

Proving Induction

ALEXANDER PASEAU
WADHAM COLLEGE,
OXFORD UNIVERSITY,
UNITED KINGDOM

alexander.paseau@philosophy.ox.ac.uk

Received by Greg Restall

Published February 15, 2011

<http://www.philosophy.unimelb.edu.au/ajl/2011>

© 2011 Alexander Paseau

Abstract: The hard problem of induction is to argue without begging the question that inductive inference, applied properly in the proper circumstances, is conducive to truth. A recent theorem seems to show that the hard problem has a deductive solution. The theorem, provable in ZFC, states that a predictive function M exists with the following property: whatever world we live in, M correctly predicts the world's present state given its previous states at all times apart from a well-ordered subset. On the usual model of time a well-ordered subset is small relative to the set of all times. M 's existence therefore seems to provide a solution to the hard problem.

My paper argues for two conclusions. First, the theorem does not solve the hard problem of induction. More positively though, it solves a version of the problem in which the structure of time is given *modulo* our choice of set theory.

I THE HARD PROBLEM

Call the task of persuading an inductive sceptic that inductive inference to new conclusions is truth-conducive the hard problem of induction. Inductive inference is here understood as deductively invalid inference.¹ A method is truth-conducive when it generally leads to the truth (applied properly in the proper circumstances—this is understood). And an inductive sceptic is someone who does not antecedently accept inductive inference as a route to new conclusions but is in other respects like us: he accepts deductive inference,

¹Inductive inference in this broad sense subsumes, but is not restricted to, enumeratively inductive inference. Note that our definition allows that one can reason inductively to necessary conclusions and deductively to contingent ones. Some philosophers contrast inductive inferences with deductive ones according to the nature of the conclusions they support. For instance, in his discussion of the problem of induction Vickers (2009, sec. 1.2) disbars deductive inferences from having contingent conclusions, by definition. This has the strange consequence that a syllogism such as 'All Frenchmen are charming; Olivier is a Frenchman; therefore Olivier is charming' is not deductive.

takes normal perceptual experience at face value, relies on memory as we do, is self-aware, and accepts any a priori truths available to him, for example of mathematics, modality and morality. This is a hard version of the problem of induction because the task is to persuade the sceptic, rather than to reassure ourselves, of the truth-conduciveness of inductive inferences.

A recent mathematical theorem due to Christopher Hardin and Alan Taylor seems to offer a solution to the hard problem. The theorem states that a function exists with the following property: whatever world we live in, the function predicts the world's present state from its previous states and does so correctly at all times, apart from a well-ordered subset. On the usual model of time, a well-ordered subset is small relative to the set of all times. The function in question is accordingly almost always correct, and thus the theorem apparently solves the hard problem of induction.

The first of this paper's two main conclusions is that, despite appearances, the theorem does not solve the hard problem. But there is a weaker version of the problem, namely, to persuade an inductive sceptic that inductive inference is truth-conducive given the structure of time. We may call this the parametrised hard problem of induction, since the time parameter is given. Our second and more positive conclusion is that the Hardin–Taylor Theorem solves the parametrised hard problem of induction *modulo* our choice of set theory. More precisely: if the development of set theory were to take a particular route, the parametrised problem would have a deductive solution. Although a more modest epistemological conclusion than that the hard problem of induction is soluble, it is nevertheless significant.

The paper is organised as follows: §2 contains some preliminaries, §3 sets out the theorem, §4 draws its apparent epistemological implications, §5 clarifies several points in this connection (for a shorter version of the paper this middle section may be omitted), and §§6–9 assess the theorem's implications.

2 TERMINOLOGY

Let T be the partial order consisting of the set of all times ordered by 'earlier than', and S the set of all states the world might be in. Call any function from T to S a *world history*.² A world history thus encodes a possible way the world might be at all times (it is in this sense a possible world). For example, if f is a world history and t is a member of T , the value of f at t , $f(t)$, represents the way the world is at t according to world history f . The set of world histories is the set ${}^T S$, that is, the set of functions with domain T and range S . Similarly, call any function from $\{t^* \in T : t^* < t\}$ to S , which encodes a possible way the world might be up to some time, a *partial world history*. We write ${}^{\{t^* \in T : t^* < t\}} S$ for the set of functions with domain $\{t^* \in T : t^* < t\}$ and range S . If a world history f agrees with a partial world history on the latter's domain—the times

²For ease of exposition, we shall often equate partial orders (which may be linear) with their domains.

up to some t —we symbolise the latter as f_t and say that f extends f_t . Likewise f_{t_2} extends f_{t_1} if the former's domain is a superset of the latter and the two functions agree on f_{t_1} 's domain.

A *predictive function* is a function P from $\bigcup_{t \in T} \{t^* \in T: t^* < t\}S$ to ${}^T S$ with the following property: given a partial world history f_t , $P(f_t)$ is a world history that agrees with f_t up to t (but not necessarily on any time later or equal to t). Given any predictive function P we may consider P 's success in predicting the present state of the world for varying presents.³ Let $W_P(f)$ be the set of times at which P predicts a state of the world represented by the world history f different from the actual state given by f . $W_P(f)$ is thus the set of times t in T such that $(P(f_t))(t) \neq f(t)$. We may think of $W_P(f)$ as the set of 'presents' at which the predictive function P predicts the wrong state of the present given the past in a world represented by the world history f . Likewise, the set $R_P(f)$ is the set of times at which P predicts the right state of the world in this same sense. P almost always correctly predicts f just when $W_P(f)$ is a small subset of T . What 'small' might mean in this context will be clarified shortly.

