Constructing the World, Lecture 4
Revisability and Conceptual Change: Carnap vs. Quine
David Chalmers

Text: http://consc.net/oxford/.
E-mail: chalmers@anu.edu.au.
Discussion meeting: Thursdays 10:45-12:45, Ryle Room.

Primary aim: Defend the a priori/a posteriori distinction from Quine's arguments in "Two Dogmas of Empiricism".

Secondary aims: Introduce intensions in the scrutability framework. Examine the interaction between meaning and diachronic principles of rationality.

1. "Two Dogmas of Empiricism"
The two dogmas are (i) belief in an analytic/synthetic distinction, and (ii) reductionism (every meaningful statement is equivalent to a statement about experience). I'll focus on Quine's most influential arguments against (i), especially construed as arguments against an a priori/a posteriori distinction.

Sections 1-4 of "Two Dogmas" argue that analyticity can be understood only via cognate notions such as meaning, definition, synonymy, etc, leading to a circle. This argument is widely rejected, as such circles are common with philosophically important notions.

Section 5 argues that sentences cannot be associated with sets of confirmatory experiences, because of the underdetermination of theory by evidence. This material is specific to the second dogma.

The most influential part of "Two Dogmas" is section 6, especially the first two paragraphs.

> "The totality of our so-called knowledge or beliefs … is a man-made fabric which impinges on experience only along the edges. … A conflict with experience at the periphery occasions readjustments in the interior of the field. Truth values have to be redistributed over some of our statements. ... But the total field is so undetermined by its boundary conditions, experience, that there is much latitude of choice as to what statements to re-evaluate in the light of any single contrary experience. 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.
>
> 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?"

The arguments of paragraph 1 are directed at an *Aufbau*-style view and the second dogma, but the arguments of paragraph 2 are directed at the analytic/synthetic distinction.

The crucial points are:

(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 take these points to suggest either that no sentences are analytic/a priori, or that no principled distinction can be drawn between those that are analytic/a priori and those that are not.

Response [Grice and Strawson]: Any *sentence* can be held true come what may, and no *sentence* is immune to revision. But this holding-true and revision will often involve changing the meaning of the sentence. If so, it has no bearing on the status of the original sentence as analytic or a priori.

Quinean response: (1) The appeal to meaning here is circular [cf. Sections 1-4]. (2) There's no *principled basis* for classifying some cases as involving conceptual change and others as involving conceptual constancy.

Project: investigate the prospects for drawing a principled distinction here, using tools drawn from (i) Carnap's "Meaning and Synonymy in Natural Language", (ii) the scrutability framework, and (iii) Bayesian confirmation theory

2. Carnap on intensions

Carnap: a "scientific procedure" for analyzing meaning and synonymy. The meaning of an expression is its *intension.* The intensions associated with a term are determined by a subject's linguistic dispositions. We investigate an intension of a term such as 'Pferd' by presenting subjects with descriptions of possible cases and asking for a judgment about what 'Pferd' applies to. We can represent intensions as functions from possible cases to extensions.

[In the scrutability framework: Given a sentence S and a possible scenario T specified using a PQTI-sentence T': subject judges (using a Cosmoscope?) whether S is true if T' is true.]

Then: A sentence is analytic if its intension is true at all possible cases. Two expressions are synonymous if they have the same intension. An expression undergoes conceptual change between t_1 and t_2 iff its intension at t_1 differs from its intension at t_2.

3. A Carnapian analysis of a Quinean case

Quine: "Any statement can be held true, come what may."

Example: Fred asserts at t_1: 'All bachelors are tidy'

Prima facie, this statement is contingent and synthetic, but Fred could hold onto it in face of any apparently countervailing evidence, by adjusting ancillary claims.

• Faced (at t_2) with a untidy unmarried man: "He's no bachelor! Bachelors must be over 30, and he's only 25."
• Faced (at t_3) with a 35-year-old with a dirty apartment: "He is tidy! Look at his well-organized sock drawer."

Question: Does the move from t_1 to t_2 [or t_3] involve conceptual change? Only if the intension for "All bachelors are untidy" changes between t_1 and t_2.

Let B be 'All bachelors are tidy'. Let C be a scenario with 25-year-old unmarried men in dirty apartments.

By Carnap's criterion, the intension of B is true at C at t_2. Is the intension of B true at C at t_1?

Diagnostic question: At t_1, does Fred judge that B is true in the possible case specified by D? E.g. at t_1, would Fred accept "If there are 25-year-old unmarried men with dirty apartments (and so on), then all bachelors are tidy?".

