How to be both a Quinean and a Bayesian

Damien Rochford

Many of us contemporary philosophers want to be Quineans, and also want to be Bayesians. But that is a difficult thing to do, for reasons I will explain in section 1. Happily, there is a solution. It involves thinking of many epistemic distinctions, such as the distinction between rational and irrational belief revision, as relative, rather than absolute. I expound the view, including formal details about how it should work, in section 2. I explain how the view allows one to reconcile Quinean and Bayesian commitments in section 3. (There is also an appendix, in which I give an account of what it is to have a credence in a merely apparent possibility — a notion I make use of in section 2.)

The view I am presenting is a version of what Daniel Greco calls 'contextualist foundationalism'.[1] As such, it will be a third way between the coherentism suggested by a straightforward interpretation of Quine, and the (absolute) foundationalism suggested by a straightforward interpretation of the Bayesian framework.

My proposal is a way of making precise Quinean scepticism concerning the *a priori*.

1 The Problem

It is hard to be both a Quinean and a Bayesian. Here are three reasons why:[2]

[1] Greco (unpublished). Greco himself recommends contextualist foundationalism as a way of reconciling Bayesian update with epistemic holism.

[2] Note that, in what follows, I will mostly assume that a Bayesian is committed to the standard Bayesian update rule — i.e., conditionalisation. But nothing turns on that; you can give a version

1.1 Frontloading

Consider the following argument, variations of which have been used in many places in the literature[3] to argue that agents have a lot of *a priori* justification for things. Let H be some proposition an arbitrary agent A has justification to believe:

P1	A is either justified in believing H *a priori* or on the basis of empirical evidence.
P2	If A is justified in believing H on the basis of empirical evidence, then there is some empirical evidence E such that she is justified in believing that if E, then H *a priori*.
C	A is *a priori* justified in believing either H or that if E, then H.

P1 appears to be trivial. Here is an argument for P2: suppose A is justified in believing H on the basis of empirical evidence E'. Then she must be justified in believing if E' then H. She is justified in believing *that* either *a priori* or on the basis of empirical evidence E''. If the later, then she must be justified in believing that if E' and E'' then H. She is justified in believing *that* either *a priori* or on the basis of empirical knowledge E'''... Repeat until you hit something the agent is justified in believing *a priori* (that must happen, as eventually you will have put all of the empirical evidence there is to put into the antecedent, so if the agent has justification for believing the conditional it has got to be *a priori*). Now just let the E in P2 stand for whatever long conjunction of evidence propositions you got by repeating this process. End of argument. This justification for P2 is why the argument above is sometimes called the 'Frontloading Argument': we keep loading evidence propositions into the front of the conditional.

As H was arbitrary, this argument works for all H. So it appears we have lots of *a priori* justification for things. What a horribly un-Quinean conclusion!

---

of each of these problems for any update rule I can think of. That's relatively obvious for the problems of subsections 1.1 and 1.3; for the problem of subsection 1.2, see footnote 7.

[3]See, for example, chapter 4 of Chalmers (2012), White (2006), Wright (2002), Schiffer (2004) and the 'Explainer' example in Hawthorne (2002).

The Quinean will want some way to resist the Frontloading Argument. But that is particularly hard to do, if our Quinean is also a Bayesian. A version of this argument appears to just fall out of the Bayesian framework. It goes like this: let H be some proposition in which an agent A is justified in having high credence...

P1 A is either justified in having a high, *a priori*, prior credence in H, or she has high credence in H on the basis of some empirical evidence.

P2 If A has high credence in H on the basis of empirical evidence, then there is some empirical evidence E such that she justified in having a high, *a priori* conditional credence in H, given E.

C A is *a priori* justified in having either high credence either in H, or high conditional credence in H given E.

P1 is indeed trivial, given the Bayesian framework. P2 seems to be an inevitable consequence of the Bayesian framework too. Here's the argument: suppose A is justified in having high conditional credence in H on the basis of empirical evidence E'. Then she must be justified in having high conditional credence in H given E'. She has justification for *that* either *a priori* or on the basis of empirical evidence E''. If the later, then she must have justification for having high conditional credence in H given E' and E''. She has justification for *that* either *a priori* or on the basis of empirical evidence E'''... Repeat until you hit something the agent has justification for *a priori*.

This argument for P2 relies on the Bayesian version of what David Chalmers has called the 'Frontloading Principle':

> Frontloading Principle: If an agent has justification for high conditional credence in P, given Q, on the basis of R, she has justification for high conditional credence in P, given Q *and* R.

If you are happy to talk about propositional justification, and you are committed to conditionalisation — i.e., the standard Bayesian update rule — then there doesn't appear to be any way around commitment to the Frontloading Principle.

That's bad news, for those of us who want to be both Quineans and Bayesians. How do we reconcile the Frontloading Principle, which we are committed to as Bayesians, with our scepticism concerning the *a priori*?[4]

1.2 Defeat

[5]Suppose I learn that

> Evidence: The majority of Australians currently intend to vote for the Liberal-National Coalition in the next election.

and on the basis of this I quite rationally raise my confidence in

> Hypothesis: The Coalition will win the next election.

On the standard Bayesian way of modelling things, what happened was this: I started off with a certain credence in Hypothesis, which we can write as '$C(\text{Hypothesis})$'. Then, when I learned Evidence, I updated my credences by conditionalising on Evidence. As a result, my new credence in Evidence was my old conditional credence $C(\text{Hypothesis}|\text{Evidence})$. Given that I raised my credence in Hypothesis, the following must be true: $C(\text{Hypothesis}|\text{Evidence}) > C(\text{Hypothesis})$.