3 A THEOREM

Christopher Hardin and Alan Taylor (2008) proved:

For any linear T and any S , there is a predictive function M with the property that, for any f , $W_M(f)$ is a well-ordered subset of T (with respect to the given ordering on T).

Recall that, in ZF , the Axiom of Choice is equivalent to the Well-Ordering Principle, and thus that ZFC proves that every set can be well-ordered. In particular, it proves that a well-ordering of ${}^T S$ exists. Pick one of these well-orderings. For every partial world history f_t in $\{t^* \in T: t^* < t\}S$ let $M(f_t)$ be the least element of the subset of ${}^T S$ which consists of all the functions in ${}^T S$ that agree with f_t on all times previous to t (the least element, that is, with respect to the chosen well-ordering on ${}^T S$). M is well-defined because ${}^T S$ is well-ordered and the subset of ${}^T S$ that consists of the functions agreeing with f_t on times previous to t is non-empty, since it contains all extensions of f_t . The predictive function we seek is M .

To show this, we must prove that $W_M(f)$ is well-ordered for any f . Fixing a world history f , suppose $t_1 < t_2$ (in the given ordering of T)⁴ and that M is wrong at both t_1 and t_2 , i.e. that $(M(f_{t_1}))(t_1)$ and $(M(f_{t_2}))(t_2)$ are wrong predictions of the state of the world, so that both t_1 and t_2 are members of $W_M(f)$. Now if a function in ${}^T S$ agrees with f_{t_1} prior to t_1 then it agrees with

³Hardin and Taylor (2008) take partial world histories as equivalence classes of world histories that agree up to but not including some time.

⁴We use the same symbol ' $<$ ' for the given ordering relation T and the chosen ordering relation on ${}^T S$; it will be clear from the relata which is which. Subsets of T or ${}^T S$ are assumed to inherit these respective orderings.

f_{t_2} prior to t_1 as well since f_{t_2} extends f_{t_1} . Moreover, $M(f_{t_1})$ and $M(f_{t_2})$ are by assumption distinct, since $(M(f_{t_1}))(t_1) \neq f(t_1)$ and $(M(f_{t_2}))(t_1) = f(t_1)$, as $M(f_{t_2})$ agrees with f on all times prior to t_2 , including t_1 . Thus $M(f_{t_1}) < M(f_{t_2})$ in the chosen well-ordering of ${}^T S$. In other words, for all t_1 and t_2 in $W_M(f)$, if $t_1 < t_2$ then $M(f_{t_1}) < M(f_{t_2})$. Given the assumption that $W_M(f)$ is linearly ordered, the converse follows similarly. The set of times $W_M(f)$ at which M delivers the wrong prediction is therefore order-isomorphic to a subset of ${}^T S$ under the chosen well-ordering. Hence $W_M(f)$ is also well-ordered with respect to the given ordering on T . No specific facts about f were used in this proof, so the result generalises for any world history f .

4 EPISTEMOLOGICAL SIGNIFICANCE

What gives the Hardin–Taylor theorem great interest in connection with the hard problem of induction is that for some T , well-ordered subsets of T are small compared to T itself. In particular, suppose T is the set of real numbers.⁵ Any subset of the real numbers well-ordered under their standard order is countable and has measure zero.⁶ Hence if T is the set of real numbers, standardly ordered, $W_M(f)$ is small compared to T for any f whatsoever.⁷ Likewise if T is an interval of real numbers, or a union of intervals of real numbers, or any subset of the reals containing a real interval.⁸ Since time is usually modelled as the real numbers, or more generally spacetime is modelled as a geometry based on the real numbers, the Hardin–Taylor theorem therefore seems to show that the hard problem of induction is solvable on our usual physical assumptions.

It is worth stressing just how strong the theorem is. Suppose T is the set of real numbers and that the actual world's world history is $f_{@}$. A consequence of the theorem is that there is a predictive function M such that $M(f_{@}(t))$ agrees with $f_{@}(t)$ for all but a small subset of T . This fact might be thought easy to contrive: simply take M to be the function which outputs the constant function $f_{@}$ for any input. However, the theorem demonstrates that the same function M whose existence it guarantees would work even if our actual world history were different: M is almost always correct for any choice of world history. On the assumption that T is the set of real numbers, the theorem is the strong claim that there is a predictive function that is almost always correct

⁵For brevity, I usually write ' T is the set of real numbers' in place of ' T is modelled as the set of real numbers', etc.

⁶By 'countable' we mean either countably infinite or finite. If X is a well-ordered subset of reals under their standard order then there is a rational between its α^{th} and $\alpha + 1^{\text{th}}$ members, so the cardinality of X is no greater than that of the rationals, which are countable. A countable set has measure zero since singletons have measure zero and measures are countably additive.

⁷Hereon we suppress the qualification that the order of the reals is the standard one and assume all familiar domains have their standard order unless otherwise specified.

