

Review of David Chalmers's *Constructing the World*

Louis deRosset

August 20, 2014

David Chalmers writes big books. *Constructing the World* is no exception. The book is organized into eight substantial chapters and seventeen excursuses. The chapters pursue the main line of inquiry, and the excursuses respond to some objections and offer interesting developments of some of Chalmers's proposals. The book is a sustained exploration of the case for and philosophical applications of *scrutability theses*, which say that all truths can be determined (or "sried", to use Chalmers's preferred term (30)) from a relatively austere class of truths. Different scrutability theses are obtained by specifying, *e.g.*, the sorts of truths that figure in the austere basis and the form the scrying takes. The scrutability theses that play the most important role in the discussion concern scrutability of all truths from the union *PQTI* of:

P: The class of physical truths, including macrophysical truths,¹ microphysical truths, and laws;

Q: The class of truths concerning what it's like to be in this or that conscious state, together with the psychophysical laws that connect consciousness with physical states;

I: A class of indexical truths, most prominently including truths which specify where one is and which time is now; and

T: A 'that's all!' truth saying that *P*, *Q*, and *I* are exhaustive, in a sense to be explicated.

¹The inclusion of macrophysical truths in *P* is explicit in the book (110, 261), but marks a departure from some of Chalmers's earlier defences of scrutability theses; see esp. [Chalmers and Jackson, 2001].

The first two chapters and the beginning of the third are devoted to introducing scrutability theses and isolating the one that figures most centrally:

A Priori Scrutability For all truths M , the material conditional ($PQTI \Rightarrow M$) is a priori.²

The fifth and sixth chapters are devoted to countering objections to A Priori Scrutability, and the seventh and eighth and related excurses to considering applications of A Priori Scrutability and certain strengthenings of it. The applications include the explication of Fregean notions of semantic and mental content, some limited arguments against external world skepticism, explicating the relation between the metaphysically fundamental and the derivative, vindicating the unity of science, and enabling the dissolution of some metaphysical disputes through the epistemic analysis of language. These explorations are interesting and enlightening, but they all depend on A Priori Scrutability and its strengthenings, so the heart of the book is the third and fourth chapters, in which Chalmers makes the case for A Priori Scrutability.

Imagine, urges Chalmers, that you have the use of a *cosmoscope*, a machine with four features. First, it can show you a map at any scale of the distribution of matter and energy in any region of (actual) spacetime.³ Second, if you select an actual subject at an actual time, then the cosmoscope can induce in you an imaginative experience with the same phenomenal character as whatever experience that subject may be undergoing at that time. Third, the cosmoscope can calculate, given a specification of a physical setup, what the counterfactual consequences of the laws would be were that setup to be realized. Finally, the cosmoscope can specify where you are and what time it is now. If I had a cosmoscope, it seems that in theory I could figure out whether the weather here will be hotter tomorrow than it is today. Assume the standard identification of temperature in fluids with mean molecular kinetic energy (MMKE).⁴ Since I know that temperature = MMKE, I could direct the cosmoscope to compare

²Read ' $(PQTI \Rightarrow M)$ ' as indicating a material conditional whose antecedent is the (perhaps infinitary) conjunction of the truths in $PQTI$; I will take similar liberties with the use-mention distinction below.

³Actual spacetime does not contain a cosmoscope, so Chalmers is asking us to imagine a fictional situation; see (114-5). Some modifications of the cosmoscope thought experiment may have to be made to accommodate the presentation of the microphysical truths of quantum mechanics, which would not take the form of a specification of the distribution of mass and energy in spacetime (115, 293-8).

⁴For the sake of brevity, in what follows I will write 'temperature = MMKE' even though, strictly speaking, it is only temperature in fluids which can be identified with MMKE.

tomorrow's MMKE in this region to today's. If I happen not to know that temperature = MMKE, then I could discover what temperature is by, as it were, looking over the shoulders of the scientists as they did the experiments whose results enabled them to know that temperature = MMKE. I could either follow their reasoning myself or use the cosmoscope's computational features to check their reasoning for me. They knew that temperature = MMKE, so, it seems, the cosmoscope puts me in a position to know that temperature = MMKE. Or rather, since the situation in which I have the use of a cosmoscope may be physically very different from the actual situation, I am in a position to know that *if* the world is as the cosmoscope represents it to be, then the weather here will be (as it turns out) hotter tomorrow.