[4]Just to show how hard this problem is, here is a small digression. Any view on which you can learn empirically about evidential relationships is going to want to resist the Frontloading argument, as it is exactly these evidential relationships that the argument purports to show are *a priori*. A view on which one can learn empirically about evidential relationships that has got a lot of recent attention is Dogmatism, á la Pryor (2000). Though he doesn't exactly put it this way, White (2006) effectively points out that this feature of Dogmatism is incompatible with Bayesianism, as it is usually conceived. I know of two responses to this problem on behalf of Dogmatism. One is due to Brian Weatherson (Weatherson (2007)); the other is due to Pryor himself (Pryor (unpublished)). Neither clearly escape all versions of the Frontloading Argument. Weatherson would need to say more about how update is supposed to work in his model before we could tell, but it is hard to see any way of filling in the details that gives you a way out of all versions of the Frontloading Argument (even if, perhaps, it gets you out of the version built on conditionalisation; you can give an analogous argument for pretty much any update rule). Pryor is much more forthcoming about the details of update on his model, but it is unclear what the relationship is between his update rule and rationality (in particular what rational constraints there are on an agent's range of 'inheritance rules', as Pryor calls them); in any case, he says nothing which suggests a way out of (all versions of) the Frontloading Argument.

[5]My presentation in this section is much indebted to Greco (unpublished).

There are many ways in which you can make me re-lower my confidence in Hypothesis — i.e., you can *defeat* my evidence, as epistemologists say. The most straightforward way is to provide some countervailing evidence which outweighs Evidence; you could, for instance, tell me

> Rebutter: Unbeknownst to the electorate, the Coalition is about to introduce legislation that will be very unpopular.

This can defeat my evidence because $C(\textsc{Hypothesis}|\textsc{Evidence} \wedge \textsc{Rebutter})$ need not be greater than or equal to $C(\textsc{Hypothesis}|\textsc{Evidence})$; it could be substantially lower. So when I conditionalise first on Evidence, then on Rebutter, I can first raise then lower by credence in Hypothesis. The fact that this is possible is sometimes called the 'non-monotonicity of confirmation', and it is one of the things the Bayesian apparatus is particularly useful for understanding.

But here's something which the standard Bayesian model does not handle so well. Another way to defeat my evidence is to attack Evidence itself. Maybe, you might say, I shouldn't be so confident that the Coalition will win the next election because I shouldn't be so confident that the majority of Australians currently intend to vote for them, because

> Undercutter: The polling data comes from a source with a flawed methodology.

It would be nice to treat this case like the previous one — that is, understand the defeat as resulting from the fact that $C(\textsc{Hypothesis}|\textsc{Evidence} \wedge \textsc{Undercutter})$ is lower than $C(\textsc{Hypothesis}|\textsc{Evidence})$. But $C(\textsc{Hypothesis}|\textsc{Evidence} \wedge \textsc{Undercutter})$ is *not* lower than $C(\textsc{Hypothesis}|\textsc{Evidence})$. Even if the polling data is not credible, it remains true that the majority of Australians now intending to vote for the Coalition makes it more likely that the Coalition will win.

If Undercutter really does defeat my evidence for Hypothesis, then it looks like my evidence needs to be something other than what we originally said it was; it needs to be something like

EVIDENCE′: The Coalition is ahead in the polls.

But this, it seems, can easily be undercut by something like

UNDERCUTTER′: The relevant news article reported the poll results incorrectly.

By a similar argument to the above, it seems EVIDENCE′ can't be my evidence after all; it's got to be something more like

EVIDENCE″: The newspaper reported that the Coalition is ahead in the polls.

We can keep doing this for a long time. If we are to stick to our Bayesian guns, it seems that we must treat all rational defeat as rebutting defeat, in which case it seems that this process of undercutting purported evidence propositions must eventually end with some special set of propositions that a) are the only propositions on which ideally rational agents ever conditionalise, and b) are immune to rational undercutting.

Some philosophers[6] think that's right; popular candidates for such special propositions include things in the vicinity of 'my visual field has such-and-such arrangement of colours'. But we Quineans believe there is no natural place to stop this process; that *all* evidence is potentially undercut-able, including the example just given. So we Quineans have a problem, if we want to be Bayesians: how do we reconcile the fact that all evidence is potentially undercut-able with the commitment to un-undercut-able evidence propositions the Bayesian apparatus appears to force on us?[7]

[6]The view is part of a package traditionally associated with Descartes. White (2014) is a contemporary endorsement of the view, and I think is representative of many contemporary epistemologists.

