

Evidential support and its presuppositions

Luis Rosa*

Forthcoming in *Hinge Epistemology and Religious Belief*

Danièle Moyal-Sharrock and Nuno Venturinha (eds.)

1 Introduction

Philosophers sometimes impose requirements that seem hard or even impossible for a single theory to meet. And so it happens that, concerning certain topics in ethics, epistemology and metaphysics, no theory that has been fleshed out on those topics does meet all of them, and the prospects that there will be one such theory are bleak.

Fortunately, however, this appearance of impossibility isn't always accurate. Sometimes there is a way of accommodating all the disparate requirements—only it has escaped our attention so far. It might even be a rather simple one.

This paper features a case in point. The requirements that we will look into here are requirements on theories of *epistemic rationality*, and the paper tries to show that it is possible for a single theory to satisfy all of them. The target requirements go as follows. A theory of rationality should render the following sentences true, respectively:

- (a) Sometimes subjects are in possession of *conclusive evidence* for the truth of a proposition p (good enough to conclude that p).
- (b) But even in those cases the subject's evidence may not *completely* rule out the possibility that p is false.
- (c) The rationality of belief does not always coincide with the rationality of *high but non-maximal* credence.
- (d) Neither does the rationality of belief always coincide with the rationality of *maximal* credence.

*CONCEPT, University of Cologne.

The surface grammatical forms of these requirements perhaps suggests that they cannot all be satisfied at the same time. However, we claim, there is a candidate semantics for the target sentences that makes all of them true.

All of (a)–(d) can be satisfied at the same time, we claim, if only the sentences whereby we talk about sufficient evidential support and epistemic rationality are context-dependent in a specific way—namely, in much the same way that many natural language modals are taken to be, in the tradition of Kratzer (1977). In fact, constructions of adverbs and verbs such as ‘sufficiently supports’ and ‘conclusively supports’ will here be taken to behave similarly to necessity modals such as ‘must’ and ‘necessarily’.

Context will put *restrictions* on the kinds of possible worlds or scenarios that the contributions of those expressions are supposed to quantify over. In the canonical account, roughly put, ‘It must be that p ’ or ‘Necessarily p ’ are true just in case ‘ p ’ is true in all the possible worlds within some suitably restricted set of possibilities. Similarly, ‘ S ’s evidence conclusively supports p ’ will come out true according to the present proposal just in case ‘ p ’ is true in all the possible worlds within a suitably restricted set where S ’s evidence is true.

In our case, the relevant restrictions stem ultimately from the *presuppositions* that speakers make in the context they are in when they try to determine how a subject’s evidence bears on the truth-values of different propositions (a *speaker* is part of the context of her own utterance, together with a *time*, *location*, *possible world*, and possibly more items—this is the technical use that we are going to make of ‘context’ here).¹ This approach takes inspiration on Coliva’s (2016) suggestion that *hinge propositions* play the role of guiding us in assessing what counts as evidence for what. Only our ‘hinge propositions’ here are nothing other than the presuppositions that speakers make when they talk about evidence and evidential support.

That being said, we do not expect our notion of *presupposition* to satisfy all or even most of the things that philosophers have said about hinge propositions. It is totally okay (or maybe even expected) if it doesn’t.²

We start our investigation by looking at how both (a) and (b) could come out true without contradiction under such a construal of the semantics of the target sentences.

¹These items of a context are borrowed from Kaplan (1989). Lewis (1996) famously offers a contextualist account of knowledge ascriptions where presuppositions also play a role. See Blome–Tillmann (2014) for more on the relationship between knowledge ascriptions and presuppositions.

²Thanks to Jakob Ohlhorst for feedback on this point.

2 The compatibility of conclusive evidential support with the possibility of falsehood

Consider a case such as the following. The police is trying to determine *whether Mr. White robbed the bank*. While Mr. White waits at the secure interview room, Detective Constanza gains access to footage from the bank's surveillance camera. Previously, Constanza has spent about an hour talking to Mr. White at the interview room. So Constanza definitely knows *how Mr. White looks like*.