The cosmoscope contains and presents the information in *PQI* and provides some computational support for reasoning from that class of truths. It thus illustrates the fact that a suitably idealized version of me is in a position to know many conditionals of the form 'if *PQI*, then *M*', where *M* is a truth not in *PQI*. But, Chalmers notes, the cosmoscope does not put me in a position to infer every such truth from *PQI*. For instance it does not put me in a position to infer that there are no ghosts; the cosmoscope gives me no information one way or the other about the existence of ghosts (151), since ghosts, we may assume, are neither physical nor phenomenal entities. In general, a cosmoscope will not allow a normal user to infer every *negative truth* from *PQI*. A *positive truth* is a truth which cannot conceivably be rendered false by adding something to a world in which it is true (111, 153); a truth is *negative* just in case it is not positive, and the notions of *positive* and *negative claims* are defined similarly, except that there is no requirement that they be true. For example, 'there are no ghosts' can be rendered false by adding ghosts to a world, and so is clearly negative.

So, Chalmers points out, *PQI* does not provide us with enough information to scry all of the negative truths. We will need to appeal to *T*, the 'that's all!' truth, to secure a fully general scrutability thesis. But, in what sense is *PQI* all? *PQI* is not *syntactically* complete, since there are truths such that neither they nor their negations are in *PQI*. The case of the non-existence of ghosts also shows that *PQI* is not *logically* complete, *analytically* complete, nor *modally* complete on plausible assumptions, given the most straightforward ways of making sense of those ideas. Chalmers critically discusses a number of

proposals (151-5), but the version he uses is

T All positive truths are entailed a priori by *PQI*.

Let *A* be the claim ‘there are ghosts’. *A* is positive. If an ideal reasoner can, by a priori methods, recognize both that *A* is positive and that (*PQI* \Rightarrow *A*) is not a priori,⁵ then, Chalmers notes (154), a simple argument from *T* will yield (*PQTI* \Rightarrow $\neg A$).

Chalmers’s cosmoscope thought experiment makes a plausible case for the claim that any truth *M* that anyone has ever been in a position to know is *conditionally scrutable* from *PQTI*: a cosmoscope-equipped, ideal reasoner who is otherwise ordinary is in a position to know the conditional, ‘If *PQTI*, then *M*.’ The cosmoscope puts us in a position to “look over the shoulder” of the knower as she gathers the relevant evidence for *M*. Since that evidence enables her to know *M*, it could enable us to do likewise. Similarly, the thought experiment makes a plausible case that any truth *M* that someone could be in a position to know is conditionally scrutable from *PQTI*: the cosmoscope puts us in a position to gather the evidence that the possible investigator might have gathered. If we aren’t outright skeptics and so concede that lots of things have been known or are knowable, then Chalmers’s thought experiment makes a plausible case for the claim that rather a lot – maybe everything – is conditionally scrutable from *PQTI*. But, does this justify the conclusion that the corresponding material conditionals are a priori, as A Priori Scrutability requires? Almost all of the applications of scrutability theses in the excursuses and later chapters depend on an affirmative answer to this question, since almost all of them depend on A Priori Scrutability or some strengthening thereof. The burden of the fourth chapter is to defend an affirmative answer to this question.

There is reason, though, to doubt that Chalmers’s defense succeeds. Assume that we somehow know a priori that temperature is whatever phenomenon best plays the temperature role. Then *PQI* by itself does not put me in a position to know the conditional, ‘If *PQI*, then temperature = MMKE’; in particular, it does not put me in a position to infer the conditional ‘if *PQI*, then temperature \neq caloric fluid’, where *caloric fluid* is a non-physical substance which typically mediates the causal relations between MMKE and its characteristic physical effects and otherwise plays the temperature role better than MMKE.⁶

⁵This second condition relies on the validity of an S5-like principle for a priority (154).

⁶A substance is *non-physical* in the sense intended in this stipulation if it is a substance

In general, because the cosmoscope gives me no information one way or the other about the distribution of non-physical substances, it does not put me in a position to know anything about the distribution of caloric fluid. On our assumption that we know a priori that temperature is whatever best plays the temperature role, ‘temperature = MMKE’ turns out to be a negative truth, as does ‘there is caloric fluid.’ For example, one can make ‘temperature = MMKE’ false by adding a liberal quantity of caloric fluid whose behavior suits it for the temperature role better than MMKE. We might suppose, *e.g.*, that there are anomalous concentrations of caloric fluid induced by the presence of capsaicin. So, caloric fluid, on this hypothesis, explains why we get the sensation of (high) temperature when we come into contact with hot sauce and muscle balm, even in the absence of high MMKE. Chalmers contends (124) that adding the ‘that’s all!’ claim T helps us scry ‘temperature = MMKE’ a priori from $PQTI$. But it won’t help in the simple way it helped with ‘there are no ghosts’, since ‘temperature \neq MMKE’ is *also* negative on the assumption that it is a priori that temperature is whatever phenomenon best plays the temperature role.