[7]A natural suggestion is to abandon the standard Bayesian update rule in favour of something like Jeffrey conditionalisation. A feature of standard conditionalisation is that all rationally learned propositions have credence 1, and no further evidence can make a rational agent lower her credence in a learned proposition. That can look like the source of the above problem. But in fact, it is not, as Weisberg (2009) points out. There is a lot to say about why, exactly — too much

1.3 Conceptual Change vs. Change of Opinion

In section 1.2, I said that Quineans are committed to all propositions being rationally undercuttable (or, at least, all propositions that can be rationally learned, and used as evidence, being rationally undercuttable). Many anti-Quineans are willing to concede that there is a *sense* in which everything is rationally undercuttable: all *sentences* an agent is willing to assert she may later become unwilling to assert (call this *revising a sentence*), and this need not involve any irrationality. But, say the anti-Quineans, we must distinguish between two kinds of revision: revision because an agent changes her opinion with regard to matters of fact, and revision because of *conceptual change*, where this means something like changing the meaning you associate with a given sentence in your idiolect. The anti-Quinean concedes that all sentences are rationally revisable via conceptual change, but maintains that many are immune to rational revision of the first, matter-of-fact kind.[8]

Quine, presumably, was not too impressed by this response, as he doubted that one could, really, make sense of the distinction between 'changing opinions concerning matter of fact' and 'changing a concept'. And I think we latter-day Quineans should be similarly sceptical that, in general, this distinction amounts to a real difference. But if David Chalmers is right,[9] we can't both be sceptical about this distinction and be Bayesians, as the Bayesian apparatus provides all the resources you need to make the distinction between revision via conceptual change and revision via changing you opinion on matters of fact.[10]

---

to say here. I will just report the pertinent details. Two things are true about Undercutter: a) if the agent learns it out of the blue, without learning Evidence, then it has no impact on the agent's credence in either Hypothesis or Evidence , and b) updating on Undercutter after updating on Evidence cancels the effect of updating on Evidence on Hypothesis. There is no way to make those things both true if the relevant update is Jeffrey Conditionalisation, or anything that looks remotely like a normal Bayesian update rule.

[8]This is, famously, the response to Quine (1951) given in Grice & Strawson (1956).

[9]See Chalmers (2011).

[10]You might, in fact, think that it is extremely easy to make the opinion-revision/conceptual-revision distinction using the Bayesian framework, just because the Bayesian framework presupposes that propositions are the objects of credences, rather than sentences. The idea is as follows:

What Chalmers points out is that, on the standard Bayesian assumption that rational agents update by conditionalisation, and the pretty trivial assumption that having a credence in a sentence is just having a credence in the sentence the proposition expresses, the following version of conditionalisation for sentences follows:

(CS): If an agent is fully rational, and what she learns between t_1 and t_2 is specified by sentence r, *and the content of* s *does not change between* t_1 *and* t_2, then for all s, $C_2(s) = C_1(s|r)$.[11]

The important difference between conditionalising on sentences and conditionalising on propositions is that there is only one way an agent can fail to conditionalise when she changes her credence in a proposition: be irrational. But there are *two* ways an agent can fail to conditionalise when she changes her

---

an agent 'believes a sentence' (or 'is committed to a sentence', or 'is willing to assert a sentence') when she has high credence in the proposition expressed by that sentence in her idiolect. An agent will have 'revised a sentence', in the relevant sense, when she goes from having a high credence in the proposition expressed by that sentence to a low credence in the proposition expressed by that sentence. Revising a sentence because of conceptual change is just revising a sentence because the proposition it expresses in one's idiolect changes. Revising a sentence because of change in opinion concerning the facts is keeping the proposition expressed by the sentence fixed, but lowering one's credence in that proposition.

I think that is too quick. A Quinean might take on all of this Bayesian machinery, and be happy to model agents using it, but insist that the distinction between revising because of conceptual change and revising because of change in opinion is an artefact of the model. The Quinean could maintain that the very same agent could be modelled in either way.

On the other hand, the criterion Chalmers suggests is not an artefact of the model, as long as it's not an artefact of the model who the rational agents are, what it is they learn, and what sentences they are willing to assert. One way of thinking about the view I will go on to recommend in section 2 is as a view on which who the rational agents are is, in part, an artefact of how they are modelled, rather than a fact that is entirely independent of the model. You could also think of my view as saying that what, exactly, an agent learns at any one time is in part an artefact of the model. Don't worry too much about that for the moment; it will all become clearer in section 2.

[11] '(CS)' stands for 'Conditionalisation for Sentences'. This formulation of (CS) would have to be made more complicated if we were to take into account the context-sensitivity of sentences (you can see the problem by considering sentences that include indexicals like 'now'). A way out of this, which I will take as sufficient for the purposes of this paper, is to limit (CS) to something like what Quine called 'eternal sentences' — sentences with all the indexical elements explicitly articulated, so they have constant truth-values across contexts (within a world).

credence in a statement, according to (CS): she can be irrational, or the content of the relevant sentence (in the agent's idiolect) can change. This second way of failing to conditionalise on a sentence is what Chalmers identifies as conceptual change.

(CS), then, gives us a sufficient condition for belief revision involving conceptual change: if, after learning r, a rational agent's $C_2(s)$ does not equal her $C_1(s|r)$, then her belief revision involved conceptual change. And if we're Bayesians, we're committed to (CS), it seems. So how can we Bayesian Quineans reconcile our commitment to (CS), and the sufficient condition for conceptual change it provides, with our scepticism concerning the conceptual-change/change-in-opinion distinction?

2 The Solution

2.1 Minkowski Epistemology

I think there is one solution to all these problems, for someone who wants to be both a Quinean and Bayesian. It involves denying that the distinction between rational and irrational belief-revisions is absolute; it is, instead, fundamentally relativistic, or contextual.

Now that sounds interesting, but what does it mean?

Suppose, back in the beginning of the twentieth century, someone told you there was no absolute distinction between events that are simultaneous with a given event and events that are not; that, instead, simultaneity was fundamentally relativistic, or contextual. That sounds interesting, but what does it mean? A fully satisfying answer to that question will tell us the following:

- what the non-relative, absolute phenomenon is;
- how one combines the absolute phenomenon with some contextually supplied parameter to yield a verdict about which events are simultaneous; and

- what it is about a context that supplies this parameter.