The footage that is now being inspected by Constanza shows a Mr. White resolutely holding a gun and pointing it at the employees as he enters the bank's building. One of the bank employees nervously hands Mr. White some money. After collecting the money, Mr. White heads to the exit door and disappears behind it. When it comes to determining *whether Mr. White robbed the bank*, Constanza knows nothing else that counterbalances that evidence regarding the matter.

Concerning a case such as this one, we find it proper to utter sentences such as 'Now Constanza has conclusive evidence that *Mr. White robbed the bank*'. Conclusive evidence that p is evidence that is good enough to conclude that p (to believe that p and settle the question of *whether p*). Before seeing the footage, Constanza might have had some evidence to suspect that Mr. White did it—e.g. there was no one else to testify about Mr. White's alleged whereabouts at the time of the robbery, Mr. White looked like he was lying, etc. But now... it's on film, what else could we ask for? It is now rational for Constanza to conclude that *Mr. White robbed the bank*.

On the other hand, we also find it proper to utter sentences such as 'It is of course possible that *Mr. White did not rob the bank*, given Constanza's evidence' or 'Constanza's evidence doesn't completely rule out the possibility that *Mr. White didn't rob the bank*'. We find it proper to say these things concerning Constanza's situation even after becoming aware that he has seen the footage from the bank's surveillance camera.³

If we press Constanza in the right way, he will even agree with us on these remarks. Could it not be that the officer who transferred the video file from the the bank to the police station edited the video with some uncanny and extremely fast technology before reporting to the Bureau? That is *incredibly unlikely*—other officers would have to have cooperated with him, the time-interval was very short, etc.—but it cannot be *completely* ruled out.

³Relatedly, on the connection between infallibilism and contextualism about knowledge ascriptions, see Venturinha (2020).

Not, at least, by Constanza's total evidence.

How can 'Constanza has conclusive evidence that *Mr. White robbed the bank*' and 'It is possible that *Mr. White did not rob the bank* given Constanza's evidence' both be true at the same time and possible world, as uttered from the same context of utterance?

In a nutshell, they could both come out true if this was a case of multiple modalities: 'is conclusive evidence that...' and 'It is possible that... given the evidence' do not always quantify over one and the same space of possibilities. Context might put different restrictions on them (in particular, the one modal isn't just the *dual* of the other).

To work this out in more detail, let us think of the presuppositions that restrict the space of possibilities that 'Constanza's evidence is conclusive evidence that...' quantifies over as being *lifted* in the determination of the space of possibilities that 'it is possible that... given Constanza's evidence' quantifies over.

The former space can be *properly contained* within the latter space, even though both of them are constituted only by possibilities where Constanza's evidence is true/accurate. That is, we think of the context where those expressions are used as providing a set of propositions that are assumed/known to be true by the interlocutors of that context (the presuppositions of the conversational background),⁴ which then determines a modal base—some set of possible worlds—relative to which we assess the truth or falsehood of the sentences where they occur. Then some propositions that are assumed to be true in order to determine the modal base of 'Constanza's evidence is conclusive evidence that...' are not assumed to be true in order to determine the modal base of 'it is possible that... given Constanza's evidence', even in cases where the placeholder '...' is filled with the very same declarative sentence.

The truth of claims about *evidential support*, then, can hinge on presuppositions that the truth of claims about *possibility given the evidence* do not hinge on.

⁴See Stalnaker (1970) on this type of context dependence of modal expressions in general, as well as Kratzer (1981). The way van Fraassen puts it is particularly congenial to our approach here: 'The context will generally select the proposition expressed by a given sentence A via a selection of referents for the terms, extensions for the predicates, and functions for the functors (i.e. syncategorematic words like 'and' or 'most'). But intervening contextual variables may occur at any point in these selections. Among such variables there will be the assumptions taken for granted, theories accepted, world-pictures or paradigms adhered to, in that context. A simple example would be the range of conceivable worlds admitted as possible by the speaker; this variable plays a role in determining the truth-value of his modal statements in that context, relative to the pragmatic presuppositions' (van Fraassen 1980, p. 137).

