitions about this are so foggy that they are not much help.
  - I myself do think that there is an exceptionally strong intuition that beliefs are closed under propositional parthood (in Yablo and Fine's sense).
    - E.g. if you believe Fido is a big bad wolf, you also believe Fido is a big wolf.
    - Closure under conjunction elimination is an instance of this.
  - Actually one of the niggles I have about Adam and Agustín's account is that it doesn't bear out this generalisation. (W.r.t. metalinguistic beliefs, conjunction elimination is only respected for conjunctions with "salient conjuncts"...)
    - But I'm not sure how compelling that objection really is. On the whole, the intuitive entailment data are unreliable, since they put pressure in the direction of complete deductive closure, which is exactly what we're trying to avoid.
    - The trouble is that our truth-value judgments about belief reports are not sensitive to the difference between "free entailments", which beliefs are automatically closed under, and entailments that are not quite free, but still very cheap.
- This lack of constraints makes it difficult to see which hyperintensional distinctions correspond to real distinctions between mental states, and which ones we've just made up. We're reduced to making some arbitrary choices.
  - And that is exactly what you see in solutions to the semantic problem of logical omniscience on its own terms: these accounts of belief are rife with arbitrary cutoffs — in the length of formulas, the number of reasoning steps, etc.
- **There are just too many hyperintensional distinctions we *could* draw, and on its own terms the semantic problem of logical omniscience does not give us any guidance to determine which ones we *should* draw.**
- A second, related difficulty with trying to solve the semantic problem in isolation is that it seems cheap, and theoretical bad practice, to **postulate arbitrary mental distinctions to save the semantics, without independent, psychological motivation**. It would be much better to have an account of what *makes* certain intensionally equivalent doxastic states distinct.
  - "One needs an account of what states of belief, desire and intention *are* that explains how the fine-grained structure of some notion of proposition contributes to distinguishing between different states of belief, desire or intention." (Stalnaker, *Context and Content*, p. 27).
- The decision theoretic standpoint gives us a principled way to resolve the arbitrariness problem and to meet Stalnaker's demand.
  - When belief states are systematically associated with distinct behavioural manifestations, that offers a compelling theoretical reason to distinguish them.

- It is relatively uncontroversial that beliefs are individuated at least as finely as truth conditions. One important reason for that is that truth-conditionally distinct beliefs are systematically associated with distinct choice dispositions/betting patterns.
- On the other hand, if a distinction between two belief states is so subtle as to have no behavioural manifestation at all, under any circumstances, then that's a good reason to be skeptical that we are really talking about distinct mental states.
  - You don't have to be a dyed-in-the-wool behaviouralist to feel this pressure. The sorts of marked failures of deductive closure that motivate the problem of logical omniscience are very clearly not of this kind.
  - In any case, drawing such super-subtle hyperintensional distinctions is not going to help with the decision-theoretic problem of logical omniscience.
- To sum up: the decision-theoretic approach to the problem of logical omniscience gives us a principled way to justify a certain particular way of individuating doxastic states. Direct intuitions about the truth values of, and entailment relations between belief reports can't do that. That's why Adam and Agustín's approach to the problem is, I believe, fundamentally the correct one.

#### 5. **Is the Problem of Logical Omniscience really about Metalinguistic Beliefs?**

- Now we get to the more critical portion of my comments. My first objection concerns the scope of Adam and Agustín's proposed solution.
- Adam and Agustín focus entirely on *metalinguistic* beliefs. That is, beliefs that are about the truth value of sentences.
  - Being locally and/or globally Boolean are properties that entirely supervene on an agent's credences about such metalinguistic propositions.
  - Likewise for the local and global sentential coherence properties.
- Correspondingly, they also focus on choice situations in which an agent (Watson) has to solve distinctively linguistic puzzles — i.e. to figure out whether a given complex formula is a tautology or a contradiction.
- Figuring out what's going on in this sort of situation could amount to a complete solution of the logical omniscience problem only if the use of deductive reasoning were restricted to its effects on linguistic choices.
  - Admittedly, standard examples of deductive ignorance do seem to involve linguistic actions: e.g. writing down the answer to a maths exercise, solving a logic puzzle.
- Still, it does not seem to be true in general that the use of deductive inferences is restricted in this way. Here are a number of examples of tasks that require abstract deductive reasoning to accomplish a non-linguistic goal through non-linguistic means:

- Solving a sudoku
- Unscrambling a Rubik's cube
- Solving a disentanglement puzzle, e.g. [this one](#) :
- A rat figuring out a maze.
- A crow working out a complicated mechanism in order to get a treat (like [this one](#))