We *have* a fully satisfying account of the relativity of simultaneity. The absolute thing is how space-time is, as modelled by a Minkowski space-time. The contextually supplied parameter is a choice of orthogonal space-like and time-like axes, which you can combine with a Minkowski model to yield a verdict about simultaneity. The fact about context which supplies the choice of axes is (roughly) the state of motion of the thing we are interested in.

What the Bayesian Quinean wants, I say, is to give a similarly satisfying account of what relative rationality is. She wants a way to characterise the non-relative phenomenon, she wants to show how to combine that non-relative thing with a contextually supplied parameter and get verdicts about which revisions are rational and which aren't (verdicts which involve Bayesian machinery, to maintain the 'Bayesian' part of her title), and she wants a story about how context supplies that parameter. In short, she wants to do Minkowski epistemology.

In the remainder of this section, I am going to show one way of doing Minkowksi epistemology. I am confident it is not the best way. But I think it isn't a bad start, and have some hope that it could provide the basis of a more complete, more fully satisfying account of relative rationality. I am not wedded to the details of the following account; I think the important thing is to show that such an account is possible, and to give you some confidence that the idea of relative rationality makes sense.

2.2 The Absolute Phenomenon

What is non-relative, on the account being pursued here, is what the agent takes to be possible, at a given time. It is these apparent possibilities that are the objects of credence, on my account. Also non-relative: the agent's credences in those possibilities apparent to her, at a given time.

What is an apparent possibility? It is *not* a special kind of possibility; it may not be a possibility at all. It is the most common thing in the world for an agent

to think that something is possible when it isn't really. I think maybe you can checkmate me within four moves; actually you can't. I think I could have existed in world without Darth Vader; actually I couldn't have, for I am Luke Skywalker, and he is my father. I think it is possible that both a) Newtonian Mechanics is true and b) that bowling ball, rolling down the lane, is going to knock over all the pins; actually those things could not both be true. In each case something is apparently a possibility, but what is apparent is not a real possibility.[12]

So what *is* an apparent possibility, when it's not a real possibility? And what is it to have a credence in one of these things? I have an account. I develop it in more detail elsewhere. I have some of the more pertinent details in an appendix; below I will give you the gist of those details. I include an account in this paper because I think it is important to see that these ideas are not hopelessly mysterious. But note that my particular account of apparent possibilities and the agent's who believe in them is separable from the rest of what I say; if you have a preferred account of these things, feel free to substitute it here. The one thing you need to know, if you want to skip my account below, is that I end up identifying apparent possibilities with sets of 'notional worlds', as I call them — things that, like apparent possibilities, may not really be worlds at all, but the agent takes them to be worlds anyhow.

The account I will give takes *real* possibilities, and the idea of having a credence that one obtains, as primitive; merely apparent possibilities, and credences in them, are explained in terms of real possibilities and credences in *them*. In particular, it involves identifying a credence in a merely apparent possibility with a kind of fragmented belief in many incompatible real possibilities. Let me explain.

As I said, my account involves identifying an apparent possibility with a set of *notional worlds*, as I call them. Formally, a notional world is just a tuple of real worlds.[13] To have a credence in an apparent possibility is to have credences

[12]As I understand Kripke (1980), my use of 'real possibility' in precisely the same as Kripke's use of 'metaphysical possibility', and my use of 'apparent possibility' is the precisely the same as Kripke's use of 'epistemic possibility'. I avoid that terminology only because the waters have been so muddied since Kripke (1980) that those phrases can easily be misunderstood.

[13]Precedent for using tuples of worlds to represent something like a world can be found in

in the notional worlds that make it up.

To have a credence in a notional world is to believe, perhaps mistakenly, that there is some possibility that all the real worlds in that notional world obtain at the same time. When you *are* mistaken, and those worlds cannot all be actual (because they are not all identical), your belief-state is to be explained in terms of *fragmentation*.

Fragmentation is a certain approach to modelling imperfectly rational agents. The idea of fragmentation models is that one can treat even imperfect agents as rational locally, in confined contexts; the irrationality only manifests on a global level. I spell out an agent's credence in a merely notional world as a kind of fragmented state, in which the agent has credences in many real worlds that do not match up in the way they should.

2.3 The Parameter

2.3.1 Special Worlds

Notional worlds are a bit like what Chalmers calls 'epistemically possible scenarios'.[14] To put things in my terms, an epistemically possible scenario is a notional world that would seem like it might really be a possible world to an ideally rational agent who had no empirical information — i.e., it is a notional world that an ideally rational agent cannot rule at as really possible *a priori*.[15] This characterisation of epistemically possible scenarios presupposes that the same apparent possibilities are apparent to all ideally rational agents who lack *a posteriori* knowledge. That is, it presupposes that certain notional worlds are very special: they are *the* notional worlds that characterise the *a priori* position of ideally rational agents. (Let 'special worlds' henceforth be shorthand for 'worlds that characterize the *a priori* position of ideally rational agents'.)

Rescher & Brandom (1980).

[14] In many places; notably Chalmers (2006) and Chalmers (2012).

[15] This isn't quite right, because, strictly speaking, epistemic scenarios are more like centred worlds than worlds. But the reasons for that are not important here.