3 Truth-conditions under which (a) and (b) are made mutually consistent

In order to make the proposal more precise, let us first construe of a subject's total body of evidence at a possible world w as a set of propositions E_w .⁵ That is, E_w contains a proposition p if and only if it is part of the subject's evidence that p at w . To that set there will correspond a set of possible worlds $e(w)$, namely, the set of all possible worlds where all the propositions that belong to E_w are true. So $e(w)$ represents the ways in which the world could be as far as the subject's evidence goes.

We further let H_c be the set of propositions that are presupposed to be true at a context $c = \langle \text{interlocutors, time, place, world, ...} \rangle$. To that set there will also correspond a set of possible worlds $h(c)$, namely, the set of worlds where all the presuppositions that are operative in context c are true. So $h(c)$ represents the ways in which the world could be as far as what is taken for granted by the interlocutors of c goes.

Given that much, we can now say:

(CE) 'The subject's evidence is conclusive evidence that p ', as uttered in c , is true at w if and only if (i) p is true at every possible world that belongs to both $e(w)$ and $h(c)$, and (ii) p is not a member of H_c .

We can think of it as follows. Take first the possible worlds that belong to $e(w)$, or the worlds that are compatible with the subject's total evidence. Now create a new set of possible worlds out of the elements of $e(w)$, namely, those that also belong to $h(c)$. In the end, then, we have a set of possible worlds—the ones belonging to the *intersection* of $e(w)$ and $h(c)$ —that are compatible not only with the subject's total evidence, but also with all the presuppositions H_c that are made by us in context c . That is, the set we end up with is the set of possibilities that are left open by the subject's evidence *when restricted by the target presuppositions*.

Finally we check, first, if p is true at all the members of that set and, second, if p is not itself one of the presuppositions in H_c . If both of these conditions are satisfied, then the claim that the subject's evidence is conclusive evidence that p , as made in context c , comes out true; otherwise it comes out false. (Note that, where p is one of the presuppositions in H_c ,

⁵For arguments to the effect that evidence should be propositional in order to play the roles it does, see Williamson (2000, Ch. 9). We are thinking of possible worlds as *centered* on time here, so that we don't have to flesh out that parameter every time in our formulation (though sometimes we do). Also, we are assuming E_w to be minimally *consistent*.

'The subject's evidence is conclusive evidence that p ' as uttered in c must be false. Consider, for example: 'Constanza's evidence is conclusive evidence that *his eyes are functioning properly*'—more on this below). Here, conclusive evidential support comes off as a kind of *epistemic necessity*.

We talked above of presuppositions being lifted when we make the switch from evidential support-talk to relative possibility-talk (possibility relative to the evidence). To this lifting of presuppositions there corresponds a *widening* of $h(c)$, that is, a superset of $h(c)$ which includes more possible worlds than the ones that were already included in $h(c)$ (possibility is *easy*). There should be rules determining which widenings are permissible in context c , or rules as to which presuppositions can be lifted when talk of relative possibility kicks in at context c . We are going to assume that there are such rules here, without so much as trying to flesh them out here.

Now we can also state:

(CP) 'It is possible that p , given the subject's evidence', as uttered in c , is true at w if and only if (iii) there is a possible world v that belongs to both, $e(w)$ and a permissible widening of $h(c)$, such that p is true at v , and (iv) p is not a member of H_c .

We can think of it as follows. Take first the intersection of $e(w)$ and $h(c)$, namely, the set of possible worlds where the subject's evidence is true, further restricted by the presuppositions of H_c . Possibly, that set is even smaller than $e(w)$ itself, or it is a proper subset of $e(w)$. Now think of an operation of constructing a bigger subset of $e(w)$ —bigger than the previous subset of $e(w)$ —by including some possibilities in it that are excluded by some of the presuppositions in H_c . Finally we check, first, if there is a p -world inside that bigger subset of $e(w)$ and, second, if p is not itself one of the presuppositions of H_c . If both of these conditions obtain, then the claim that it is possible that p given the subject's evidence (or that it might be that p given the subject's evidence), as made in context c , comes out true; otherwise it comes out false.

