

Revisability and Conceptual Change

David J. Chalmers

1 Introduction

W.V. Quine's article "Two Dogmas of Empiricism" is one of the most influential works in 20th-century philosophy. The article is cast most explicitly as an argument against logical empiricists such as Carnap, arguing against the analytic/synthetic distinction that they appeal to along with their verificationism. But the article has been read much more broadly as an attack on the notion of the a priori, and on the program of conceptual analysis.

I will address Quine's article construed as a critique of the notions of analyticity and apriority. I will not try to give a positive account of the analytic or the a priori, beyond the standard definition of analyticity as truth in virtue of meaning and of apriority as knowability with justification independent of experience. I am much more inclined to defend the notion of apriority than the notion of analyticity, so I will focus more on the former, but the response that I will develop can be used to defend either notion from Quine's arguments.

I will focus especially on the most influential part of Quine's article: the arguments in the final section concerning revisability and conceptual change. In addressing these arguments, I will adopt a line of response grounded in Carnap's underappreciated article "Meaning and Synonymy in Natural Language". I will argue that an analysis inspired by this article, when conjoined with tools drawn from contemporary two-dimensional semantics and from Bayesian confirmation theory, provides just what is need to reject Quine's argument.

2 The Arguments of "Two Dogmas"

In sections 1 through 4 of "Two Dogmas", Quine argues that if one tries to make sense of the notion of analyticity, one ends up moving in a circle through cognate notions (synonymy, definition, semantic rules, meaning), and one cannot break out of the circle. Most philosophers have

not been greatly moved by this worry, as it seems that one finds a similar circle for all sorts of philosophically important notions: consciousness, causation, freedom, value, existence. So I will set these criticisms aside here.

In section 5 of the article, Quine makes points that are addressed specifically at Carnap's logical empiricism, criticizing his construction of physical concepts from phenomenal concepts in the *Aufbau*, and his verification theory of meaning. I will set these points aside here, as I am not concerned to defend Carnap's construction or the verification theory of meaning.

The extraordinary influence of Quine's article can be traced in large part to the short final section of the article. Part of this influence stems from the positive picture that Quine offers in the first paragraph of the section, characterizing the totality of our knowledge as a "man-made fabric which impinges on experience only along the edges", in which "no particular experiences are linked with any particular statements in the interior of the field, except indirectly through considerations of equilibrium, affecting the field as a whole". This picture serves as an powerful alternative to the verificationist picture provided by some logical empiricists.¹ It does not contain any direct argument against the analytic/synthetic distinction or related notions of apriority, however.

The most influential arguments against an analytic/synthetic distinction are found in the second paragraph, which I quote here in full:

"If this view is right, it is misleading to speak of the empirical content of an individual statement—especially if it be a statement at all remote from the experiential periphery of the field. Furthermore it becomes folly to seek a boundary between synthetic statements, which hold contingently on experience, and analytic statements which hold come what may. Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?"

I will focus on these critical arguments here. There are two crucial points.

(Q1) "Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system."

(Q2) “No statement is immune to revision.”

Many have taken these points to suggest either that no sentences are analytic, or that no distinction can be drawn between analytic and synthetic sentences. One relevant idea is that analytic sentences are those that can be held true come what may, and are likewise those that are immune to revision. If so, (Q1) suggests that by the first criterion, all sentences will count as analytic. And (Q2) suggests that by the second criterion, no sentence will count as analytic. Either way, there is no useful distinction between analytic and synthetic sentences to be had. Similarly, if we assume that a priori sentences are those that can be held true come what may, or those that are immune to revision, (Q1) and (Q2) suggest that there is no useful distinction between a priori and a posteriori sentences to be had.

One common way of responding to the argument from (Q2) is to suggest that revisability is quite compatible with apriority (or analyticity), on the grounds that a priori justification (or the justification we have for believing analytic sentences) is *defeasible*. For example, I might know a mathematical claim a priori, but my justification might be defeated if I learn that a leading mathematician thinks that the claim is false. I think that this response is correct as far as it goes, but it concedes a great deal to Quine. On a common traditional conception, at least some a priori justification (and some justification for believing analytic truths) is *indefeasible*. One might reasonably hold that some a priori justification (in logic or mathematics, say) yields not just knowledge but certainty, at least on ideal reflection. These claims are not obviously correct, but they are also not obviously incorrect, and I do not think that Quine’s argument establishes that they are false. So I will take another line of response.

The response I will develop takes off from the response given by Grice and Strawson at the end of their article “In Defense of a Dogma”. This response holds that (Q1) and (Q2) are compatible with an analytic/synthetic distinction, for a reason quite different from the one given above. Here is a passage addressing the argument from (Q2):

Now for the doctrine that there is no statement which is in principle immune from revision, no statement which might not be given up in the face of experience. Acceptance of this doctrine is quite consistent with adherence to the distinction between analytic and synthetic statements. Only, the adherent of this distinction must also insist on another; on the distinction between that kind of giving up which consists in merely admitting falsity, and that kind of giving up which involves changing or dropping a concept or set of concepts. Any form of words at one time held to express

something true may, no doubt, at another time, come to be held to express something false. But it is not only philosophers who would distinguish between the case where this happens as the result of a change of opinion solely as to matters of fact, and the case where this happens at least partly as a result of a shift in the sense of the words. Where such a shift in the sense of the words is a necessary condition of the change in truth-value, then the adherent of the distinction will say that the form of words in question changes from expressing an analytic statement to expressing a synthetic statement. ... And if we can make sense of this idea, then we can perfectly well preserve the distinction between the analytic and the synthetic, while conceding to Quine the revisability-in-principle of everything we say.

Here the central point is that our judgments about any *sentence*, even an analytic sentence, will be revisable if the meaning of the words change. For example, if 'bachelor' changes from a term for unmarried men to a term for sociable men, then we will no longer judge that 'All bachelors are unmarried' is true. But this observation is just what an adherent of the analytic/synthetic distinction should expect. Analytic sentences should instead be understood as those sentences that are immune to revision *while their meaning stays constant*.

Following standard practice, we can say that when the meaning of a sentence changes, there is *conceptual change*: some expression in the sentence at first expresses one concept and later expresses another. When the meaning of a sentence stays the same, there is *conceptual constancy*: the expressions in the sentence will express the same concepts throughout. Then Grice and Strawson's point could be put by saying that an analytic sentence is one that is immune to revision *without conceptual change*. More cautiously, the point could be put by saying that the fact that a sentence is revisable under conditions of conceptual change does not entail that it is not analytic. Something similar applies to apriority.²

At this point, Quine has two obvious replies. The first reply is to say that the appeal to meaning in characterizing the class of analytic sentences is circular, as the notion of meaning as poorly understood as the notion of analyticity. The same could be said for the appeal to concepts and to

²We might also allow that there is conceptual change in this sense when the proposition expressed by an utterance of a sentence changes because of a shift in context. For example, 'Someone is bald iff they have no hairs' might be accepted in one context and rejected in another. It is not clear that a mere contextual shift could change the status of a sentence as analytic, as arguably the meaning of such a sentence stays constant throughout. But if we say that a sentence is a priori if it expresses a proposition that is knowable a priori, then it is natural to hold that in these conditions a sentence might be a priori in one context but not in another.

propositions. This reply would be in the spirit of the first four sections of “Two Dogmas”. But then this argument will not be much of an advance on the arguments in the first four sections, and anyone who is not moved by those arguments will not be moved by this one.