If there were such special notional worlds, they could do a lot of theoretical work for us. For instance, we could use them to factor the modelled agent's epistemic imperfection into two kinds: her *ignorance* and her *irrationality* — at least a certain kind of irrationality, to do with inconsistency. In so far as an agent had failed to rule out notional worlds in the special set (other than the one corresponding to the actual world), she would be ignorant of what the world is like; her remedy would be empirical, *a posteriori* information. In so far as an agent had failed to rule out notional worlds outside the special set, she would be irrational, as she would have failed to notice certain inconsistencies in her beliefs; her remedy would be more *a priori*, rational insight (maybe achieved through more time to sit and think, or more processing power, or something).

This would make an important difference to how we interpreted what an agent was doing when she learned. Suppose an agent learned something which allowed her to rule out special notional worlds. That would count as learning empirical information. If the agent updated her belief state on such information, and ended up going from a high credence to a low one in a certain proposition, that would count as revising a belief because of a change in opinion concerning matters of fact.

On the other hand, suppose an agent learned something that enabled her to rule out some non-special notional worlds. That would count as learning via rational insight. A perfectly rational agent would never learn in that way, of course; it's all old news, for the perfectly rational agent. But the rest of us would. And sometimes that would mean we would go from a high credence to a low one in a certain proposition on the basis of rational insight. That would not be revising a belief because of a change in opinion concerning a matter of fact, exactly. It is something more like changing your opinion concerning what follows from the matters of fact you already believe. Is that conceptual change? Unclear. (CS) doesn't settle this question for us, as it only concerns 'fully rational' agents, who would never update in this way. In any case, it is more like conceptual change than change in opinion concerning a matter of fact.

A learned proposition will determine a set of notional worlds: those worlds that are consistent with the proposition, for all that a fully rational agent can tell. It follows just from the definition of special worlds that all such worlds will be special, if there really are special worlds. The information a learned proposition carries *just is* that the actual world is among those special worlds with which it is consistent. Rational update involves having 0 credence in all worlds which one can rationally know to not be actual, given the proposition learned; that trivially includes all non-special worlds.

That means we can use our set of special worlds to help distinguish the fully rational belief revisions from the rest. Rational belief revisions *never* leave non-special worlds with non-zero credence. It is sufficient for your belief revision to be less than fully rational that you end up with non-zero credence in a non-special world.

Finally, as you may have already guessed, we can use special worlds to distinguish *a priori* from *a posteriori* propositions. *A priori* propositions are just the propositions that are true at all special worlds; *a posteriori* propositions are the ones that are true at only some special worlds.

2.3.2 No Worlds are Special Absolutely

The Quinean Bayesian position is that there are no notional worlds that are special absolutely. The very idea of the *a priori* position of the ideally rational agent is misguided, according to us Quinean Bayesians. You can idealise agents in lots of ways, but there is no special, natural way to idealise *all* agents such that they all converge on the same epistemic position, characterisable with the same set of notional worlds.[16]

That being so, there is no absolute, context-free way to factor an agent's epistemic imperfection into her ignorance and her irrationality, and there is no absolute way to distinguish learning empirical information and learning by ratio-

[16]For some examples of agents who it is natural to idealise in ways that do not converge, see section 2.4.

nal insight. What is true instead is that context picks out some set of worlds as special. And *relative to that contextually supplied set of worlds*, one can make all of those distinctions.

This also means that there is no absolute way to construe the information carried by a learned proposition, as there is no absolute fact of the matter as to which worlds one can know by rationality alone to be inconsistent with a learned proposition. The information carried by a proposition will itself be relative to a choice of special worlds.

[Aside: This does not mean we cannot identify learned propositions across agents: facts about what *real* worlds are consistent with a given proposition are absolute, and you can use those facts to individuate propositions. It is just that there is no absolute fact about how much insight rationality alone gives you into the absolute consistency facts.]

Relatedly, if there are no absolutely special worlds, there will be no absolute fact of the matter concerning which belief revisions are rational and which are not; that too will be relative to a choice of special worlds.

Finally, if there are no absolutely special worlds, then there is no absolute distinction among propositions into the *a priori* ones and the the *a posteriori* ones. This is why the no-special-worlds position is Quinean.

2.4 How Context Supplies the Parameter

This is a question on which I have no settled view. There are many options, corresponding to the analogous options philosophers have pursued in their study of contextualism about knowledge. Here are two options I find attractive (and that may, in fact, be compatible).

First option: it is facts about the way in which the relevant agent *encodes* information that determines which set of worlds to count as special. On this view, rationality has to do with the *vehicles* of representation, rather than their content. Two different agents that encode the same content with very different vehicles may be subject to very different norms of rationality.

For example, you can imagine agents who are hard-wired in such a way that if they know the velocity with which a ball is thrown through the air on their home planet, then all they need to do to figure out the path it will take is sit and think for a while. For all I know, human beings are, in fact, such agents. For such an agent, it makes sense to count notional worlds in which the ball is thrown with such-and-such a velocity through the air on their home planet and does *not* take the correct path as *non*-special, and to count the agent as *ipso facto* irrational if she has non-zero credence in such a world. But you can also imagine agents for whom there is no way to figure out the path of a ball with such-and-such velocity on the very same planet without further information; perhaps, unlike the first kind of agent, they have no tacit, hard-wired knowledge of what gravity is like on that planet, and need to discover that empirically. It makes sense to count the notional world that is not special for the first agent as special for the second one.