In this way, our first pair of claims (a) and (b) come out true, with no inconsistency between them:

- (a) Sometimes subjects are in possession of *conclusive evidence* for the truth of a proposition p (good enough to conclude that p).
- (b) But even in those cases the subject's evidence may not *completely* rule out the possibility that p is false.

Consider again the case of Detective Constanza from Section 2. The sentence ‘Constanza’s evidence is conclusive evidence that *Mr. White robbed the bank*’ as uttered in our imagined context c can be true at our imagined world w even though the sentence ‘It is possible that *Mr. White didn’t rob the bank* given Constanza’s evidence’, as uttered in the very same context c , is also true at that world w . For the fact that *Mr. White robbed the bank* is true in all the worlds at the intersection of $e(w)$ and $h(c)$ (which makes the former sentence true) is compatible with the fact that *Mr. White robbed the bank* is false at some worlds at the intersection of $e(w)$ and a permissible widening of $h(c)$ (which makes the latter sentence true).

We now move on to generalize this framework to talk of evidential probability and ascriptions of rational credence.

4 The relationship between rational high credence and rational belief

Our next pair of claims is, remember:

- (c) The rationality of belief does not always coincide with the rationality of *high but non-maximal* credence.⁶
- (d) Neither does the rationality of belief always coincide with the rationality of *maximal* credence.⁷

Let us say that rational credences are credences that mirror evidential probabilities.⁸ More carefully, ‘It is rational for the subject to have credence/degree of confidence x that p ’ (where ‘ x ’ can either stand for a real number or an interval of real numbers), as uttered in c , is true at w if and only if ‘The probability that p is the case given the subject’s evidence is/lies within x ’, as uttered in c , is true at w .

Furthermore, we are also assuming that ‘It is rational for the subject to believe that p ’, as uttered in c , is true at w if and only if ‘The subject’s evidence is conclusive evidence

⁶See Friedman (2013) for the related claim that one is permitted to suspend judgment about p even in cases where one’s evidence vindicates very high/very low credence in p . See also Buchak (2014) and Staffel (2016), who defend the view that it can be rational for one not to believe that p when one’s evidence vindicates high credence in p .

⁷Accordingly, Worsnip (2016) makes the point that it is perfectly coherent for one to believe something one is less than maximally confident of.

⁸See Eder (2019) for more on the relation between evidential probabilities and rational credences.

that p' , as uttered in c , is true at w (an evidentialist notion of rational belief). So the idea that we are exploring here is roughly the following. Just like *conclusive evidential support* is compatible with the *possibility of falsehood*, so the *rationality of belief* is compatible with the *rationality of less than maximal credence*.⁹

Crucial to our strategy is that evidential probability talk will be constrained by the same presuppositions that are operative in relative possibility talk (relative to the subject's evidence). The semantic contributions of expressions such as 'probable', 'likely', etc.—as deployed in a particular context—will be probability functions that measure the probabilities of propositions conditional on the subject's body of evidence. Those functions don't need to assign probability 1 to the presuppositions that belong to the set H_c , however, whose job is again to restrict the set of possibilities that are left open by the subject's evidence to determine what it conclusively gives support to. That is, where a is a member of H_c , it may happen that the the relevant function assigns probability $x < 1$ to a .

In this framework, it becomes possible for 'It is rational for the subject to invest very high credence in p' ' to be true even though 'It is rational for the subject to believe that p' ' is false. That can happen, say, in contexts where we evaluate standard lottery situations.

Let p be the proposition that *the subject's ticket is a loser*. Within the set of possibilities $e(w)$ that are left open by the subject's evidence, further constrained by all our contextual presuppositions—that is, those members of $e(w)$ that are also members of $h(c)$ —there are possible worlds where p is false (possible worlds where the subject's lottery ticket is the winner). So the sentence 'It is rational for the subject to believe that p' ' comes out false. And, yet, there are enough of those possible worlds to make it very likely that p on the subject's evidence. In fact, it can be *very probable* that p conditional on the subject's evidence *even without* further constraining by some of our contextual presuppositions, under some permissible widening of $h(c)$. So the sentence 'It is rational for the subject to invest very high credence in p' ' comes out true.