If "yes": then there is no change in the intension of B between t_1 and t_2 (just a nonstandard intension all along). This is a *prefigured* case: the judgment at t_2 is reflected in the speaker's dispositions to make conditional judgments regarding possible cases at t_1.

If "no" [more plausibly]: then the intension of B changes between t_1 and t_2. This is a *postfigured* case: the judgment at t_2 is not reflected in the speaker's dispositions at t_1.

Carnapian criterion: postfigured judgments entail conceptual change.

4. Refining Carnap's Account

Problems and refinements:

  (i)   Subjects can make mistakes. Response: define intensions in terms of idealized judgments (what the subject would judge on ideal reflection).

  (ii)   Intensions are a posteriori (Kripke). Response: define intensions by considering scenarios as epistemic possibilities (cf. scrutability, primary intensions).

  (iii) What are possible cases and in what vocabulary are they specified? Response: Maximal specifications in a generalized scrutability base. [Requires a version of the generalized scrutability thesis.]

Then: The (primary) intension of an expression E for a subject is a function that maps scenarios to extensions, mapping a scenario w to what the subject ideally should judge to be the extension of E under the supposition that D is actually true, where D is a canonical specification of w (in the language of a generalized scrutability base).

Residual issues: (i) The account assumes meaning and conceptual constancy for basic vocabulary, so no account of conceptual change for that vocabulary (cf. indeterminacy of translation). (ii) Huge idealization. (iii) Do we have to use the a priori in characterizing this account? [E.g., requiring that 'If D then S' be a priori.]

If the answer to (iii) is yes: then one has analyzed the Quinean phenomena of holding true and revisability on non-Quinean terms, showing that they are compatible with substantive notions of analyticity and apriority, but one hasn't broken out of the Quinean circle. If no: then we've gone further to break out of the Quinean circle, grounding a notion of conceptual change without assuming these notions.

Defining intensions without apriority, using conditional scrutability: the (primary) intension of S for a subject is true at a scenario w iff the idealized rational conditional probability $p^*(S|D)$ for the subject is high, where D is a canonical specification of w.

5. A Bayesian analysis of holding-true

Alternative: A Bayesian analysis in terms of conditional probabilities, hypotheses and evidence (rather than intensions, scenarios and apriority) throughout. This can help with issues (i) and (ii) as well as (iii).

Assume a Bayesian framework on which sentences S and T are associated with unconditional credences $p(S)$, $p(T)$ and conditional credences $p(S \mid T)$ at times. Also assume the principle of conditionalization: if a subject has credence $p_1(S \mid E)$ at t1 , and acquires total evidence specified by the evidence sentence E at between t_1 and t_2 , then the subject's credence $p_2(S)$ at t_2 should be equal to $p_1(S \mid E)$.

Let E be the total relevant evidence that Fred acquires between t_1 and t_2: e.g., that there is a 25-year-old unmarried male with such-and-such living situation.

Question: what is Fred's conditional probability $p_1(B \mid E)$ at t_1?

In the "prefigured" version of the case, $p_1(B \mid E)$ is high. Then Fred's accepting B in light of E can be seen as standard updating of belief by conditionalization.

In the "postfigured" version of the case, $p(B \mid E)$ is low. Then Fred acquires E as total relevant evidence, but still accepts B. This is a violation of conditionalization.

Such violations can occur when:
        (i) E isn't the total relevant evidence – but we can stipulate that it is.
        (ii) The subject isn't fully rational at stage 1 or 2 – but we can stipulate rationality.
        (iii) The meaning or content of B changes between stages – bingo.

Conditionalization for Sentences: 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 $p_2(S) = p_1(S \mid E)$.

So: given a sentence S that is rationally held true "come what may", i.e. in light of potentially conflicting evidence E:

        (1) If $p(S \mid E)$ is initially low, this will be a case of conceptual change.
        (2) If $p(S \mid E)$ is initially high, this need not be a case of conceptual change.

This gives us some independent grip on the distinction between cases involving conceptual change and those that do not.

To establish that every sentence can be held true come what may *without conceptual change or irrationality*, Quine would need to argue that for all sentences S and all potential evidence E, $p(S \mid E)$ is high (or at least, is not low). But this is obviously false.

This suggests that it is not true that any sentence can be held true come what may, without conceptual change or irrationality.