The second, more interesting reply is to challenge Grice and Strawson to provide a *principled distinction* between cases of revision that involve conceptual change and those that involve conceptual constancy. Quine might argue that cases that are purported to be on either side of this division are in fact continuous with each other, and that there is no principled distinction to be had. Something like this thought might even be read into the last sentences of the paragraph from Quine quoted above.

Now one might suggest that Grice and Strawson are not obliged to provide a *reductive* characterization of the distinction—that is, one that does not use ‘meaning’ and cognate notions—any more than they are required to provide a reductive definition of meaning or analyticity to answer the challenge in the first four sections. Again, this suggestion seems correct as far as it goes. But nevertheless, if Quine’s opponent cannot say much to characterize the principled distinction here, he or she is at least in the awkward dialectical position of leaving a challenge unanswered, and of leaving doubts about the distinction unassuaged.

My own view is that there is much that can be said to flesh out a principled distinction here. I think that the tools for doing so can be found in Carnap’s “Meaning and Synonymy in Natural Language”.

3 Carnap on intensions

Carnap is of course Quine’s major target in “Two Dogmas of Empiricism”. It is not always appreciated that “Meaning and Synonymy in Natural Language” can be read as a sustained response to Quine, perhaps because Carnap spends little time discussing Quine in the article. Nevertheless, Carnap says enough to make clear that a response to “Two Dogmas” is intended.

Carnap’s article sets out to provide an analysis of the notion of meaning and of related notions such as synonymy. His aim is to provide a “scientific procedure” by which meaning and synonymy can be analyzed in broadly naturalistic terms. Importantly, he aims to explicate not only the notion of extension, but the notion of intension (the “cognitive or designative component of meaning”), which he notes has been criticized by Quine as “foggy, mysterious, and not really understandable”.

Carnap’s key idea is that we can investigate the intension that a subject associated with an expression by investigating the subject’s *judgments about possible cases*. To determine the intension

of an expression such as 'Pferd' for a subject, we present the subject with descriptions of various logically possible cases, and we ask the subject whether he or she is willing to apply the term 'Pferd' to objects specified in these cases. If we do this for enough cases, then we can test all sorts of hypotheses about the intension of the expression.

In this article Carnap takes the term 'intension' as a primitive, and does not build possible cases into the very nature of intensions. But for our purposes it is useful to adopt a suggestion that Carnap makes elsewhere, and simply define an intension as a function from possible cases to extensions. For a term like 'Pferd', the intension will be a function from possible cases to objects characterized in those cases. For a sentence such as 'Grass is green', the intension will be a function from possible cases to truth-values. Then Carnap's procedure above can be regarded as a way of directly ascertaining the values of the intension that a subject associates with an expression, by presenting the subject with a possible case and noting the extension that the subject associates with the case.

Of course one cannot actually present a subject with all possible cases to determine every aspect of an intension. But Carnap suggests that the intension that a speaker associates with an expression is determined by the speaker's *linguistic dispositions*. For a given expression E used by a given speaker, the speaker will have the disposition to associate a given extension with E , when presented with a possible case. For example, if S is a sentence, the speaker will have the disposition to judge the sentence as true or false of a possible case, when presented with that case. The intension of an expression can then be seen as a function that maps possible cases to the extension that the speaker is disposed to identify, when presented with that case.

In this way, Carnap defines an expression's intension in naturalistic and even operational terms. We can go on to define synonymy: two expressions are synonymous (for a speaker at a time) when they have the same intension (for that speaker at that time). And we can define analyticity: a sentence is analytic (for a speaker at a time) when its intension has the value "true" at all possible cases (for that speaker at that time).

With this definition in hand, we can go on to provide a principled criterion for conceptual change over time. An expression E undergoes change in meaning between t_1 and t_2 for a speaker iff the speaker's intension for E at t_1 differs from the speaker's intension for E at t_2 . If we accept Carnap's dispositional account of intensions, it follows that E undergoes change in meaning between t_1 and t_2 iff there is a possible case such that the speaker is disposed to associate different extensions for E when presented with the case at t_1 and t_2 .

Of course there are many immediate questions about Carnap's account. What is a possible

case? In what vocabulary are these cases specified? How can we determine whether the meaning of this vocabulary has changed? Cannot speakers make mistakes about intensions? Cannot they change their mind about a case without a change in meaning? Can meaning really be operationalized this easily? And so on. Carnap's account may need to be modified or at least refined to answer these questions.

Before addressing these matters, though, I will illustrate how Carnap's account might be used to address the challenge in section 6 of "Two Dogmas" directly. In my view, the essential aspects, if not the specifics, of the resulting response are sound. These essential aspects can then be carried over to more refined analyses couched in terms of two-dimensional semantics and Bayesian confirmation theory.

4 A Carnapian response

In "Meaning and Synonymy in Natural Language", Carnap does not mention the arguments in section 6 of "Two Dogmas", and he does not address considerations about revisibility or conceptual change at all. Nevertheless, his framework can be used to give a response to these arguments that is broadly in the spirit of Grice and Strawson's response, fleshed out with a principled criterion for conceptual change.

We can start with Quine's observation that any statement can be held true come what may. This seems correct. Even a paradigmatic synthetic sentence such 'All bachelors are untidy' can be held true in the face of apparently countervailing evidence, if we allow sufficient adjustment of ancillary claims. The question is whether such adjustments will involve conceptual change, and whether we have a principled criterion for determining this.

We might as well start with a case. At t_1 , Fred asserts 'All bachelors are untidy'. At t_2 , Fred is presented with evidence of a tidy unmarried man. Fred responds: 'He's no bachelor! Bachelors must be over 30, and he's only 25'. At t_3 , Fred is presented with evidence of a 35-year-old with a spotless apartment. Fred responds: 'He's not tidy! Look at the mess in his sock drawer.' In this way, Fred holds the sentence true throughout, and through similar maneuvers he may hold it true come what may.

Does this case involve conceptual change? We can apply Carnap's analysis to see whether the Fred's intension for 'All bachelors are untidy' changes over the relevant timespan. Suppose that C is a detailed possible case in which there is an unmarried 25-year old with a tidy apartment. At t_2 , when Fred is presented with the information that C obtains, he responds that 'All bachelors are

untidy' is true with respect to C. By Carnap's criterion, Fred's intension for B is true with respect to C at t_2 .

What about Fred's intension for B at t_1 ? The key question is: if Fred had been presented with a description of C at t_1 , before he had evidence that the case was actual, would he have judged that 'All bachelors are untidy' was true with respect to C?

If the answer is yes, then Carnap's criterion suggests that there is no relevant conceptual change between t_1 and t_2 . In this case, Fred will simply have had an unusual intension for 'bachelor' all along.