Here it also becomes possible for 'It is rational for the subject to believe that p' ' to be true even though 'It is rational for the subject to invest maximal credence in p' ' is false. That can happen, say, in contexts where we evaluate propositions that are known without being accompanied by rational subjective certainty (belief on them is not grounded on purely statistical evidence, and yet the evidence itself does less than entail them). A candidate example is where we evaluate the subject's opinions about the lottery after she has read in

⁹That would be yet another version of *dualism* about belief and credence, in the sense of Jackson (2022).

the newspaper that a ticket has already been drawn, and she notices that the ticket that has been drawn is not her ticket.

Let p again be the proposition that *the subject's ticket is a loser*. Within the set of possibilities that are left open by the subject's evidence, further constrained by all our contextual presuppositions, there are only possible worlds where p is true. That is, the intersection of $e(w)$ and $h(c)$ is made of p -worlds only. So the sentence 'It is rational for the subject to believe that p ' comes out true. And yet some worlds that belong to $e(w)$ and a permissible widening of $h(c)$ are worlds where p is false (the permissible widening stems, say, from the lifting of the presupposition that the newspaper is making an accurate report about the lottery). So the sentence 'It is rational for the subject to invest maximal credence in p ' comes out false.

And so it is that, according to the present proposal, not only (a) and (b) come out true, but so do (c) and (d). And these are all mutually consistent.

5 Irrationality versus arationality

The semantic theory put forward above allows us to truly say some of the things that we wanted to say, though those very things seemed at first not to be mutually satisfiable. We can endorse all of (a)–(d) at the same time without inconsistency. But at what price? It bears on us to at least make some of the consequences of the theory explicit, insofar as they can be perceived as undesirable consequences.

Our theory pairs sentences of types 'It is rational for the subject to believe that p ' and 'The subject's evidence is conclusive evidence that p ' respectively through their context-dependent truth-conditions. Now remember the conditions that need to be satisfied in order for (tokens of) sentences of those types to be true, as uttered in a context c , at world w : (i) p is true at every possible world that belongs to both $e(w)$ and $h(c)$, and (ii) p is not a member of H_c . It follows from the second condition that we cannot truly say that it is rational for a subject to believe that p when p is part of the presuppositions that we use to determine what her evidence gives support to.

For example, in Detective Constanza's case from above, as we uttered the sentence 'Constanza's evidence is conclusive evidence that *Mr. White robbed the bank*', we have presupposed that Constanza's eyes were functioning properly when he inspected the footage

(the footage, that is, from the bank's camera). In that same context of utterance, our theory says again that 'It is rational for Constanza to believe that *Mr. White robbed the bank*' is true. But it also says that 'It is rational for Constanza to believe that *his eyes were functioning properly when he inspected the footage*' is false.

The proposition embedded in the latter ascription is rather one of the assumed preconditions that we rely on when we try to determine what Constanza's evidence gives support to. We have something of a bug in our thought, then, when we look at an ascription of rationality for Constanza to believe those very preconditions.

This is a quite general phenomenon: while 'it makes sense' to say that it is rational for the subject to believe that *the cat is on the mat* (given her perceptual evidence), it 'doesn't make much sense' to say that it is rational for her to believe that *there is space* (a precondition for her perceptual evidence to connect up to the proposition that *the cat is on the mat* in the right way). Where pre-theoretical thought struggles with making sense of the latter sort of claim, our theory ascribes falsehood straightaway.

Does the theory entail, then, that at the target possible world the sentence 'It is irrational for Constanza to believe that *his eyes were functioning properly when he inspected the footage*', as uttered in the same context as before, is true? That would be a bad result. Granted, it is again weird to say that it is *rational* for Constanza to believe the relevant proposition—but that doesn't mean that it is *irrational* for her to believe it, either.