Underlying principle: there is a constitutive connection between *rational inference* and conceptual constancy. If [A, A→B, therefore B] is a principle of rational inference, then anyone who violates it (diachronically, for sentences A and B) without change in the meaning of A or B is irrational. Likewise, anyone who rationally violates it is engaged in conceptual change.

6. A Bayesian analysis of revisability

Example: C = 'All cats are animals' [Putnam].

This might seem paradigmatically analytic/a priori. But let E specify evidence confirming that 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.

Diagnostic question: What is our initial conditional probability $p_1(C|E)$?

If $p_1(C \mid E)$ is low [Putnam gives us reason to accept E→C *now*]: then this is a prefigured judgment, compatible with conditionalization. In this case, C is not analytic/a priori to start with (at least not in a sense requiring the possibility of certainty on ideal reflection).

If $p_1(C \mid E)$ is high, but we reject C upon obtaining total relevant evidence E: this is a postfigured judgment that violates conditionalization. So this is a case of conceptual change or irrationality.

To maintain [within a Bayesian framework] that any statement is open to revision *without conceptual change or irrationality*, Quine needs the claim that for any sentence S, there is some possible evidence E such that a subject's rational conditional probability $p_1(S \mid E)$ is low. This claim is not as obviously false as the analogous claim about holding true come what may, but it is not clear what the grounds are for accepting it.

> (i) Quine's official support involves underdetermination, ancillary claims, etc – but this sort of revision often involves violations of conditionalization, so it does not support the claim about conditional probability.
>
> (ii) Almost any claim could be rationally rejected given testimony of an apparent epistemic superior. But that a claim can be rejected in this way is no evidence that it is not a priori. Also, this does not obviously yield ideally rational revisability.
>
> (iii) This strategy gets no purchase against material conditionals of the form 'If D, then S', where D specifies a full scenario complete with total evidence.
>
> (iv) This strategy still allows us to use these conditionals to define intensions which can serve as a relevant sort of meaning, with a corresponding analytic/synthetic distinction. At worst, it follows that few statements are analytic.

In any case, conditionalization again gives us a grip on the distinction between revisions that involve conceptual change vs those that do not.

7. Objections

(1) The Bayesian analysis begs the question
(2) Rationality presupposes apriority
(3) A principled line between conceptual change and irrationality cannot be drawn.
(4) The argument requires constancy in evidence sentences.
(5) There can be rational revision by resetting priors.
(6) Subjects need not always have rational conditional credences.

8. Conclusion

Quine is right that any statement can be held true come what may, and that no statement is immune to revision. But these phenomena are quite compatible with a robust analytic/synthetic distinction and a robust notion of meaning. 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 involving conceptual change or irrationality using either conditional conceptual analysis or Bayesian analysis.

We can see this response to Quine on two levels:

(i) Defending the a priori on its own ground. If rely on a priori scrutability, we presuppose a notion of apriority in characterizing the conditionals, and argue that such a framework can accommodate all Quine's data. This does not provide an independent grounding for the notion of the a priori (although it delimits its grounding role). But for the same reasons that most philosophers reject Quine's arguments in sections 1-4 of "Two Dogmas", no such independent grounding is required.

(ii) Defending the a priori on partly independent grounds. If we use a Bayesian analysis, we need only assume a notion of conditional probability and of rationality. This assumes normative notions, but does not obviously assume the notion of apriority, so gives some independent purchase on the cases. In effect, constitutive connections between rational inference and conceptual change are used to make inroads into the Quinean circle.

We should not make this claim too strong. We have not grounded the notion of apriority in wholly independent terms. One might be tempted to define a (conclusively) a priori statement as a statement S for which the ideal conditional probability $p(S \mid D) = 1$ for all scenario specifications D. But there will be residual issues:

>   (i) Can one define the class of scenario specifications without using the notion of apriority?

>   (ii) Can we deal with potential exceptions to the thesis (e.g. due to scenarios involving misleading evidence, cognitive deficit, etc)?

>   (iii) Can the notion of ideal conditional probability be understood in a way wholly independent of the a priori?

Still, even without a reductive account of apriority, we have enough of an antecedent grasp on the relevant notions that these notions provide at least an illuminating tool for analysis. And our grip on principles of conditional rational inference helps us diagnose cases of conceptual change.

All this suggests that Quine's arguments from revisability and from holding-true do not threaten the analytic/synthetic distinction or the a priori/a posteriori distinction, and that the scrutability framework might be used to vindicate a broadly Carnapian approach to meaning.