If the answer is no, then Carnap's criterion suggests that there is relevant conceptual change between t_1 and t_2 . The intension of 'All bachelors are untidy' will have changed during this time, probably because the intension of 'bachelor' has changed during this time.

The same applies more generally. If a speaker's judgment concerning a case at t_2 is reflected in the speaker's dispositions to respond to such a case at t_1 , we can say that the speaker's judgment concerning that case is *prefigured*. If a speaker's judgment concerning a case at t_2 is not reflected in the speaker's dispositions at t_1 , we can say that the speaker's judgments concerning the case is *postfigured*. On Carnap's account, postfigured judgments but not prefigured judgments involve conceptual change.

In any case, we have what is needed. Carnap's framework allows us to see how any sentence can be held true come what may, while at the same time allowing a principled way to distinguish between those cases of holding true that involve conceptual change and those that do not. Something similar applies to cases of revisability, though I will not go through the details here.

5 Refining Carnap's account

Carnap's account of meaning is remarkably simple, and one might reasonably wonder whether such a simple account can be correct. I think that while there are problems with the account, the problems can be addressed in a way that preserves something of the spirit of the account, if not the letter.

Perhaps the most obvious problem is that subjects can make mistakes. A subject might miscalculate and judge that $36+27=73$, and they might even be disposed to judge this to be true with respect to all possible scenario. On Carnap's account, it will follow that ' $36+27=73$ ' is analytic for the subject. But this seems the wrong result: on the face of it, the sentence is not even true. Similar mistakes seem possible for non-ideal subjects in all sorts of domains.

To handle cases of this sort, we can modify the account to appeal not to what the subject *would* say in response to the case, but to what the subject *should say*, or what they would say given ideal reasoning. We might say that the intension of *E* maps a possible case *C* to the extension that the subject would identify for *E*, if they were to be presented with *C* and were to reason ideally. An effect of this change will be that the account will no longer yield an operational definition of meaning, at least unless we can find an operational criterion for ideal reasoning. But this is not a bad thing for those who are inclined to reject behaviorism in any case.

Another problem is that on the contemporary understanding, intensions are often not accessible to a subject, even by ideal reasoning. For example, if Kripke (1972) is right, the intension of ‘water’ picks out H₂O in all possible worlds, even for subjects who do not know that water is H₂O. Such subjects will not be disposed to identify H₂O as the extension of ‘water’ when presented with a possible case, so Carnap’s definition will get the intension wrong.

To handle this problem, we can take a leaf from two-dimensional semantics, which recognizes two sorts of intension. Even in light of Kripke, Carnap’s account might still apply to one sort of intension, although not the other. Kripke’s point applies to secondary intensions, which govern possible cases considered as counterfactual: if there *had been* XYZ in the oceans and lakes, water would still have been H₂O. For the purposes of Carnap’s account, though, we can focus on primary intensions, and stipulate that subjects consider the possible cases as actual. For example, we can ask them to suppose that XYZ *is actually* in the oceans and lakes in the actual world, and ask them for their verdict about the extension of ‘water’ under that supposition. Subjects will plausibly hold that ‘water’ picks out XYZ if that hypothesis is correct. This mirrors the familiar suggestion that the primary intension of ‘water’ picks out XYZ in a Twin Earth scenario, although the secondary intension of ‘water’ picks out H₂O there. So it is not out of the question that a Carnap-style account might work for primary intensions, which are often held to be the sort of intensions that are most closely tied to apriority and analyticity in any case.

A third issue is the nature of possible cases. For our purposes they should be something akin to possible worlds. They might be centered metaphysically possible worlds (worlds marked with an individual and a time), with the centering required to handle intensions for expressions such as ‘I’ and ‘now’. They might also be regarded as epistemically possible worlds, or epistemically possible scenarios, which might be modeled by maximal consistent sets of sentences that cannot be ruled out a priori. I will not try to settle this issue here, but I will use the word “scenario” as a generic term for the entities involved. Of course a large idealization is required in order to suppose that subjects can reason about entire scenarios, but here as before we can appeal to the idealization

of what the subject would say given ideal reasoning.

A fourth issue is the vocabulary in which a scenario is specified. Such a vocabulary will need to be rich enough that a full enough specification using this vocabulary plus ideal reasoning determines judgments about other expressions' extensions, without being so rich that a specification builds in all those expressions directly. It is a substantive claim that some such vocabulary can be found, but proponents of two-dimensional semantics have offered arguments for this claim along with suggested vocabularies (see, for example, Chalmers and Jackson 2001). The details of such a vocabulary will not matter for our purposes here.

A fifth issue is the worry that subjects might change their mind about a possible case without a change of meaning. Here, one can respond by requiring, as above, that the specifications of a scenario are rich enough that judgments about the scenario are determined by its specification and by ideal reasoning. If so, then if the subject is given such a specification and is reasoning ideally throughout, then there will not be room for them to change their mind in this way. Changes of mind about a fully specified scenario will always involve either a failure of ideal reasoning or a change in meaning. Of course this claim requires a version of the substantive claim in the previous paragraph. I will return to a version of this issue later.

The model we then reach is something like the following. The (primary) intension of an expression for a subject is a function that maps scenarios to extensions, mapping a scenario w to what the subject would judge to be the extension of E under the supposition that w is actual, were they ideally rational. This is not a perfect definition, but it is good enough for our purposes. This remains very much in the spirit of Carnap's definition, although the invocation of rationality makes it a sort of normative version of Carnap's account.

Importantly, we can use this account to give a version of the Carnapian response to Quine's arguments given in the previous section. Conceptual change (of the relevant sort) will occur precisely when an expression's primary intension changes across time. This will happen precisely when the subject's dispositions to judge the expression's extension in a possible case (given ideal reasoning) changes. As in the last section, we can find cases of holding-true where the dispositions change in this way, and cases where they do not. What matters is that we have a principled distinction.

A residual issue concerns the meaning of the basic vocabulary. If cases are specified in this vocabulary, then we need to ensure that the basic vocabulary does not change in meaning throughout the process. If we do not require this, the resulting condition for meaning change will be inadequate: a subject's dispositions to judge that S obtains with respect to a case specified by D

might change over time, not because the meaning of S changes but because the meaning of terms in D change. If we require this, however, then it appears that we need some further criterion for meaning change in the basic vocabulary items used in D , as the dispositional method would yield trivial results here. So it appears that the dispositional method for determining meaning change, even when idealized, is incomplete.³

A second residual issue concerns the role of the a priori in characterizing this account. It is natural to suggest that the ideal reasoning in question must be restricted to ideal *a priori* reasoning. In fact, some two-dimensional accounts (Chalmers 2004) use the notion of apriority in defining primary intensions: the primary intension of a sentence S evaluated at a world w is true precisely if a material conditional ‘If D , then S ’ is a priori, where D is a canonical specification of S . If so, then we have arrived at a principled distinction only by helping ourselves to the contested notion of apriority along the way.