That's an example in which a real — i.e., metaphysical — possibility fails to be special for one kind of agent but is special for another. You can also imagine examples in which it would make sense to count a merely notional world as special for some agents but not others. For instance, you can imagine an agent that encodes the information about temperature in such a way that she could, if she sat and thought about it, work out that 30 degrees Celsius is 86 degrees Fahrenheit — the equation for conversion is implicit in the architecture of her representational apparatus, somehow. For such an agent, it would make sense to count merely notional worlds in which something was 30 degrees Celsius but not 86 degrees Fahrenheit as non-special, and to count the agent as irrational if she has non-zero credence in such a world. On the other hand, you can well imagine an agent who encodes information about the the temperature in such a way that it comes as an empirical discovery about the world that 30 degrees Celsius is 86 degrees Fahrenheit — perhaps the the two measures are associated with two very different ways of measuring the temperature, and it is only by discovering that the different ways of measuring temperature line up in the way they do that

the agent is able to determine that 30 degrees Celsius is 86 degrees Fahrenheit. For such an agent, it makes sense to count a merely notional world in which something was 30 degrees Celsius but not 86 degrees Fahrenheit as special, and to not necessarily count her as irrational if she has non-zero credence in such a world.

Though I like this option, I think in any realistic case it will not deliver everything you might have hope for from a theory of rationality. I suspect that, whatever your story of how you go from the way an agent encodes information to the set of worlds that are special for her, any realistic case will leave lots of indeterminacy as to exactly which set of worlds is special. Perhaps this indeterminacy is resolved by some additional features of the context; all the features that traditional contextualism about knowledge points to as making certain possibilities 'live' or 'salient' could be relevant.

That's option one. Here is option two: it is not the context of the *agent* that is important, but *our* context as theorists modelling the agent. In particular, the right set of worlds to choose as special depends on what we theorists are interested in. Maybe we are interested in recommending a remedy for the agent's epistemic imperfections — should she do some sitting and thinking, or should she do some experiments? If that is what we are interested in, then we should choose as the special set of worlds those notional worlds our particular agent needs to go and do experiments to be able to rule out (that is, we should choose as special those worlds recommended by the previous option). Maybe we are interested in predicting the behaviour of the agent, in which case we should choose as special notional worlds that are live possibilities to the agent, in some sense of 'live' relevant to decision making. Maybe we are interested in something else. In short: the choice of which worlds to count as special is a practical choice, to be guided by considerations of usefulness, rather than an epistemic choice, to be guided by considerations of truth.

2.5 Quinean Bayesianism: an Alternative Interpretation

I have been expounding this Quinean Bayesian view as if, on the assumption that it is correct, the right way to construe rationality is a relativistic way. That is more or less what I think. But things are slightly complicated here. What I *really* think is that, on the Quinean Bayesian theory, there are a couple of different candidate translations of the absolutist's terms 'rationality', and prior usage does not settle exactly which one of these must be adopted. The right translation is not a fact we need to discover; it is a decision we need to make, to be guided by considerations of theoretical utility.[17]

Now, I think that the way of translating 'rationality' I have been assuming — the relativistic way — is the most natural and useful, given the Quinean Bayesian view. It uses the term for a notion that can do a lot of the theoretical work we want our notion of rationality to do. But there is another reasonably natural translation available: it is something like "rational (in the relativistic sense) given the *real* (i.e., metaphysically possible) notional worlds as the choice of special worlds". This way of using the word 'rational' is actually quite intuitive. It is natural to think of the *real* (rather than merely apparent) entailment relationships that hold between propositions as those that hold when you restrict attention to the real worlds, and it is natural to think of *real* (rather than merely apparent) rationality as rationality given the real entailment relations.

Perhaps that *is* what we usually mean by 'rationality'. People certainly thought so until the 70s. But, as we learned in the 70s, if this is what 'rational' means, then being rational has very little to do with the belief-states of agents, even under idealisation of their cognitive capacities. I take that to be one of the lessons Saul Kripke taught us in *Naming and Necessity*.[18] There just isn't any reason to think that all agents, idealised in whatever way you like, should be able to work out that water is H_2O just by thinking it through. Yet failing to believe that water is

[17] Cf. Field (1973) on the relationship between 'mass', as it appears in Newtonian theory, and Relativity.

[18] Kripke (1980)

H_2O is a failure of rationality, in this sense of 'rationality'.

I think of my view, articulated using this translation of 'rationality', as in the spirit of things Timothy Williamson says. In particular, I think of this views as in keeping with things Williamson says about knowing what is rational not, in general, being any easier than knowing any other true thing.[19]

I think there is some insight to be had into the Quinean-Bayesian view by keeping this alternative interpretation in mind. On the alternative way of thinking about it, when an agent rules out merely apparent notional worlds, she is learning what constitutes rationality, in much the same way that, when she rules out a real world, she learns what the world is like. That sounds exactly like something a Quinean should say.

3 How the Solution Solves the Problem

Conceptual Change vs. Change of Opinion The Chalmers strategy for distinguishing between conceptual change and change of opinion relied on being able to distinguish the rational belief revisions from other kinds of revisions. The sufficient condition for conceptual change was revision of a sentence that was both rational and did not follow the rule for rational revision of a proposition.