How, then, does the use of a cosmoscope enable me to scry ‘temperature = MMKE’ by “looking over the shoulders” of the scientists who initially discovered what temperature was? Plausibly, it’s because I am entitled to appeal to simplicity considerations in my reasoning.⁷ In particular, I am entitled to assume that the causal links between temperature and the phenomena comprised by the temperature role have a simple structure; this is how we rule out hitherto undetected causal mediation by caloric fluid, widespread overdetermination by MMKE and something else of the characteristic effects comprised by the temperature role, *etc.*

Chalmers’s arguments suggest two independent epistemic principles which might help establish that the claim regarding simplicity of causal structure is a priori scrutable from $PQTI$. First, Chalmers relies in his discussion of our version of T on the principle that all truths are a priori entailed by some combination of positive and *anti-positive* truths, where a truth is *anti-positive* just

that is not denoted by any term required to state the macrophysical truths. Chalmers urges us (110) to think of this vocabulary as including expressions for (relativistic) spatiotemporal properties, and quantum-mechanical properties. It will not include any expression for mud. Thus, mud is *non-physical* in this sense, so long as ‘mud’ cannot be defined in the sort of purely physical vocabulary Chalmers describes. Chalmers emphasizes that the case for scrutability is consistent with there being no way to define most or all of our macroscopic predicates in terms of the vocabulary of $PQTI$ (12-6).

⁷I am here following a line of thought suggested by [Block and Stalnaker, 1999].

in case its negation is positive (152, 154). To illustrate, ‘there are no ghosts’ is anti-positive, since its negation, which is a priori equivalent to ‘there are ghosts’, is positive. ‘The tallest person is male’, however, is neither positive nor anti-positive. It is negative, since it can be falsified by adding a particularly tall woman to a world. Its negation is also negative, however, since the negation can be falsified by adding a particularly tall man. In general, sentences containing unembedded occurrences of superlative definite descriptions will typically be neither positive nor anti-positive.

‘Temperature’ sentences appear to pose counterexamples to Chalmers’s claim that all truths are scrutable from positive and anti-positive truths in a way that undermines his case for A Priori Scrutability. On our assumptions, the claim of simplicity of structure is neither positive nor anti-positive. The simplicity claim can be rendered false by the liberal addition of caloric fluid; its negation can be rendered false by adding to a situation in which the causal structure in question is complex in such a way that the erstwhile complex structure turns out to be an insignificant anomaly in an ocean of otherwise simple causal structure. Similarly, ‘temperature = MMKE’ is neither positive nor anti-positive. The same goes for any claim (and for the negation of any claim) containing an unembedded occurrence of ‘temperature’. All such claims are a priori equivalent to a claim (or its negation) containing an unembedded occurrence of ‘the phenomenon that best plays the temperature role.’ The meaning of this superlative definite description ensures that we can add to worlds verifying a given claim in which it occurs unembedded in ways that falsify the claim. We seem, then, to have a family of ‘temperature’ and ‘caloric fluid’ truths that are neither positive nor anti-positive, with no clear way to scry such truths by appeal exclusively to truths outside the family. Chalmers’s first epistemic principle appears not to secure A Priori Scrutability.

Chalmers’s second epistemic principle, crucially deployed in chapter four, is that all knowable truths about the macroscopic natural world can be known by appeal to evidence from introspection of phenomenal states and evidence concerning the distribution of primary and secondary qualities in the environment (131-4, 162). It is unclear, however, how this principle can help establish the a priori scrutability of ‘temperature = MMKE’. After all, on the blatantly fictional assumption that there is caloric fluid, it is the typical cause of our sensations of temperature, and so is perceptible. Hence, the presence of caloric

fluid is a primary or secondary quality, and ‘temperature = MMKE’ will be scrutable from, among other things, the total absence of caloric fluid from the universe. But, as we have seen, this won’t help on the plausible view that ‘there is no caloric fluid’ is neither positive nor anti-positive, and so not itself a priori scrutable from *PQTI*.