⁸An (open) interval of real numbers is here taken to be a set of the form $\{x : p < x < q\}$, usually written (p, q) , where p and q are real numbers such that $p < q$, or p is $-\infty$ and q is real, or q is $+\infty$ and p is real, or p is $-\infty$ and q is $+\infty$.

whatever the world is like. It is not the trivial claim that, given a particular world, there is an almost always correct predictive function for it.

Another aspect of the theorem's strength is its independence from probability. The Achilles heel of Bayesian attempts to solve the hard problem of induction is that they have to assume that the sceptic's prior probability of $f_{@}$ being our actual world history is non-zero (where $f_{@}$ is our actual world history).⁹ In contrast, the Hardin–Taylor theorem does not depend on particular world histories having positive prior probability. It is independent of all probability assignments.

Now if T is the set of rational numbers, $W_M(f)$ could be the same size as T since the rationals and therefore their well-ordered subsets are countable. For example, the set of natural numbers considered as a subset of the rational numbers is well-ordered and countable. However, $W_M(f)$ might still be considered a smaller set than T in senses of 'small' other than cardinality. For example, one has to live through an infinity of time instants not in $W_M(f)$ to reach the next time instant in $W_M(f)$. If the sceptic applies M in a world whose time instants are rational then every time M leads him astray it will subsequently be correct for an infinity of times. Or consider that $W_M(f)$ is nowhere dense in T .¹⁰ In this topological sense, well-ordered sets of rationals are small relative to all the rationals or to any interval of rationals. Well-ordered subsets of the rational numbers have mathematical properties typical of smallness, even if they are of the same cardinality as the set of rationals itself.

A third possibility is that T is a subset of the integers, e.g. T might be the natural numbers or some finite segment thereof. In that case, $W_M(f)$ can be equal to T , and so $W_M(f)$ is not guaranteed to be smaller than T , either in the sense of cardinality or in any other sense. And in the case in which T is the set of integers, at any given instant the set of future times in $W_M(f)$ can be equal to the set of all future times in T . In other words, M may be perfectly falsehood-conducive in the present and future.

There are of course many other possibilities for T than these salient ones.

5 COMMENTS

The Hardin–Taylor theorem apparently solves the hard problem of induction. Roughly, the reason is that the sceptic can avail himself of the predictive function M to predict the state of the world at any given time. Given the standard scientific assumption that the set of times T in our world is an interval of real numbers, the set of times at which M gives a wrong prediction is small. In other words, the sceptic can be persuaded that there is a predictive function which predicts the present state of the world given past states, and does so correctly for all but a small number of 'presents', with 'small' understood in the

⁹Howson (2000, ch. 4) offers an overview of Bayesian arguments.

¹⁰ X is nowhere dense in T if any non-empty interval in T contains a non-empty interval disjoint from X . (An open interval in the rationals is the rational analogue of a real interval.)

sense of cardinality. The purpose of this section, which can be skipped on first reading, is to sharpen this rough explanation of the theorem's epistemological import in a numbered series of comments.

- (1) As noted, the theorem essentially depends on the Axiom of Choice, which is used to well-order the set ${}^T S$. Choice refuseniks will accordingly not accept the theorem. This should not trouble the rest of us.
- (2) The set T need not be linear. For instance, T might be a branching tree representing a world in which time splits into two branches every second. The result applies to such branching time scenarios too. Our application of the theorem uses the fact that $W_M(f)$ is a linearly ordered set, which follows from the assumption that the sceptic's path through time is linear. This seems a reasonable assumption so long as the sceptic is assumed to be a unified consciousness. If this a contingent feature of sceptics in our world, so be it: the solution, if successful, is contingently successful. For similar reasons, the result applies in a relativistic setting.
- (3) For appropriate T , the Hardin–Taylor theorem establishes that $W_M(f)$ is small relative to T . Yet consider the subset O of T at which the sceptic observes the state of the world. What if O is not T ? For example, what if the sceptic's existence occupies a cosmically short interval of time? Can the theorem still reassure him that his observations will be correct on all but a small number of occasions?

One way to sidestep this worry is to think of the theorem as applying to a *subjunctive* version of the hard problem. The theorem reassures the sceptic that, for appropriate T , were he to live and observe the states of the world for longer (in particular, were he to observe the state of the world throughout T), the predictive function M would only let him down a small number of times.

In the *indicative* version of the hard problem, in contrast, the sceptic wonders how often use of a predictive function will let him down during his actual observational lifespan. The answer to this question depends not just on T but also on O , as the following two scenarios demonstrate. In the first scenario, T is the set of real numbers and O is a well-ordered sequence of real intervals $(a_1, b_1), (a_2, b_2), \dots$, where $a_i < b_i < a_{i+1}$. For example, O might be a finite sequence of intervals $(a_1, b_1), \dots, (a_n, b_n)$ during which the sceptic is awake and paying attention to the world around him. Now if a_1 is a real number (not $-\infty$) the sceptic has not observed all of the past. So, you might ask, how is the sceptic supposed to know the partial world history up to a_1 in order to feed it into M ? The answer concedes that he doesn't know what $f a_1$ is, but that he can nonetheless apply the result by taking T to be (a_1, b_1) . The Hardin–Taylor theorem reassures him that $W_M(f)$ is small relative to this interval for any f . He can then repeat the application with T as (a_2, b_2) , etc. Or instead of taking a different T each time and applying the theorem anew to each of his observational intervals, he can take T itself to be O , that is, the sequence of real intervals $(a_1, b_1), \dots, (a_n, b_n)$, and apply the theorem in one fell swoop to

this choice of T . Either way, the theorem reassures him that M lets him down on a small number of occasions only.