- In each of these tasks, ordinary agents show trial-and-error behaviour even when they have all the information, which clearly evinces a certain lack of logical omniscience.
- But since the agents involved are not plausibly construed as reasoning about the truth of sentences here, Agustín and Adam's account of deduction has nothing to say about the deductive reasoning going on in cases like these.
  - Stalnaker: "I think it helps to see the problem [of logical omniscience] in a more general setting: to see linguistic action as a special case of action, and the use of knowledge to say how things are as a special case of the use of knowledge to make our actions depend, in ways appropriate to our ends, on the state of the world." (Logical Omniscience II, p. 272-3)
- The point of deductive inquiry is not just to improve one's linguistic skills (even if it does often do that as well).
  - Deductive inquiry can lead to non-trivial discoveries, scientific and otherwise. It can lead to knowledge of how to do something. It can allow you to navigate the world in all sorts of ways.
  - Sure, deductive inquiry can make us (linguistic creatures) better at language games, but that seems like it is just one very small part of how deductive accomplishments can be useful.
- One could reply that even in the sorts of cases I have mentioned, the deductive reasoning is really reasoning about "sentences" in the language of thought, or something like that.
  - But at that point, we are falling into a linguistic picture of thought that Stalnaker sought to banish, and which I don't think Adam and Agustín want to bring back.

## 6. Deductive Reasoning as the Accumulation of "Logical Information"

- As noted, Adam and Agustín shift their focus to metalinguistic credences. They use the notation  $P[A]$  to abbreviate  $P("A" \text{ is true})$ .
  - Since the meaning of sentences, and of logical connectives, is assumed to be contingent, the truth of any two sentences is basically independent.
  - So this allows us to indirectly assign probability values to sentences without regard for probabilistic coherence. That is to say, the "meta-credence" function  $A \mapsto P[A]$  need not be

a probability function.

- **Observation 1:** If we conditionalize  $P$  on semantic information about the connectives, then  $A \mapsto P[A]$  does become a probability function.
  - To be precise, we need to conditionalise on the info that  $\top$  is true, that every negation is true iff it's prejacent isn't, that every conjunction is true iff its conjuncts are both true, and that every disjunction is true iff either disjunct is.
- **Observation 2:** If we conditionalize  $P$  on a select part of that semantic information, then  $A \mapsto P[A]$  becomes more like a probability function.
- The proposal, then, is to model deductive reasoning as an update by partial semantic information, which brings the agent closer and closer to having a probabilistically coherent meta-credence function.
  - For instance, the agent would update with the information that certain particular negations are true iff their prejacent aren't, certain particular conjunctions are true iff both conjuncts are, and so on.
- More specifically (since we are dealing with fragmented agents), the updates are conceived as global updates, which lead to the conditionalisation of every credence fragment the agent has.
  - It's a little bit weird, in a fragmented setting, that such universal updates are even possible.
    - After all, the whole idea of the fragmentation picture is that the agent stores information in different places, and does not have the resources to distribute all their information across all of those locations.
    - But if "universal updates" were possible, then this would not be an issue...
  - Relatedly, an oddity of the approach is that fragmentation doesn't actually do any work in accounting for Watson's logical limitations, as far as I can see. All the work is done by the probabilistic incoherence of whatever meta-credence function is active.
  - When they talk about deductive reasoning, fragmentation theorists typically focus on the role of deduction and information processing in putting together information from different sources and fragments.
    - For instance, one fragment might contain the information that Jill is 8 and another fragment the info that Jane is 11. To infer that Jill is younger than Jane, you have to put that information together.
    - But conditionalizing on logical information is never going to lead to this sort of combination of information.
    - That raises a question for Adam and Agustín: do you agree that this is a role that deductive reasoning can play? And if so, how does it fit into your framework for thinking about deductive reasoning?

- I also have some reservations about the idea that what is “obvious” to an agent is context-dependent in the way Adam and Agustín say it is.
  - As mentioned above, I tend to think that the inference from any conjunction to its conjuncts should always be obvious, in any context. But according to Adam and Agustín, even the inference from a salient conjunction to its conjuncts need not count as obvious, unless those conjuncts are themselves also salient.
  - On the other hand, the obviousness function also seems to overshoot a bit, including things that are not immediately obvious in any context — as is admitted in footnote 20 (the poker example).
- I’m also wondering if more can be said about what it takes for a sentence to be salient in a choice situation.
  - In particular, in the Watson example on p. 15, it seems intuitively like the dilemma Watson renders all the sub-formulae of the long sentence salient.
  - But if that were so, no reasoning would be necessary. The answer should be obvious to Watson from the start...

## 7. In sum...

- Though I am a little skeptical about some of the details of their proposal, I think Adam and Agustín are exactly right about what a solution to the logical omniscience problem needs to do.
- In addition, the paper convincingly demonstrates through its own example that a solution of this shape is within our reach.
- Now that we know what our target should have been, I reckon discussions of logical omniscience may be about to get a whole lot more productive!