Rational belief-revisions bring credence to non-special worlds to 0, among other things. So if there is no absolutely special worlds, as the Quinean Bayesian holds, then there is no absolutely rational belief-revision. You *can* make sense of the distinction between the rational belief revisions and the rest, relative to a context. So you can use the criterion Chalmers suggests to distinguish between conceptual change and change of opinion, relative to a context. But there will be no context-invariant fact concerning which belief-updates are the rational belief-updates. Hence there will be no context-invariant fact concerning which revisions of sentences were rational, but not conditionalisation.

[19] Williamson says things like this in many places, but Williamson (2000) is probably the most definitive.

Hence the Bayesian apparatus does *not* provide the resources to make an absolute conceptual change/change of opinion distinction, when interpreted the relativistic way. I think this is a vindication of Quinean scepticism concerning that distinction.

Defeat We Quineans believe there is no natural place to stop the process of undercutting; that *all* evidence is potentially undercut-able. We can give a nice account of how this works using the Quinean Bayesian apparatus developed above.

On the account above, the informational content of a learned proposition is not intrinsic to the proposition. It is, instead, a matter of both the proposition and the contextual parameter — i.e. the choice of special worlds. This is because what you can use a learned proposition to rule out depends on the choice of special worlds, as all non-special worlds are rule-out-able on the basis of *any* evidence. This means that the very same learned proposition might be construed as Evidence or Evidence′ or Evidence″, depending on which choice of special worlds is the correct one, in the given context. Are there no special worlds in which this particular poll says that the majority of Australians intend to vote for the Coalition but they don't? Then the evidence is Evidence. Are some such worlds special, but there are no special worlds where the newspaper reports that the Coalition is ahead in the polls and they aren't? Then the evidence is Evidence′. Are some of *those* worlds special, but there are no special world that are weird in some other way? Then the evidence is Evidence″. And so forth.

For all evidence to be potentially undercut-able, then, is for there to exist a context in which updating on the conjunction of the content that evidence has *in that context* with some undercutter leaves you with a lower credence than updating on the content of the evidence *in that context* alone. This is the Quinean Bayesian claim.

Now, while I think I have said enough to show how to reconcile Quinean and Bayesian commitments on defeat by undercutting, what I have said is the

merest beginning of a proper theory of defeat by undercutting. For a start, a proper theory would specify how context determines a set of special worlds. We would have to be much more committal about that than I was in section 2.4 to get a full theory. More importantly, a very difficult issue on which my account is completely silent becomes particularly salient in the context of a theory of defeat; namely, the dynamics of apparent possibility. Situations in which defeat by undercutting occur are very naturally thought of as situations in which the possibilities apparent to an agent change. Giving an account of how that is supposed to work, and how that interacts with the kind of context change I have am concerned with in this paper, is a problem that strikes me as very hard. In any case, it is a problem on which I currently have very little to contribute.

So what I present here is more of a promissory note than a theory. But I like this promissory note.

Frontloading One effect of conditionalising is to bring credence in non-special worlds to 0. If Quinean Bayesianism is true, then whether a certain belief update counts as conditionalising or not, and hence, whether a certain belief-update counts as rational or not, is not an absolute matter, but is context relative, as what worlds are special is context relative.

The Frontloading Principle will be true relative to a given context, as long as one understands the appearances of 'justified' in the principle in the appropriate, context-relative way. There will be a context-relative sense of *a priori* that goes with this context relative Frontloading Principle, and a context-relative version of the Frontloading argument will be sound. How interesting it will be for a credence to be *a priori* justified in this sense will depends on exactly how context makes the relevant worlds special, and hence decides which updates are rational updates. The very best case for the fan of the *a priori* is that it is a relatively deep fact about the kind of agent A is that makes the relevant worlds special; maybe such worlds are special for all human agents, for instance, because of the way our brains work. As I said before, I am highly doubtful that there is any plausible

story of how to go from the way our brains work to a set of special worlds that makes the choice of such worlds fully determinate. But even supposing I'm wrong, then the conclusion, while *pretty* interesting, is not as interesting as the epistemologist might have hoped. Exactly which credences are going to be the *a priori* justified ones, in this sense, is going to be a highly contingent on the details of human psychology, in exactly the way Frege had hoped to purge from epistemic theorising. I gather this is *not* a welcome outcome for most fans of the *a priori*.

On the other hand, if Quinean Bayesianism is true, there is no context invariant version of conditionalisation, and hence no context-invariant Frontloading Principle that is true, and hence no context-invariant version of the Frontloading Argument that is sound. The sense of 'justified' such an invariant argument presupposes doesn't exist, according to the Quinean Bayesian, and the premises are false (or at least not true) by way of presupposition failure.

Appendix: Credence in Apparent Possibilities

The account I will give takes real possibilities, and the idea of having a credence that one obtains, as primitive; merely apparent possibilities, and credence in them, are explained in terms of of real possibilities and credences in *them*.

We start with the 'fragmentation' approach to modelling imperfectly rational agents. The idea of the fragmentation model is that one can treat even imperfect agents as rational locally, in confined contexts; the irrationality only manifests on a global level. An agent's belief state is not modelled as a single credence function, but as many credence functions; each represents a certain *fragment* of the agent's not-completely-integrated belief-state. Each credence function is by itself probabilistically coherent, but they can fail to cohere with each other.

It will be helpful for what follows to think of each of the agent's fragments as a credence function over its own copy of the set of propositions. We can index each credence function $C_1, C_2, \ldots$, and give each copy of the set of propositions

a corresponding index — so C_i is the credence function over P_i, for all i. I'll use 'p_i' to denote the copy of p in P_i.