Another scenario is when T is the set of real numbers and the sceptic only observes the state of the world during some finite set of times in T , say the integer times from 0 to 100. In that case, the result is compatible with each of his 101 predictions being wrong. Whether $W_M(f)$ is small relative to the sceptic's life span is therefore not determined by the fact, if it is one, that $W_M(f)$ is small relative to T .

In sum, if the theorem is to apply to the indicative version of the hard problem of induction, the set O must be of a certain kind. Roughly speaking, O must not be too small relative to T for the result to bite. For the sake of simplicity, we suppress mention of O , it being understood that discussion of the theorem's application to the hard problem in its indicative version depends not just on T but on O as well.

Note finally that the theorem's application depends on idealising the sceptic. For instance, taking $T = \mathbb{R}$ and $O = (0, 1)$ assumes that the sceptic has observed all the states of the world during the latter time interval. Yet no human sceptic appears capable of recording this many observations throughout $(0, 1)$. Perhaps this is what Hardin and Taylor mean by saying: "We should emphasize that these results do not give a practical means of predicting the future" (2008, p. 92).¹¹ However, this idealisation does not damage the theorem's application to the problem of induction, which usually entertains an idealised sceptic.

(4) The Hardin–Taylor theorem does not draw a distinction between observable and unobservable states. Yet its epistemological application to the hard problem of induction must do so. If the sceptic is incapable of knowing what f_t is at any given t , or of verifying whether $M(f_t)(t)$ is in fact the correct state of the world at t , he won't be able to use the theorem to reassure himself that inductive inference can be truth-conducive. If S consists of states unobservable by the sceptic, the function M makes predictions about unobservable states, and the sceptic cannot check the state of the universe to narrow the range of world histories as time progresses. More precisely, if S consists of unobservable states then for a given world history f , the set $R_M(f)$ of instants at which M gives the right answer might, for all the sceptic knows, be null. Correspondingly, the argument for the well-ordering of the set $W_M(f)$ of instants at which M gives the wrong answer breaks down: if t_1 is prior to t_2 and the world histories $M(f_{t_1})$ and $M(f_{t_2})$ give unobservable predictions at t_1 and t_2 respectively, it does not follow that $M(f_{t_1})$ and $M(f_{t_2})$ must be distinct, since they could agree at t_1 if $M(f_{t_1})(t_1)$ is not equal to $f(t_1)$ but both are unobservable states of the world. $M(f_{t_1})$ and $M(f_{t_2})$ might thus both be wrong at t_1 , for all the sceptic can tell. Updating one's partial world history on the basis of experience

¹¹It is unlikely they are alluding to the point to be made in §6, since they do not raise the issue of definability.

is crucial to the application of the Hardin–Taylor theorem to the hard problem, in order to eliminate incorrect candidate world histories. This cannot be done if the states of the world according to those candidate world histories are not observed by the subject attempting the updating.

If in contrast S consists of states of the world observable by the sceptic then the epistemological application of the Hardin–Taylor theorem to the hard problem of induction proceeds without hindrance. For example, suppose that the sceptic has noticed that in the past the law of gravity has held in his environs.¹² The Hardin–Taylor theorem reassures him that the function M , which counsels him let us suppose to assume that the law of gravity holds in his environs at t if it has held in his environs at all times previous to t , can only let him down on a well-ordered set of times.¹³ Given that in this instance S is a set of observable states, the theorem does not justify the belief that the law of gravity holds more generally beyond his environs. This is not at all damaging. The hard problem of induction is to convince the sceptic of the truth-conduciveness of inductive inference to new conclusions (in some circumstances), for example, the inference to the nature of his future environment from his past experience of it. It is *not* to convince him of facts we generally believe about the world, for instance that it extends far beyond our perceptual range and that the law of gravity holds throughout it (more on this in the sixth comment).

In sum, although the Hardin–Taylor theorem itself draws no distinction between observable and unobservable states, its application to the hard problem naturally proceeds by construing S as consisting of states observable by the sceptic.

(5) Suppose that T is the set of real numbers.¹⁴ Then at any given time t , the possible world histories that agree with the sceptic’s partial world history are uncountable. Yet only one of these uncountably many possible world histories is the correct one for that world. In one sense, then, the sceptic who uses the Hardin–Taylor theorem never learns from experience: at any given time t , the uncountably many rivals to the actual world history that agree with his partial world history are just as ‘live’ for him as the actual world history itself. Although any given rival to the actual world history will be ruled out at some time by observation, at any given time the predictive function M has zero chance of outputting the actual world history. Even on the assumption that he is omniscient about the past, the existence of M does not change the fact that at any given time the sceptic is none the wiser as to which of the uncountably many possible extensions of his partial world history is the correct one.

Viewed from this perspective, the Hardin–Taylor theorem may seem paradoxical. Yet the semblance of paradox is dispelled by keeping the bigger pic-

¹²By the sceptic’s environs at any given time we mean what he observes at that very time, irrespective of whether it has taken some time for that information to reach him.