We need not have the expressions 'rational' and 'irrational' behave as contradictories, however. The proposal we have fleshed out above is not committed to the claim that a sentence of type 'It is irrational for the subject to believe that *p*' is true whenever a sentence of type 'It is rational for the subject to believe that *p*' is false. We can perfectly well combine that proposal with the thesis that, in some cases (for some contexts and possible worlds), neither a sentence of type 'It is rational for the subject to believe that *p*' nor a sentence of type 'It is irrational for the subject to believe that *p*' are true, even assuming that the subject exists, etc.

Where it is neither rational nor irrational for a subject to believe that *p*, we might say that it is *arational* for her to believe that *p*, or that believing *p* is beyond the pale of assessment regarding rationality for that subject (as assessed from that particular context).¹⁰ In particular, we can say this from a context where *p* is one of the presuppositions that we

¹⁰Following a proposal recently put forward by Lopez (2023), our presuppositions would then have more things in common with so-called 'hinge commitments', this time on account of their arationality.

rely on to determine which bodies of evidence give support to which propositions.

(Relatedly, the propositions that are part of the subject's evidence set are to be distinguished from the ones that belong to our set of presuppositions. Only the former ones count as *evidence*. The presuppositions establish what counts as evidence for what, without themselves being evidence).

Another important observation that bears on the present concern is the following. That a token of 'It is rational for the subject to believe that p ' comes out false, and this is due to the fact that p is one of the presuppositions that we rely on in the context where we utter it, does not entail that it must come out false in other, perhaps only slightly different contexts.

To continue with the example from above, it is not impossible for us to produce a true utterance of 'It is rational for Constanza to believe that *his eyes were functioning properly when he inspected the footage*', if only we switch to a context where we do not presuppose that *Constanza's eyes were functioning properly when he inspected the footage* anymore. As seen from the latter context, it won't be arational for Constanza to believe that *his eyes were functioning properly when he inspected the footage*. Rather, it will be either rational or irrational for him to believe that. Verdicts of arationality are context-sensitive, as much as verdicts of rationality are.

Our theory, then, does not ascribe *irrationality* to beliefs toward the presuppositions from which we determine how a certain body of evidence bears on the truth-values of propositions. And their *arationality* is not written in stone, either.

6 Closure failure

Another possibly objectionable feature of our account is that it doesn't allow rationality to be closed under entailment, in the following sense. Even though ' p ' entails ' q ', there can be a context c and a possible world w such that 'It is rational for the subject to believe that p ', as uttered in c , is true at w , even though 'It is rational for the subject to believe that q ', as uttered in c , is not true at w (notice that the failure holds even relative to a single context).¹¹

¹¹That is an important difference between the form of contextualism about rational belief-ascriptions and evidential support talk put forward here and the typical contextualist account of knowledge-ascriptions. Contextualists about knowledge-ascriptions are usually keen on holding on to closure principles, and they take that to be an advantage of their view—see

The failure holds on account of the fact that, when we are trying to determine how the subject's evidence bears on the truth-values of propositions, we may presuppose things that follow from what is rational for that subject to believe. It can then be *rational* for her to believe a proposition while it is *arational* for her to believe something that follows from that proposition. For example, it may be rational for the subject to believe that *she had eggs for breakfast yesterday* while it is arational for her to believe that *there is a past*.

How worrisome is it for our theory to have this feature? Here are some reasons not to worry too much about it. There is evidence, there is what it supports, and then there is that on which the support relation between them hinges. Our theory doesn't allow us to truly ascribe rationality to believe the latter—at least not when we are using its evidentialist notion of epistemic rationality (rational belief is belief whose content is conclusively supported by the evidence). But that doesn't mean that there can't be other epistemic statuses, including ones that might be ascribed through words such as 'rational', 'entitled', 'justified', etc., for which some form of closure principle holds true.¹²

Furthermore, and similarly to what we already saw in the previous section, that it is arational for the subject to believe a conclusion that follows from premises that are rational for her to believe—because that conclusion is a presupposition in virtue of which it is rational for her to believe the premises—does not entail that it cannot be rational for that subject to believe that conclusion. Moving to a context where the target conclusion is not presupposed to hold true (in order to determine which propositions are supported by the subject's evidence) may allow us to truly assert that it is rational for that subject to believe that conclusion, too.