In this model, propositions are identified with sets of possible worlds — i.e. *really* possible worlds. So each P_i can be thought of as a power set of the set of possible worlds.

Now, I will tell you how to transform this model, which appeals only to real possibilities, into a model that contains objects that play the role of apparent possibilities. Suppose the agent has n fragments.[20] Let $\mathbf{W^n}$ be the Cartesian product of n copies of the set of possible worlds. Call the elements on $\mathbf{W^n}$ *notional worlds*. Just as an agent's belief-state can be represented as a set of n probabilistically coherent credence functions over n copes of (power sets of) real worlds, it can be also be represented as one probabilistically coherent credence function over (the power set) of one set of notional worlds, each of which is an n-tuple of real worlds.

Let $\mathbf{p_i^n}$ be the set of all notional worlds $\langle w_1, w_2, \ldots, w_n\rangle$ such that $w_i \in p_i$. You can think of $\mathbf{p_i^n}$ as another way of representing p_i. For given $C_1, C_2, \ldots C_n$, there is always a probabilistically coherent credence function $\mathbf{C^n}$ over the power set of notional worlds such that, for all p_i, $\mathbf{C^n}(\mathbf{p_i^n}) = C_i(p_i)$. For instance, there is always the $\mathbf{C^n}$ such that $\mathbf{C^n}(\langle w_1, w_2, \ldots, w_n\rangle) = C_1(w_1) \times C_2(w_2) \times \ldots \times C_n(w_n)$. So you can represent any given fragmented agent as having a coherent credence function over notional worlds.

In fact, there is always more than one $\mathbf{C^n}$ which obeys the above constraint, so $\mathbf{C^n}$ contains more information than $C_1, C_2, \ldots C_n$. The extra information tells you something about how the different fragments *interact*. The exact form of $\mathbf{C^n}$ tells you how updating in one fragment constrains updates that occur in other fragments.

The natural interpretation of a notional world (that an agent has non-zero credence in) is as an apparent possibility (apparent to the agent in question, that

[20]The letter 'n' suggests that the agent has only finitely many fragments, but I'm pretty sure everything I say carries over to trans-finite cardinalities.

is) of a certain kind. An agent who has non-zero credence in $\langle w_1, w_2, \ldots w_n \rangle$ takes it to be a live possibility that all of $w_1, w_2, \ldots w_n$ are true at the same time. If there are any two w_i in that notional world that are not identical, the agent is wrong; there is no real possibility that corresponds to the apparent possibility that notional world represents. If all the w_i *are* identical, the agent is right: the apparent possibility is, indeed, a possibility.

To summarise: on my account, a merely apparent possibility is a set of notional worlds, where a notional world is a tuple of real worlds. To have a credence in an apparent possibility is to have credences in the notional worlds that make it up. To have a credence in a notional world is to believe, perhaps mistakenly, that there is some possibility that all the worlds in that notional world obtain at the same time. When you *are* mistaken, having a credence in a notional world involves being fragmented in a certain way.

References

Chalmers, David (2012). *Constructing the World*. Oxford University Press.

Chalmers, David J. (2006). "The Foundations of Two-Dimensional Semantics." In *Two-Dimensional Semantics: Foundations and Applications*, Oxford University Press, pp. 55–140.

Chalmers, David J. (2011). "Revisability and Conceptual Change in "Two Dogmas of Empiricism"." *Journal of Philosophy*, 108(8): pp. 387–415.

Field, Hartry (1973). "Theory Change and the Indeterminacy of Reference." *Journal of Philosophy*, 70(14): pp. 462–481.

Greco, Daniel (????).

Grice, H. P., and P. F. Strawson (1956). "In Defense of a Dogma." *Philosophical Review*, 65(2): pp. 141–158.

Hawthorne, John (2002). "Deeply Contingent a Priori Knowledge." *Philosophy and Phenomenological Research*, 65(2): pp. 247–269.

Kripke, Saul A. (1980). *Naming and Necessity*. Harvard University Press.

Pryor, James (????). "Uncertainty and Undermining."

Pryor, James (2000). "The Skeptic and the Dogmatist." *Noûs*, 34(4): pp. 517–549.

Quine, Willard V. O. (1951). "Two Dogmas of Empiricism." *Philosophical Review*, 60(1): pp. 20–43.

Rescher, N., and R. Brandom (1980). *The Logic of Inconsistency*, vol. 31. Basil Blackwell.

Schiffer, Stephen (2004). "Skepticism and the Vagaries of Justified Belief." *Philosophical Studies*, 119(1-2): pp. 161–184.

Weatherson, Brian (2007). "The Bayesian and the Dogmatist." *Proceedings of the Aristotelian Society*, 107(1pt2): pp. 169–185.

Weisberg, Jonathan (2009). "Commutativity or Holism? A Dilemma for Conditionalizers." *British Journal for the Philosophy of Science*, 60(4): pp. 793–812.

White, Roger (2006). "Problems for Dogmatism." *Philosophical Studies*, 131(3): pp. 525–57.

White, Roger (2014). "What is My Evidence that Here is a Hand?" In D. Dodd, and E. Zardini (eds.) *Scepticism and Perceptual Justification*, Oxford University Press, pp. 298–321.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.

Wright, Crispin (2002). "(Anti-)Sceptics Simple and Subtle." *Philosophy and Phenomenological Research*, 65(2): pp. 330–348.