¹³Well-ordered, to repeat, with respect to the given ordering of time.

¹⁴And that S has more than one element.

ture in mind, as Hardin and Taylor stress in their paper. Locally speaking—at any given time t —the predictive function M applied to f_t does indeed have no more chance of being right about the observable state of the world at t than any other method (necromancy, reading tea leaves, guessing, etc.). The local perspective makes M 's success a mystery. Viewing things globally, however, we see that M can only go wrong on a well-ordered set of times, which must be small if $T = \mathbb{R}$. M can be successful in this global sense without being any better locally than its rivals.

(6) The predictive function M may not be a function we would recognise as representing a correct means of projecting from past experience. For instance, M may not counsel the sceptic to infer that all emeralds he will ever observe are green from the fact that all emeralds he has previously observed are green. Despite his having only observed green emeralds, it may counsel him that the emeralds he will henceforth observe are red; re-applied an hour later it may counsel him that they are yellow; and two hours later that they are blue with pink dots. As Alexander George (2007) points out, there is no reason to think that M vindicates our inductive practices; it may instead vindicate 'gruesome' inductive inferences.

This is not a criticism of the application of the theorem to the hard problem of induction, which is just one version in the cluster of problems known loosely as 'the problem of induction'. The problem of persuading the sceptic of the truth-conduciveness not just of any old inductive inferences but of our own dear and accustomed ones is an even harder task than solving the hard problem of induction. A solution to the hard problem of induction need not also be a solution to this further problem.

(7) Lots of functions other than M will do the predictive job. Suppose that instead of taking $M(f_t)$ to be the least element of the subset of ${}^T S$ consisting of all the functions that agree with f_t on times previous to t we take $M(f_t)$ to be this subset's second least element.¹⁵ As can be verified by running the proof with this choice of M , the resulting set $W_M(f)$ is also well-ordered (with respect to the given ordering of T). And of course there are many well-orderings of ${}^T S$ one could choose in the first place; M depends upon this initial choice of well-ordering as well.

In attempting to persuade the sceptic of the truth-conduciveness of a predictive function we are thus faced with an *embarras de choix*. There are many such predictive functions M . However, this does not threaten the Hardin–Taylor theorem's application to the hard problem of induction. On the contrary, if the argument is sound it shows that the inductive sceptic can be reasonably persuaded to choose one from several almost always correct predictive functions. So much the better.

(8) We come finally to the question of *when* the sceptic is supposed to apply

¹⁵This is a different choice assuming that S has at least two members.

the function M . The various problems that come under the heading of the problem of induction are traditionally concerned with how to justify a claim about unobserved states of the world on the basis of its observed states. Yet the first thing that strikes everyone about the Hardin–Taylor theorem’s application to the hard problem of induction is that it does not seem to be an instance of this general category, because the predictive function M predicts the present rather than the future from past observations. Since the sceptic observes the present state of the world, applying M to the present instant seems not to afford him any more information than is experientially available to him. The observed present state of the world is by definition already part of his past and present experience. So in what sense does the Hardin–Taylor theorem offer a potential solution to the hard problem of induction?

The answer is not entirely straightforward and will occupy the rest of this section.¹⁶ A first stab at answering the question is to exploit the theorem by applying the prediction to the next moment in time. Suppose t is the present and t^+ the next moment in time following the present. Since we are assuming that the sceptic observes the present and that he has observed the past, he knows what f_{t^+} is (recall that f_{t^+} is the partial world history up to but not including t^+). Hence he can apply M to f_{t^+} to output a prediction for the observable state of the world at t^+ .

A problem with this strategy to be discussed in §7 is that the sceptic cannot assume that any given moment in time has a successor, even if this has proved true in the past. A second problem is that the worlds in which every moment in time has a successor are not worlds in which a well-ordered subset of T is small compared to T . Any linearly ordered set in which every element has an immediate successor or is the set’s maximum element is made up of some copies of the integers or the natural numbers, with a copy of the negative numbers or a finite ordered set possibly tagged on at the end.¹⁷ Yet a well-ordered subset of the integers or the positive numbers or a finite ordered set need not be small in any pertinent sense of ‘small’. It is true that an element of a well-ordered subset of the integers can only have finitely many predecessors; yet it is also true that all the integers after it may be members of the well-ordered subset in question. Thus if the sceptic can come to justifiably believe that every time instant has an immediate successor, he can no longer use the theorem to reason that the set of times at which M fails is small compared to T .

An alternative answer to our question of when and how to apply M is that

¹⁶Suppose the theorem could do no more than show that the present state of the world is predictable correctly for most presents. It would belittle the theorem’s epistemological value to claim it as nil because the present is already observed. The theorem would offer an inferential warrant on top of an observational warrant for the world’s current state.

¹⁷Let $S(a)$ be the successor of a and $P(a)$ the predecessor of a . Then the set of points finitely accessible from a by applications of S or P is either (i) $\{\dots, PPa, Pa, a, Sa, SSa, \dots\}$, or (ii) $\{P^m(a), \dots, PPa, Pa, a, Sa, SSa, \dots\}$, or (iii) $\{\dots, PPa, Pa, a, Sa, SSa, \dots, S^n(a)\}$, or (iv) $\{P^m(a), \dots, PPa, Pa, a, Sa, SSa, \dots, S^n(a)\}$.