As before, it is not clear how bad these residual problems are. One might still see the intensional analysis as demonstrating that the Quinean phenomena of holding-true and revisability are quite compatible with the intensional framework and have no power to refute it. Even if one has to assume some independent grip on the notion of apriority, and on the meaning of expressions in the basic vocabulary, one can still use the framework to provide a reasonably enlightening analysis of relevant cases. Still, we have not broken out of the Quinean circle. It would be nice to be able to characterize the relevant distinctions without such a direct appeal to the contested notions.

I think that such a characterization can be found. The key idea is to cast things in terms of conditional probability, rather than in terms of apriority.

An initial observation is that something very much like a primary intension can be characterized without appealing to apriority, by appealing to conditional probability instead. In particular, one can define the intension of a sentence S at a scenario w , for a subject, in terms of the subject’s rational conditional credence $cr'(S|D)$, where D is a canonical specification of w . We can say that the intension of S is true at w iff $cr'(S|D)$ is high, and false at w iff $cr'(S|D)$ is low. Here we require an idealization, so that $cr'(S|D)$ is the conditional credence that the subject would have

³This objection is related to Quine’s argument from the indeterminacy of translation in *Word and Object*. Quine took Carnap’s account to be a serious challenge to his arguments in “Two Dogmas”, and the indeterminacy argument can be seen in part as a response to it. Here, Quine argues that no dispositional analysis can settle facts about meaning, because multiple assignments of reference will always be compatible with a subject’s behavioral dispositions. This applies even to Carnap’s account, if we allow multiple potential assignments of reference to the basic vocabulary. In effect, Carnap’s account assumes that the meaning of the basic vocabulary is fixed, but it is not clear why such an assumption is legitimate, and it is not clear how this meaning might itself be grounded in dispositional facts.

given ideal reasoning, or something along those lines.

If we do this, then we will have a principled criterion for conceptual change that does not appeal to apriority. On this criterion, a subject's intension for S will change between t_1 and t_2 iff there is a scenario w with canonical specification D such that $cr'(S|D)$ changes from high to low or vice versa. One could then run the arguments of the previous section once again using this notion. This will provide a reply to Quine's challenge that gets around the second residual issue above (regarding apriority), though it may still be subject to a version of the first issue (regarding the basic vocabulary).

At this point, however, I think an alternative analysis involving conditional probability is available. This analysis is closely related to the one just mentioned, and is a descendant of the Carnapian analysis in the previous section, but it does not require the apparatus of possible cases and intensions. Instead it can proceed using only standard Bayesian considerations about evidence and updating. In addition to the advantage of familiarity, this approach has other significant advantages in responding to Quine's challenge. By avoiding the need for canonical specifications of complete possible scenarios, it avoids the large idealization needed to handle enormous specifications. It also has the potential to avoid or minimize both residual issues above.

6 A Bayesian analysis of holding-true

Let us assume a standard Bayesian model, on which sentences are associated with unconditional and conditional credences for subjects at times. That is, for a given subject and a given time, a sentence S will be associated with an unconditional credence $cr(S)$, and a pair of sentences S and T will be associated with conditional credence $cr(S|T)$. (These ordinary credences $cr(S|T)$ should be distinguished from the idealized rational credences $cr'(S|T)$ from the previous section.) Credences are standardly taken to be real numbers between 0 and 1, but for our purposes exactitude is not required. It is enough that some credences be high and others low.

I will also assume a version of the standard Bayesian principle of conditionalization: if a subject has credence $cr_1(S|E)$ at t_1 , and acquires total evidence specified by the evidence sentence E at between t_1 and t_2 , then the subject's credence $cr_2(S)$ at t_2 should be equal to $cr_1(S|E)$. I will give a more precise version of this principle below. The nature of evidence sentences will be discussed later in this article, but for now we can think of them either as specifying that certain experiences obtain, or as specifying that certain observable states of affairs obtain.⁴

⁴The arguments I present here can also be run using the principle of Jeffrey conditionalization (Jeffrey 1984), which

We can start with a typical case whereby an apparently synthetic sentence is held true in face of apparently countervailing evidence, by appeal to appropriate ancillary theses. As in section 1, suppose that at t_1 , Fred asserts ‘All bachelors are untidy’. At t_2 , Fred acquires evidence indicating that there is a tidy, unmarried 25-year old man, and responds by denying that the man is a bachelor, as bachelors must be over 30.

Let B be ‘All bachelors are untidy’, and let E be Fred’s total relevant evidence acquired between t_1 and t_2 . Let $cr_1(S)$ stand for Fred’s credence in S at t_1 , and $cr_2(S)$ stand for Fred’s credence in S at t_2 . Then $cr_1(B)$ and $cr_2(B)$ are both high.

The crucial question is: What is $cr_1(B|E)$, Fred’s conditional credence in B given E at t_1 , before Fred acquires the evidence in question?

If $cr_1(B|E)$ is high, then Fred’s judgment at t_2 reflects a conditional credence that he already had at t_1 . In this case, the judgment at t_2 is *prefigured*, in a sense analogous to the sense discussed earlier. Here, Fred’s accepting B in light of E can be seen as according with the principle of conditionalization.

If $cr_2(B|E)$ is low, then Fred’s judgment at t_2 fails to reflect the conditional credence that he already had at t_1 . In this sort of case, the judgment at t_2 is *postfigured*, in a sense analogous to the sense discussed earlier. Here, Fred’s accepting B in light of E appears to violate the principle of conditionalization.

Now, on standard Bayesian assumptions, there are two central ways in which one can obtain apparent violations of conditionalization for sentences. First, this can happen when the subject is not fully rational throughout the process: perhaps at t_1 they have not thought things through properly, or at t_2 they make some sort of reasoning error. Second, the content of the key sentence B can change between t_1 and t_2 . This may happen in cases involving indexicals, which are not relevant here, or in cases of conceptual change. In these cases, it remains possible that the subject’s credences in relevant *propositions* obey conditionalization, but that their credences in associated sentences do not, because the association between sentences and propositions changes over time.⁵

allows conditionalization on evidence of which a subject is not certain.

⁵A potential third way that conditionalization can be violated arises on views where sentences express certain sorts of relativistic contents: for example, a view on which utterances of the sentence ‘It is raining’ always express the same temporal proposition *It is raining*, which can be true at some times and not at others. On Saturday, I might have a low conditional credence in *It is raining* given *The weather forecast says rain on Sunday*, then on Sunday I might acquire evidence that the weather forecast says rain on Sunday, resulting in high credence in *It is raining*, without irrationality. On a more standard view on which the content of ‘It is raining’ uttered at t is *It is raining at t*, this will be classified as a change in content, but on the temporal view the content stays the same. For present purposes, we can either count

We might formulate this as a version of the Bayesian principle of conditionalization, for sentences:

(CS) If a subject is fully rational, and if the subject acquires total evidence specified by E between t_1 and t_2 , and if the content of sentence S does not change between t_1 and t_2 , then $cr_2(S) = cr_1(S|E)$.

Perhaps the most familiar version of the principle of conditionalization is cast in terms of propositions: if a fully rational subject acquires total evidence specified by proposition e between t_1 and t_2 , then $cr_2(p) = cr_1(p|e)$. (CS) follows from this claim in conjunction with the plausible claims that when sentence S expresses proposition p for a subject at that time, $cr(S) = cr(p)$ at that time, and that the content of a sentence is the proposition it expresses.