7 Conclusion

We wanted a theory of epistemic rationality that renders the following claims true:

- (a) Sometimes subjects are in possession of *conclusive evidence* for the truth of a proposition p (good enough to conclude that p).

Lasonen-Aarnio (2017) for more on this.

¹²Say, for example, a principle of knowledge-expansion under competent deduction such as the one put forward in Hawthorne (2004) might still be true (when you know the premise, and you competently deduce the conclusion from the premise while still maintaining your knowledge of that premise, you thereby come to know the conclusion, too).

- (b) But even in those cases the subject's evidence may not *completely* rule out the possibility that p is false.
- (c) The rationality of belief does not always coincide with the rationality of *high but non-maximal* credence.
- (d) Neither does the rationality of belief always coincide with the rationality of *maximal* credence.

We did that by fetching evidentialist notions of rational belief, rational credence and epistemic possibility that are context sensitive. In particular, they are sensitive to the presuppositions that ascribers make/do not make when they are trying to determine how the subject's evidence bears on different propositions. We pay a certain price to endorse the theory, of course. For the reasons exposed above, however, we don't think that it is such a big price to pay, in exchange for its own theoretical virtues.

References

- Blome-Tillmann, Michael (2014). *Knowledge and Presuppositions*, New York: Oxford University Press.
- Buchak, Lara (2014). 'Belief, Credence, and Norms', *Philosophical Studies* 169(2): 285–311.
- Coliva, Annalisa (2016). 'Which Hinge Epistemology?', in A. Coliva and D. Moyal-Sharrock (eds.) *Hinge Epistemology*, Leiden: Brill, pp. 6–23.
- Eder, Anna-Maria (2019). 'Evidential Probabilities and Credences', *British Journal for the Philosophy of Science*: 1–21.
- Friedman, Jane (2013). 'Rational agnosticism and degrees of belief', In T. S. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology*, Vol. 4, Oxford: Oxford University Press, pp. 57–81.
- Hawthorne, John (2004). *Knowledge and Lotteries*, New York: Oxford University Press.
- Jackson, Elizabeth (2022). 'On the Independence of Belief and Credence', *Philosophical Issues* 32(1): 9–31.
- Kaplan, David (1989). 'Demonstratives', in Almog, Perry, and Wettstein (eds.) *Themes from Kaplan*, Oxford: Oxford University Press, pp. 481–563.

- Kratzer, Angelika (1977). 'What 'Must' and 'Can' Must and Can Mean', *Linguistics and Philosophy* 1: 337–355.
- Kratzer, Angelika (1981). 'The Notional Category of Modality', in H. J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts*, Berlin: W. de Gruyter, pp. 38–74.
- Lasonen-Aarnio, Maria (2017). 'Contextualism and closure', in J. J. Ichikawa (ed.) *The Routledge Handbook of Epistemic Contextualism*, New York: Routledge, pp. 156–167.
- Lewis, David (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74(4): 549–567.
- Lopez, Aliosha B. (2023). 'Hinge commitments as arational beliefs', *Synthese* 201(109): 1–26.
- Staffel, Julia (2016). 'Beliefs, Buses and Lotteries: Why rational belief can't be stably high credence', *Philosophical Studies* 173 (7): 1721–1734.
- Stalnaker, Robert (1970). 'Pragmatics', *Synthese* 22(1/2): 272–289.
- van Fraassen, Bas (1980). *The Scientific Image*, Oxford: Oxford University Press.
- Williamson, Timothy (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Worsnip, Alex (2016). 'Belief, Credence, and the Preface Paradox', *Australasian Journal of Philosophy* 94(3): 549–562.
- Venturinha, Nuno (2020). 'Non-Skeptical Infallibilism', *Analysis* 80(1): 186–195.