M should not be applied to predict the present. It should be used instead to predict the state of the world at some fixed time in the future. For example, M may be used to predict the observable state of the world one second after the present. With the same choice of M , the prediction is $M(f_t)(t + 1)$ rather than the prediction $M(f_t)(t)$.

The problem with this suggestion is that the Hardin–Taylor theorem no longer applies. If $t_1 < t_2$ then it follows, as in §3, that $M(f_{t_1})$ cannot be any greater than $M(f_{t_2})$ in the well-ordering of T S. But that $(M(f_{t_1}))(t_1 + 1)$ and $(M(f_{t_2}))(t_2 + 1)$ are wrong predictions does not entail that $M(f_{t_1})$ and $M(f_{t_2})$ are distinct. For if t_2 is less than a second after t_1 the partial world history f_{t_2} contains no information about the observable state of the universe at $t_1 + 1$. Hence the world history $M(f_{t_2})$, which extends f_{t_2} , may have the same value as $M(f_{t_1})$ at $t_1 + 1$. In other words, all times within a one-second interval of each other may be in $W_M(f)$, and consequently the predictive function M need not almost always be correct even if T is the set of reals. The objection generalises to the proposal that M should be used to predict the future ϵ time units later for any $\epsilon > 0$.

A proposal in the same vein is to use M to predict the observable state of the world at one-second intervals. If T is the set of reals and the sceptic starts observing at 0, M might be used to predict the observable state of the world at times 0, 1, 2, 3... The problem with this variant of the previous suggestion should be obvious. If the sceptic applies M to predict the state of the world at one-second intervals, the set of times at which predictions are made is thereby guaranteed to be well-ordered. For his predictions will be at staggered time instants: 0, 1, 2, 3, ... This sequence is by definition well-ordered, so it could be equal to $W_M(f)$: the sceptic could go wrong at all these times. Neither of these ways of implementing the second answer therefore overcomes the problem.

Here is my suggested answer to the question. Assume for the rest of this section that T is the set of reals. A well-ordered set of reals has the following property: for any real number x , be it in the set or not, there is a positive ϵ_x such that the interval $(x, x + \epsilon_x)$ is disjoint from the set.¹⁸ Now for any f the set $W_M(f)$ is well-ordered, so this property of well-ordered sets of reals implies that the set of times at which the predictive function M wrongly predicts the immediate future is null. ‘The immediate future’ is here understood not as the time instant immediately succeeding the present (there is no such time when $T = \mathbb{R}$), but rather as any time interval starting from but not including the present. Suppose then that the sceptic uses the function M to predict his world’s world history so that at any given time t he predicts that the world history is $M(f_t)$. Whether or not this prediction lets him down at the present time t , it will stand him in good stead in the immediate future, that is, through-

¹⁸Consider the subset of the well-ordered set consisting of elements greater than x . If this subset is empty, the result follows. If not, it has a least element, y . Pick ϵ_x such that $0 < \epsilon_x < y - x$.

out a time interval $(t, t + \epsilon_t)$ for some positive amount ϵ_t . This is the sense in which the Hardin–Taylor underwrites an inference about future observable states of the world, and hence from observed states to unobserved states of the world.

Thus understood, the predictive value of M should not be overestimated. What ϵ_t might be for any given t cannot be determined in advance. In particular, the sceptic cannot make any plans for the future on the basis of the theorem. For example, if M predicts that no predators will pounce at t , the sceptic cannot go to sleep at t safe in the knowledge that his life is not endangered. For one thing, t might be one of those few times at which M goes wrong. For another, even if M is correct at t and thus at an interval starting at t , as it almost always is if $T = \mathbb{R}$, the interval $[t, t + \epsilon_t)$ throughout which it is correct may be shorter than the time needed to renew his energies before rejoining battle with his predators. In fact, there is no positive amount ϵ such that ϵ_t must be larger than ϵ for any t : the greatest lower bound of the values ϵ_t over all times t may be 0. And the sceptic must keep on updating his partial world history at every instant in $[t, t + \epsilon_t)$; for if he doesn't—if he takes $M(f_t)$ as his best estimate of the world function throughout $[t, t + \epsilon_t)$ —he could go wrong at all these times. Nevertheless, the sceptic can rest assured that the function M will, for most times (if $T = \mathbb{R}$), correctly predict some stretch of the immediate future, of positive but unknown duration, in the sense we have carefully specified.

6 INFORMATIVENESS: FUNCTIONS AND METHODS¹⁹

We now argue for two main conclusions. The negative conclusion is that the theorem does not solve the hard problem of induction. The positive conclusion is that, given a particular choice of extension of set theory, it solves a weaker form of the problem.

So far we have gone along with the idea that M is a method or policy, which counsels various predictions. But is that so? Take the most promising application of the theorem, in which $T = \mathbb{R}$ (and S has more than one member). The proof that there is a well-ordering of the reals is nonconstructive in the following sense. A term is a formula $\{x : \phi(x, y_1, \dots, y_n)\}$ where $\phi(x, y_1, \dots, y_n)$ is a wff whose free variables are y_1, \dots, y_n (if $n > 0$). A term is said to be definable if and only if it can be expressed with no free variables. This definition accords with the intuitive idea of definability, namely expressibility by a sentence rather than an open formula. Using forcing, it can be shown that no well-ordering of the real numbers is definable in ZFC.²⁰ Applied to the case in which $T = \mathbb{R}$,

¹⁹This section overlaps with my (2008).