It follows that *if* Fred in the postfigured case above is fully rational, then this is a case of conceptual change. Of course it might be that Fred is not fully rational, but this is of no help for Quine. It is unremarkable that irrational subjects might hold on to any sentence or reject any sentence, and this observation has no consequences regarding analyticity or apriority. For Quine's observations about revisability and holding-true to have any bite, rational subjects are required. So we may as well assume that Fred is fully rational.

If we assume that the relevant subjects are fully rational, we now have a principled criterion for conceptual change in a case of holding-true. Suppose that our subjects accepts S at t_1 , acquires apparently countervailing evidence E between t_1 and t_2 , and continues to accept S at t_2 . Then we can say

- (i) If $cr_1(S|E)$ is low, this is a case of conceptual change.
- (ii) If $cr_1(S|E)$ is high, this need not be a case of conceptual change.

One can now ask: is it true that a subject can hold on to any given sentence S come what may, in light of any evidence, *without irrationality or conceptual change*? By this analysis, this claim requires that for any given sentence S and any evidence E , $cr(S|E)$ is high (or at least is not low). But this claim is obviously false. For a rational subjects and most sentences (including most paradigmatic empirical sentences), there will be evidence sentences E such that $cr(S|E)$ is low.

The moral here is that in the general case, Quinean holding-true-come-what-may requires widespread violation of conditionalization, which requires irrationality or conceptual change. But

these as changes in content in an extended sense, or we can require in principle (CS) that the content in question is non-relativistic content.

the fact that an irrational subject might reject a sentence is no evidence that it is not analytic or a priori,⁶ and the fact that a subject might reject a sentence after conceptual change is no evidence that it is not originally analytic or a priori. So Quine's argument from holding-true fails.

7 A Bayesian analysis of revisability

For our central example of revisability, we can use a familiar case of Putnam's. Let C be 'All cats are animals'. This might seem paradigmatically analytic or a priori. But let E specify evidence confirming that the furry, apparently feline creatures that inhabit our houses are actually remote-controlled robots from Mars, while the other creatures that we see are all organic. Putnam argues that if we discovered that E obtains, we would reject C . So let us suppose that Sarah accepts C at t_1 , acquires total evidence as specified by E , and rejects C at t_2 .

Here, the diagnostic question is: What is Sarah's initial conditional probability $cr_1(C|E)$?

If $cr_1(C|E)$ is low, then Sarah's judgment at t_2 reflects a conditional credence that she already had at t_1 . In this case, the judgment at t_2 is *prefigured*. Here, Sarah's accepting C in light of E can be seen as according with the principle of conditionalization.

If $cr_2(C|E)$ is high, then Sarah's judgment at t_2 fails to reflect the conditional credence that she already had at t_1 . In this sort of case, the judgment at t_2 is *postfigured*. Here, Sarah's accepting C in light of E appears to violate the principle of conditionalization.

For exactly the reasons given before, the postfigured case requires either that Sarah is not fully rational, or that her use of C undergoes conceptual change between t_1 and t_2 . Cases of this sort are of no help to Quine. Again, the fact that an irrational subject might reject a sentence is no evidence that it is not analytic or a priori, and the fact that a subject might reject a sentence after conceptual change is no evidence that it is not originally analytic or a priori.

For Quine's argument to succeed, he needs to exclude cases of this sort. That is, he needs to make the case that any sentence can in principle be rationally revised without a violation of conditionalization. This requires that for all rational subjects and for all sentences S , there exists an evidence sentence E such that $cr(S|E)$ is low.

This claim is not so obviously false as the corresponding claim about holding true come what may. For this reason, one might regard the argument from revisability as a stronger argument than

⁶Perhaps there are certain strong conceptions of analyticity on which an analytic sentence cannot be rejected by any subject, rational or irrational. But I do not think that these conceptions are standard, and in any case no such constraint applies to apriority.

the argument from holding-true. Indeed, supporters of Quine such as Putnam and Harman have concentrated on the argument from revisability, and have made claims very much like this. For example:

Still, it is not clear just what the grounds are for accepting this claim. At this point, a number of observations can be made.

First, Quine's official grounds for the revisability claim involves involves the ability to revise ancillary claims when necessary. These grounds are the same as for the holding-true claim, and it is clear that Quine sees the two as continuous. These grounds suggest that *after* obtaining evidence, a subject could use these features to revise a given sentence. But we have seen that revisions of this sort typically involve violations of conditionalization. These grounds do very little to suggest that *before* acquiring the relevant evidence, a subject's conditional credence $cr(C|E)$ will be low.

Second, it is certainly the case that almost any claim could be rationally rejected given testimony of an apparent epistemic superior. But this claim has no bearing on apriority: that a claim could be rejected in *this* way is no evidence that it is not a priori. One might also note that it is far from clear that this claim applies to ideally rational thinkers (perhaps their grounds for accepting a mathematical claim, say, will defeat the evidence concerning an apparent epistemic superior?). If not, this line of thought does not support the claim that any claim could be rejected under conditions of ideal rationality. It is also worth noting that at least in many cases (mathematical cases and the like), evidence of this sort will often be *misleading* evidence against the sentence in question. So this line of thinking does not give us reason to hold that any sentence could come to be *correctly* rejected.

Third, even if this sort of consideration applies to many apparent cases of a priori truths, there are a number against which it has no purchase. Some such cases include *material conditionals* of the form 'If D , then S ' (like those discussed in the previous section), where D is a lengthy specification of an arbitrary scenario (including a full specification of any evidence obtained in that scenario), and where S is a sentence such as 'Water is H_2O ' such that $cr(S|D)$ is high. Assuming a fully rational subject, it follows that $cr(D \supset S|D)$ is high, so that $cr(D \supset S)$ is also high. Now a quick two-case argument suggests that no evidence E could lead us to rationally reject $D \supset S$. First case: If E is not evidence specified within D , then E entails $\neg D$. In this case, $cr(\neg D|E) = 1$, so $cr(D \supset S|E) = 1$. Second case: If E is evidence specified by D , then D entails E . Now $cr(D \supset S|E)$ must lie between $cr(D \supset S|E \& \neg D)$ and $cr(D \supset S|E \& D)$. But the former is 1, and the latter is just $cr(D \supset S|D)$, which we have seen is high. So $cr(D \supset S|E)$ is high. Putting the two cases together, $cr(D \supset S|E)$ is high for all E . Importantly, material conditionals very much

like these are the a priori truths that are most important in the two-dimensional framework.

Fourth, once one notes that this argument allows some truths S such that $cr(S|E)$ is high for all E , then it is clear that there is no longer a sound principled argument that for all S , there is an E such that $cr(S|E)$ is low. As a result, we may expect to find many more exceptions to this claim. Indeed, many Quineans have conceded such objections, for example in the domains of mathematics and logic, and there is no reason not to expect many more.