Chalmers offers a response to the difficulty posed by the simplicity claim, writing that “it is very plausible that *PQTI* describes a simple world, so that ‘The world is simple’ is implied by *PQTI*” (167). This response seems inadequate. Even if it’s true that *PQTI* describes a fairly simple world, so would *PQI* conjoined to ‘there is caloric fluid’ and some suitably revised ‘that’s all!’ claim T^* . What’s needed is the claim that the world is simple *in a particular respect, viz.*, with respect to the structure of the causal links between temperature and the phenomena comprised by the temperature role.

There are lots of moving parts in play here, so Chalmers has resources at his disposal to deal with this issue.⁸ He could, for instance, appeal to some other conception of what the ‘that’s all!’ truth T says. It seems, intuitively, as if it is the business of T to turn the absence of evidence for caloric fluid into evidence for its absence. I have argued that, because ‘caloric fluid’ truths are neither positive nor anti-positive, it is mysterious how the particular version of T at hand can do this work. But one might hope that some other version will do the trick and is not otherwise unsuitable for Chalmers’s purposes. As I have indicated, Chalmers discusses alternative treatments and their attendant difficulties (151-6).

Alternatively, Chalmers could attempt to identify some positive or anti-positive truth P that is entailed a priori by the ‘caloric fluid’ hypothesis. Suppose, for instance, that C_1, C_2, \dots are all of the conceivable substances that could, consistently with what is knowable a priori, fit the description my stipulation associates with ‘caloric fluid’. Chalmers might argue that ‘there is caloric fluid’ entails a priori the disjunction ‘ C_1 exists or C_2 exists or ...’. Assuming the plausible claim that the disjunction is positive, its negation is scrutable from *PQTI* in the same way that ‘there are no ghosts’ is.

On reflection, however, it is not clear that the disjunction ‘ C_1 exists or C_2 exists or ...’ is entailed a priori by ‘there is caloric fluid’. To be sure, since by stipulation the C_n are all of the conceivable candidates for caloric fluid, any

⁸Thanks to David Chalmers for helpful suggestions for responses.

substance s that is not among the C_n is such that it is a priori that it is not a candidate. We might symbolize this as

$$(1) \quad (\forall s)(s \neq C_1 \wedge s \neq C_2 \wedge \dots \Rightarrow \Box \neg C(s))$$

where ‘ \Box ’ indicates a priority and ‘ C ’ is a predicate for being a conceivable candidate. Further, we may grant that the distinctness claims in question are all a priori, and so strengthen (1) to yield the claim that, for all substances s , it is a priori that if s is a candidate, then it is among the C_n :

$$(2) \quad (\forall s)\Box(C(s) \Rightarrow s = C_1 \vee s = C_2 \vee \dots)$$

What’s needed, however, is that it is a priori that there are no candidates other than the C_n , *i.e.*,

$$(3) \quad \Box(\forall s)(C(s) \Rightarrow s = C_1 \vee s = C_2 \vee \dots).$$

It’s not clear how the premise (2) supports the conclusion (3). In general, the pattern of inference on display here, which corresponds to the Barcan formula for a priority, is not invariably truth preserving. Let the C_n exhaust the (actual) concrete objects and ‘ C ’ stand for (actual) concreteness. So interpreted, (2) might be plausibly held to be true and (3) false. Perhaps it might be argued that, though this pattern of inference fails in many cases, its instance in the case at hand is valid. If so, that argument needs to be made.

Here I have only scratched the surface of one of the dozens of interesting and rich discussions in *Constructing the World*, offering some preliminary reasons for thinking that its principal argumentative aim has not been achieved. I have tried to show how some of the book’s central claims play out when applied to a particular example. I wish that Chalmers had done the same, especially in chapters three and four. What’s principally missing from the book are applications of the abstract claims at issue to some particular battery of test cases, illustrating in relatively concrete terms the import and plausibility of the various claims in play. This makes the book hard going, even for readers with training in the relevant areas. Still, it handsomely repays the work needed to think through the issues, especially for those with interests in contemporary epistemology, metaphysics, philosophy of language, and philosophy of mind.⁹

⁹Thanks to David Chalmers, Terence Cuneo, Mark Moyer, Derk Pereboom, and Matthew Weiner for comments on earlier drafts of this review.

References

Ned Block and Robert Stalnaker. Conceptual Analysis, Dualism, and the Explanatory Gap. *The Philosophical Review*, 108(1):1–46, 1999.

David J. Chalmers and Frank Jackson. Conceptual Analysis and Reductive Explanation. *The Philosophical Review*, 110(3):315–360, 2001.