²⁰The result was proved in Feferman (1965). Note that if we augment the language of set theory with a term ' $<_{\mathbb{R}, S}$ ' and add to the axioms the formula ' $<_{\mathbb{R}, S}$ is a well-ordering of ${}^{\mathbb{R}}S$ ', it would be easy to name a well-ordering of ${}^{\mathbb{R}}S$: its name would simply be ' $<_{\mathbb{R}, S}$ '. However, although we would then possess a name for the strategy that name would be uninformative. A

the indefinability of a real well-ordering shows that the Hardin–Taylor theorem is nonconstructive in a strong sense. The proof derives the function M from a well-ordering of the set ${}^{\mathbb{R}}S = \mathbb{R}^S$ of all functions from \mathbb{R} to S . But it would follow straightforwardly from the definability of a well-ordering of this set that the real numbers also have a definable well-ordering.²¹ Thus the well-ordering of \mathbb{R}^S which the proof uses to derive M is not definable. Hence the Hardin–Taylor proof does not give us a method or rule or means of inference for predicting the present. A method for predicting the present given the past that cannot be expressed in words does not merit the label. If we cannot in principle express or exhibit our supposed method, it is no method at all, not even an in principle method.

To elaborate on this point, the following are methods or rules or means of inference: ‘if you have observed that A_1 to A_n are B_s , infer that all A_s are B_s ’, or ‘if you have observed that A_1 to A_n are B_s , infer that no other A_s are B_s ’. But the latter is not: ‘if you have observed that A_1 to A_n are B_s , infer that all A_s are B_s if $f(t) = 1$ ’, where f is an unspecified function. This instruction is impossible to follow unless one knows what f is. In the context of the hard problem of induction, what this shows is that we cannot convince the sceptic that there is some inductive method or rule or means of inference which is truth-conducive. At most, we can convince him of the existence of a truth-conducive function. What we would like to persuade the sceptic of, however, is that such a function exists and that it can be applied, at least in principle, by human reasoners.²²

Now in some models of ZFC , there are definable well-orderings of both the set of reals \mathbb{R} and the set $\mathbb{R}^{\mathbb{R}}$ of functions from the reals to the reals. An example is Gödel’s constructible universe \mathbf{L} . If one adds the ‘axiom’ of constructibility to ZFC , which states that the class \mathbf{V} , the universe of sets, equals the class \mathbf{L} , the universe of constructible sets, then all functions become definable. The reason is that there is a global well-ordering of the universe in \mathbf{L} (every set first appears at some L_α and each set in $L_{\alpha+1}$ is defined by a formula in the language of set theory with parameter sets drawn from the well-ordered L_α), so one can always define a set with a particular property (if there is one) as the least set with that property in the global well-ordering. Yet the axiom of constructibility is not generally accepted.²³ Similarly, the inner model \mathbf{HOD} , the class of hereditarily ordinal-definable sets, is also globally well-orderable and thus con-

bare stipulated name of this kind has no informational value.

²¹For example, the real number r may be identified with its characteristic function χ_r , where $\chi_r(r) = s_1$ and $\chi_r(x) = s_2$ if $x \neq r$, for some distinct s_1 and s_2 in S .

²²In several areas of mathematics, for example game theory or descriptive set theory, the word ‘strategy’ is used to denote arbitrary functions of a certain kind, be they expressible or not. We should not confuse this usage with the standard non-mathematical one, which is the one relevant to the problem of induction. (Ditto for related uses of ‘method’, ‘rule’, ‘means of inference’...)

²³Gödel’s initial inclination to accept it notwithstanding. See Maddy (1993) for a summary of arguments against $\mathbf{V} = \mathbf{L}$.

tains a definable well-ordering of the set ${}^{\mathbb{R}}\mathbb{R}$.²⁴ Yet neither $\mathbf{V} = \mathbf{HOD}$ nor any other axiom extending ZFC that implies that real well-orderings are definable commands general acceptance.

In sum, the Hardin and Taylor theorem implies that a mostly correct predictive function exists if $T = \mathbb{R}$. However, it does not follow that this function is definable. If in a century's time say our accepted best set theory allows for definable well-orderings of such functions—for example, if some principle that entails the definability of a well-ordering of ${}^{\mathbb{R}}\mathbb{R}$ comes to be accepted—the theorem could be used to establish the existence of a definable mostly correct predictive function, assuming our world is of this type.²⁵ As far as today's mathematics goes, however, the proof does not furnish the sceptic with a truth-conducive means of inference.

7 THE SET T

How are we supposed to convince the sceptic that the set of times has a particular structure? For example, how might we convince him that $T = \mathbb{R}$? There are in fact two objections here. The first is difficult to assess, but the second is fatal. The first objection asks how the sceptic is supposed to convince himself that the times he has experienced are modelled as the reals say rather than the rationals or even the integers. Is the structure of experienced time an observed fact? Arguably, no. That past time can be modelled as an interval of real numbers (say) seems to be an inference to the best explanation of experience rather than something given by experience itself. Yet inference to the best explanation is not a means of inference available to our inductive sceptic. So by assumption the sceptic is not justified in believing that the structure of past times—the times he has already experienced—is given by the real numbers.