Fifth, it is worth stressing that even if this line of argument succeeded, it would be much more conservative than Quine's original line. It leads naturally to a view on which there is an analytic/synthetic distinction. At worst, it would be the case that most or all sentences previously regarded as analytic (a priori), such as 'All cats are animals', will be reconstrued as synthetic (a posteriori). But one could still use conditional probabilities to characterize Carnapian intensions, once one acknowledges that the intensions for sentences such as 'All cats are animals' will be false at some scenarios. One will still have a principled distinction between cases that involve conceptual change and cases that do not. In this way, the advocate of analyticity, apriority, and conceptual analysis will have much of what they want.

In any case, the Bayesian analysis has given us what we wanted: a principled criterion for identifying cases of conceptual change. It has only given us a sufficient condition, rather than a necessary and sufficient condition, but this is good enough for our purposes. With this analysis in hand, it is clear that Quine's arguments from revisability and holding-true fail.

8 Quinean Objections

(1) *The Bayesian analysis begs the question.*

It might be suggested that the Bayesian principle (CS) that I have appealed to simply assumes a notion of conceptual change without argument, and therefore begs the question against the Quinean skeptic about this notion. I do not think that this is quite right. (CS) is itself a consequence of the principle of conditionalization for propositions and of two other weak assumptions, none of which say anything about conceptual change. Still, this line of argument assumes a notion of proposition, about which a Quinean might be skeptical.

Now, Quine's doubts about propositions has been much less influential than his doubts about the analytic/synthetic distinction. But in any case, I think it is clear that Bayesian accounts of confirmation require *either* something like propositions or something like the notion of conceptual change to get off the ground. Bayesian credences will be assigned either to abstract entities such

as propositions, events, or sets, to linguistic items such as sentences, or to mental items such as beliefs. If we take the first route, then we can use these entities just as we used propositions to ground a notion of conceptual change. And if we take the second or the third routes, we need to require something like conceptual constancy in order to avoid counterexamples to principles such as conditionalization.

Of course a Quinean might simply reject Bayesianism altogether, along with the associated principle of conditionalization. This would seem rash, however, as Bayesianism is an extremely successful theory with widespread empirical applications. So by a Quinean's own lights, it is hard to reject it. Furthermore, even if one rejects Bayesianism, a successor theory is likely to have corresponding principles of diachronic rationality, governing how beliefs should be updated over time in response to evidence. And precisely the same issues will arise for these principles: if they apply to abstract items we can use these to define conceptual change, and if they apply to linguistic items or mental items, we will require a notion of conceptual change.

I think the deeper moral is that there is a constitutive link between *rational inference* and conceptual constancy. Issues such as those floated here will arise for any principle of diachronic rationality at all. If it is a principle that from A and $A \supset B$ one should infer B , and if the premises and conclusions here are sentences or mental items, then to avoid obvious counterexamples, the principle should require that A and B have the same meaning on each occasion when they occur. And if the principle applies to abstract objects such as propositions, these can themselves be used to define conceptual change. So if we are not skeptics about principles of diachronic rationality, a notion of conceptual change will be hard to avoid.

(2) *Rationality presupposes apriority.*

It might be suggested that in appealing to the notion of rationality, the notion of apriority is tacitly smuggled in. For example, someone might hold that all principles of rational inference depend on underlying principles about the a priori: for example, perhaps an inference from some premises to a conclusion is rational precisely if it is a priori that if the premises obtain, the conclusion is likely obtain. Or perhaps the distinctive idealization made by the Bayesian involves some tacit assumptions about the a priori. For example, perhaps the Bayesian requirement that rational subjects should have credence 1 in logical truths depends in some way on the belief that logical truths are a priori. If so, the appeal to rational principles here presupposes one of the key notions at issue.

The reply here is straightforward. Whether or not the objector is correct that rationality depends in some way on apriority, the appeal to rationality is innocuous in the current dialectical

context. The relevant class of opponents are those who accept the notion of rationality, but who question the notion of apriority. My argument is intended to establish that *if* one accepts certain principles concerning rationality, then one should reject Quine's argument against the a priori. If this objector is correct, then the opponent should either give up on the principles concerning rationality, or accept the notion of the a priori. Either outcome is sufficient for my purposes. I am happy to concede that if an opponent rejects the notion of rationality, or rejects all relevant principles of diachronic rationality, then the current argument has no purchase against them.

It is also worth noting that the principles of rationality that I appeal to are principles that many or most opponents of the a priori accept. Conditionalization has no obvious connection to the a priori, for example. I do not know whether the special status that the Bayesian gives to logical truths has a special connection to the a priori, but in any case this status plays no role in my argument. That is, the argument does not require the Bayesian claim that rationality requires credence 1 in logical truths. In fact, the picture I have sketched appears to be compatible with a view on which logical truths deserve rational credence less than 1, and on which they can be revised given relevant evidence. All that is required is that such a revision should obey conditionalization. Nothing here smuggles in any obvious presuppositions about the a priori.⁷

(3) *A principled line between conceptual change and irrationality cannot be drawn.*

A Quinean may suggest that our concept of rationality is not fully determinate, and that as a result a clear division between cases of irrationality and cases of conceptual change cannot be found. Some hard cases, such as revising logic in light of quantum mechanics, are not easily classified as either.

However, my reply to Quine's argument does not require drawing a line here. It suffices for the purposes of the argument that the violations of conditionalization involve *either* irrationality or conceptual change, and we do not have to classify these violations further. In any case, as long as there are clear cases of rational judgment, the existence of unclear cases entails at worst a vague distinction, not a nonexistent distinction.

(4) *The argument requires constancy in evidence sentences.*

Recall the first residual issue for the framework of intensions discussed earlier: the framework assumes conceptual constancy in the base vocabulary, so that the framework cannot explain this conceptual constancy. One might think that an analogous issue arises here, with respect to the

⁷Of course, my argument appeals to logical claims at various points, but this does not require that logical truths are a priori, or that they are unrevisable. It merely requires that they are true. Likewise, my argument does not require that the principle of conditionalization is itself a priori or that it is unrevisable. It simply requires that the principle is true.

evidential vocabulary: the vocabulary used to specify evidence sentences such as E . After all, conditionalization concerns what to do when one has a certain credence $cr(S|E)$ and then learns E . The conditional credence is in part an attitude to a sentence E , and what one learns is also a sentence E . One might think it is required that the sentence have the same meaning on both occasions. If so, then any apparent failures of conditionalization in a rational subject could be blamed on a change in the meaning of terms in E , instead of a change in the meaning of terms in S , and it is not clear that we have a principled way to choose.

As it stands, this picture is not quite right. Learning E does not typically involve the *sentence* E at all. Perhaps if learning was always by testimony, and if E is a sentence used in testimony, then the issue arises. But for our purposes we can assume that the relevant learning is by perception. Here, E will be a sentence characterizing the evidence that one learns, and the learning process need not involve this sentence at all. So there is no use of E at t_2 that needs to be aligned with the use of E at t_1 . At best we need to require that E as used at t_1 correctly applies to the evidence acquired at t_2 . But this is a much weaker requirement, concerning only the extension of E as used at t_1 , with no role for any use of E at t_2 .

Still, it can be argued that acquiring evidence requires having certain *attitudes* to the evidence. For example, the rationality of Bayesian conditionalization on new experiences arguably requires not just that one has the experiences, but that one is certain that one has them. If so, one might suggest that the framework tacitly requires that at t_2 , one is certain of the evidence statement E (saying that certain experiences obtain). This issue is starker in alternative frameworks such as Jeffrey conditionalization, which accommodate uncertainty about evidence by giving an explicit role to one's credence in evidence statements such as E at t_2 . Does this not require some sort of constancy in the meaning of E after all?

The issue is delicate. For the reasons given above, I think that the sentence E as used at t_2 plays no essential role here. However, it is arguably the case that subjects must be certain of (or have other appropriate attitudes to) certain evidential *propositions*, such as the proposition that certain experiences obtain, which were expressed by E at t_1 . Or without invoking propositions: the subject must be certain that the relevant evidence obtains (that they are having certain experiences, say), where this is the same evidence concerning which they had conditional credences at t_1 . Without this alignment, one could always respond to an apparent failure of conditionalization by saying that although the subject's initial credence was conditional on evidence e obtaining, and although evidence e later obtained, the subject in fact became certain that some *other* evidence e^* obtains. If this were so, there would be no violation of conditionalization (the subject would not acquire

the evidence e), and there would arguably be no irrationality.

This requirement of alignment provides some room for the Quinean to maneuver, but the room is extremely limited. To eliminate this room altogether, we need only suppose that we have a grip on what it is for a subject to accept or suppose that certain evidence obtains. With this much granted, we can simply stipulate that for our purposes, the conditional credences $cr(S|E)$ relevant at t_1 are credences in S conditional on the evidence that is actually obtained at t_2 . This removes any loophole, and does so without making any assumptions about constancy in the meaning of language across time. At most, we have to assume an understanding of certain beliefs and suppositions about evidence.

The required assumptions can be made even smaller by noting that for our purposes, evidence can be limited to experiences or at least to observational states of affairs. While there is a sense in which non-observational states can serve as evidence for other claims, it is plausible that knowledge of these states of affairs is itself grounded in evidence concerning experiential or observational matters. On a Bayesian view, our credences in these states of affairs must then match those determined by conditionalization on experiential or observational matters. I think it is also plausible that credences in observational states of affairs should themselves match those determined by conditionalization on experiential matters.⁸ If the latter claim is granted, then for present purposes we can restrict the relevant evidence in cases of revisability and holding-true to experiential states. And even without it, we can restrict the relevant evidence to observational states. So to answer the Quinean worry, we need only suppose that we have a grip on what it is for a subject to accept or suppose that certain experiential or observational states of affairs obtain. And this is something that Quine's arguments in "Two Dogmas" do not give us any reason to doubt.

The upshot of all this is that the residual issues about a base vocabulary are not eliminated altogether on a Bayesian approach, but they are minimized, in a way that brings out the severe costs of the Quinean position. A Quinean who rejects the notions of analyticity and apriority along present lines must also insist that there is no objective fact of the matter about whether a subject accepts or supposes that a given observational state obtains. This view would presumably go

⁸For example, if one is fully rational, one's credence that there is a red square in front of one, should match one's antecedent conditional credence that there is a red square in front of one given that one is having an experience as of a red square. (If norms of rationality do not ensure certainty about the experiences one is having, one can move to a Jeffrey-conditionalization analog.) Theses of this sort have been denied by some dogmatists about perception (e.g. Pryor 2006), and might also be denied by some who think that perceptual knowledge is more secure than introspective knowledge (eg. Schwitzgebel 2008).

along with a generalized skepticism about the contents of thought, perhaps in the spirit of Quine's skepticism about meaning developed in his arguments concerning radical translation. It would likewise require a certain skepticism about diachronic rationality, for reasons discussed earlier.

Quine himself argues both for skepticism about meaning (in *Word and Object*) and for a sort of skepticism about norms of rationality (in "Epistemology Naturalized"). But few have been prepared to follow him here, and even those who sympathize with the Quine of "Two Dogmas" have tended to reject these later views. Of course Quine's arguments for these views deserve attention in their own right, but it is clear that the arguments in "Two Dogmas" do not provide much direct support for them. Still, the current analysis suggests a deep linkage between these views. Defending the arguments of "Two Dogmas" against a certain sort of appeal to conceptual change leads naturally to skepticism about diachronic rational principles and about the content of language and thought. Contrapositively, once even minimal claims about rationality and about thought are accepted, the arguments I have considered against analyticity and apriority dissolve.

(5) *There can be rational revision by resetting priors.*

A Quinean might suggest that (CS) is false, as there is a way that fully rational subjects might violate conditionalization without conceptual change. The relevant method here is that of *resetting priors*. This method stems from the observation that most Bayesians allow that there is some flexibility in one's ultimate priors: the prior probabilities that a subject should have before acquiring any empirical evidence. (Of course these priors are something of a fiction.) For example, on Carnap's framework for inductive logic, equally rational subjects may have different values for λ , the parameter that guides how quickly the subjects adjust their beliefs in light of inductive evidence, and this difference can be traced to a difference in ultimate priors. Two such subjects might acquire exactly the same evidence over time, while being led to quite different posterior probabilities. If G is the thesis that a certain sort of global warming is occurring, for example, one subject might be led to a high credence in G , while another might be led to a low credence in G .

Now, a subject with a high credence in G might reflect and observe that their high credence is tracable entirely to the value of λ in their ultimate priors, and that this value was quite arbitrary. They may note that it would have been equally rational to start with a lower value of λ , and to end up with a lower credence in G . At this point, a bold subject might choose to change their credences wholesale. At least if they have a good enough record of their evidence, they can "unwind" back to the ultimate priors, reset λ to a lower value, and reintegrate all the evidence by conditionalization. The subject will end up with a new set of credences, including (among many other differences) a much lower value for G .

A Quinean might suggest that there is nothing irrational about doing this, and that this method might be exploited in order that a subject can hold on to almost any sentence “come what may” and to revise almost any sentence. After all, for most non-observational empirical sentences S and most paths of evidence, there is some ultimate prior that will lead to a high credence in S , and some ultimate prior that will lead to a low credence in S . None of this requires conceptual change. So violations of conditionalization in a rational subject do not provide a sufficient condition for conceptual change, after all.

This position requires a rejection or at least a revision of orthodox Bayesianism. On the orthodox view, conditionalization is a constraint on diachronic rationality, and this sort of revision will be irrational. Furthermore, the view tends to lead to an anything-goes view of rational belief. If there are no constraints on ultimate priors, the view entails that at any moment, if $cr(p) < 1$, then one’s credence can be rationally revised so that $cr(p)$ is arbitrarily close to zero. And even if there are constraints on ultimate priors, these constraints must be weak enough to vindicate the large violations of conditionalization that the Quinean argument requires, leading naturally to a view on which most beliefs can be rationally revised at any moment into disbelief. Given this much, it is not easy to see how my beliefs can constitute knowledge at all.⁹

Furthermore, it is far from clear that all beliefs can be revised in this way. For example, given that logical beliefs, mathematical beliefs, and evidential statements are constrained to have credence 1, this method will not yield revisability for these beliefs. More generally, there is not much reason to hold that it will yield revisions to those beliefs usually classified as a priori (‘All bachelors are unmarried’, say), most of which do not appear to depend on ultimate priors. So this response is weakest where it needs to be strongest.

Most fundamentally: as long as we have a conceptual distinction between cases in which beliefs are revised by this process and cases in which they are not, we still have enough to draw a distinction between those violations of conditionalization that involve conceptual change and those that do not. The Quinean will have to insist that we do not have a grip on this conceptual distinction, so that there is no distinction to be drawn between cases of resetting priors and cases of conceptual change. I do not think there is much reason to accept this. Furthermore, even if this line were accepted, it would once again lead to an across-the-board skepticism about principles of belief updating and other forms of diachronic rationality. So if principles of diachronic rationality are allowed at all—even the liberal principles suggested by the current approach—then the

⁹In addition, this method is a sort of belief revision that is not driven by evidence at all. So the sort of revisability that one gets out of this picture is contrary to the empiricist spirit of many Quineans.

distinction between conceptual constancy and conceptual change remains intact.

(6) *Subjects need not have conditional credences.*

It might be objected that the Bayesian analysis requires the assumption that for every sentence S used by a subject and every possible evidence sentence E , the subject has a conditional credence $cr(S|E)$. But this is an unrealistic idealizing assumption.

In response: The idealization is not enormous. For most S and most E , the subject will have some relevant dispositions involving S and E , for example involving their willingness to accept various bets involving S and E . In many cases, these dispositions will line up in a clear enough way that $cr(S|E)$ will be high. In other cases, they will line up in a clear enough way that $cr(S|E)$ will be low. In other cases, the dispositions may be enough of a mix that it is hard to say.

A Quinean might suggest that if $cr(S|E)$ is indeterminate in this way, and the subject later rejects S upon learning E , this should not count as a violation of conditionalization. If so, they might then suggest that for any S , there is some E such that $cr(S|E)$ is indeterminate in this way, and such that the subject could later reject S on learning E without violating conditionalization. Perhaps this sort of revisability is enough for their purposes?¹⁰

I do not think that this is enough, however. Cases of this sort seem to turn essentially on the subject's not being fully rational. If the subject is fully rational, then the subject's dispositions to accept S on *supposing* E and on *learning* E should be the same, assuming no conceptual change. That is, if a fully rational subject rejects S on learning E and thinking things through, then if the subject were to have been initially presented with the *supposition* that E and had thought things through, the subject should have rejected S conditional on that supposition. To fail to meet this condition is a failure of full rationality, just as is an ordinary violation of conditionalization. So at best the Quinean has presented us with a kind of revisability that can only be exploited by subjects who are less than fully rational. Like the sort of revisability that can be exploited only by irrational subjects, this sort of revisability has no bearing on matters of apriority.

¹⁰It is especially likely that ordinary subjects will lack credences $cr(S|D)$ involving the scenario specifications D discussed earlier, due to the enormous size of these specifications. This observation does not affect the use of conditional credences involving D to define intensions, as these credences used there are always idealized rational credences $cr'(S|D)$, for which the current issue does not arise. And where nonidealized credences are concerned, these cases will not yield cases of revisability along the lines in the text, because the subject will be incapable of learning that D .

9 Conclusion

Quine is right that any statement can be held true come what may, and that no statement is immune to revision. But as Grice and Strawson observe, these phenomena are quite compatible with a robust analytic/synthetic distinction and a robust notion of meaning. A Bayesian analysis reveals that Quine is not right that any statement can be held true come what may *without conceptual change or irrationality*, and likewise for revision. We can pin down the distinction between cases that involving conceptual change and cases that do not using either the method of intensions or Bayesian analysis.

The method of intensions characterizes intensions in terms of certain idealized dispositions, and uses this notion to draw a distinction between cases that involve conceptual change and cases that do not. In the central version that I examined, this method assumes the notion of apriority, so it does not provide an independent grounding for that notion. Still, it shows how a framework involving apriority can accommodate all of Quine's data. And for the same reasons that most philosophers reject Quine's arguments in sections 1-4 of "Two Dogmas", no independent grounding is required.

The Bayesian analysis takes things a step further and defends the a priori on partly independent grounds. This analysis assumes the notion of conditional probability and the normative notion of rationality to provide conditions for conceptual change, but it does not assume the notion of apriority. In effect, constitutive connections between rational inference and conceptual change are used to make inroads into the Quinean circle.

The conclusion should not be too strong. While I have responded to Quine's arguments against the a priori and the analytic, I have not provided a positive argument for the analytic/synthetic distinction or the a priori/a posteriori distinctions, and I have not tried to ground these notions in wholly independent terms.

One might be tempted to take things a step further still, attempting to define apriority in terms of conditional probability and rationality. For example, one might suggest that a sentence S is a priori for a subject precisely when the ideal conditional probability $cr(S|D)$ is 1 (or: is high) for all scenario specifications D . But there will be residual issues. For a start, it is not clear that one can define the class of scenario specifications without using the notion of apriority. In addition, the thesis is subject to various potential counterexamples: for example, one might argue that when S is a mathematical truth, and D specifies a scenario in which one is a poor mathematical reasoner,

$cr(S|D)$ should be much less than 1. So much more would need to be said here.¹¹

Still, we have seen that these notions can at least help us in diagnosing issues regarding meaning, conceptual change, and the a priori. And we have seen enough to suggest that Quine's arguments in the final section of "Two Dogmas of Empiricism" do not threaten the distinction between the analytic and the synthetic, or the distinction between the a priori and the a posteriori.

References

- Carnap, R. 1955. Meaning and synonymy in natural languages. *Philosophical Studies* 6(3):33-47. Reprinted as Appendix D of *Meaning and Necessity* (University of Chicago Press, 1956).
- Chalmers, D.J. and Jackson, F. 2001. Conceptual analysis and reductive explanation. *Philosophical Review* 110:315-61.
- Chalmers, D.J. 2004. Epistemic two-dimensional semantics. *Philosophical Studies* 118:153-226.
- Jeffrey, R. 1983. *The Logic of Decision*. University of Chicago Press.
- Kripke, S. 1980. *Naming and Necessity*. Harvard University Press.
- Pryor, J. 2007. Uncertainty and undermining. <http://www.jimpryor.net/research/papers/Uncertainty.pdf>
- Quine, W.V. 1951. Two dogmas of empiricism. *Philosophical Review* 60:20-43.
- Quine, W.V. 1960. *Word and Object*. MIT Press.
- Quine, W.V. 1967. Epistemology naturalized. In *Ontological Relativity and Other Essays*. Columbia University Press.
- Schwitzgebel, E. 2008. The unreliability of naive introspection. *Philosophical Review*.

¹¹I address the second issue in forthcoming work, arguing that there is an idealization on which $cr(S|D)$ should be 1 here, and that this idealization can be understood without appeal to the a priori. I think that the role of the a priori with respect to the first issue is more robust.