As I say, this objection is difficult to assess, because it is difficult to determine whether or not the structure of time is experientially given. As already acknowledged the sceptic is an idealised being, so perhaps he can experience a real time interval as a real time interval even if we more limited beings are unable to do so. There is more to say on this point, but this is not the place to do so. Either way, a second and fatal objection looms.

Suppose that time has so far assumed the structure of the real numbers, or any set whose well-ordered subsets are small, even observably. Nevertheless, the sceptic has no grounds for thinking this structure will continue into the future. It follows that he cannot be persuaded that there exists a truth-conducive function M . For example, perhaps in the past T has had the structure of a real interval; yet that does not prevent it from being discrete in future. Given his inductive scepticism, the inductive sceptic cannot brush aside this possibility. As Hume pointed out, invoking a principle of uniformity, which states that the

²⁴For a review of the relevant properties of \mathbf{L} and \mathbf{HOD} , see ch. 13 of Jech (2003).

²⁵This would be an in principle construction of course, since the definability of a well-ordering does not imply, say, its effectiveness.

future will resemble the past, begs the question against the sceptic.²⁶ This circularity objection is the besetting problem for purported solutions to the hard problem of induction. For all the sceptic knows, then, the set of present and future times he will experience is a discrete ordered set of a million instants. If so, M can go wrong at all of them, and the Hardin–Taylor theorem does not provide a mathematical recipe for doing any better.

This objection is unanswerable: the result that the set of present and future times in $W_M(f)$ is a small subset of the set of present and future times in T depends on what T is. Yet his future's topology is not a fact available to the sceptic.

8 INFORMATIVENESS: THE SET S ²⁷

Does a version of the same objection apply to the choice of S ? You might think it does—after all, why should the inductive sceptic assume that some particular set S is the set of world-states? All the world states he has so far observed may be elements of S . But that this will continue to be the case is not a fact given by experience, and there is no acceptable argument for it from the sceptic's point of view. Does this mean that we cannot get him to accept that the theorem applies to his situation, as it assumes that world states will in future continue to be drawn from S ?

No. The sceptic may take S to consist of all previously observed states s_1, s_2, \dots , together with the catch-all hypothesis 'none of the above'. For example, if the previously observed states are (the mutually exclusive) s_1, s_2, s_3 , the catch-all hypothesis would be $\neg s_1 \wedge \neg s_2 \wedge \neg s_3$; in other words $S = \{s_1, s_2, s_3, \neg s_1 \wedge \neg s_2 \wedge \neg s_3\}$. Or he might take $S = \{s_1, s_2, s_3, s_4, \neg s_1 \wedge \neg s_2 \wedge \neg s_3 \wedge \neg s_4\}$ for some state s_4 distinct from and mutually exclusive with s_1, s_2 and s_3 . Or he might take S to consist of all the states of affairs he can conceive of. Some of these he might have experienced, others he might not have. Suppose for example that the sceptic can only conceive of two particulars, a and b and two properties F and G . The conceivable atomic states of affairs are for him Fa, Fb, Ga, Gb , and the conceivable world states, which may be taken as the members of S , are their sixteen combinations $\pm Fa \wedge \pm Fb \wedge \pm Ga \wedge \pm Gb$.

Which set the sceptic adopts as S will naturally affect the predictive function M 's informativeness. For instance, taking S as $\{Fa, \neg Fa\}$ will result in a predictive function with binary output. Taking S as the sixteen-membered set consisting of the conjuncts of the atoms Fa, Fb, Ga, Gb and their negations will

²⁶On the circularity of all such attempts Hume writes: "... all our experimental conclusions proceed upon the supposition that the future will be conformable to the past. To endeavour, therefore, the proof of this last supposition by probable arguments, or arguments regarding existence, must be evidently going in a circle, and taking that for granted, which is the very point in question." (1748, pp. 35–6).

²⁷In this section we are prescinding from the objections in §6 and §7. We also assume that S consists of observable states, for reasons mentioned in the fourth comment in §5.

result in a more informative predictive function with 16 possible outputs.²⁸ The informativeness of M (*modulo* the definability objection in §6) thus depends on S . If the sceptic wishes to extract maximum information about the future state of his environs, he should choose as fine-grained a set as possible. Either way, whether his choice of S is fine- or coarse-grained, there is no corresponding objection to the one in §7. The sceptic can simply adopt any partition of world states as the set S and apply the theorem for this particular choice of S .

9 THE PARAMETRISED PROBLEM

On first appearance, the Hardin–Taylor theorem promises an a priori solution to the hard problem of induction. Alas, for the two reasons given, it doesn't deliver on that promise. The non-constructive nature of the proof seems to imply that the theorem does not establish a means of inductive inference, even an in principle means (§6). The actual proof doesn't define a predictive function, nor does it explain how to go about defining one. Perhaps the deepest objection is the second one (§7), that the argument's application to the hard problem of induction assumes a time structure that the sceptic has no reason to accept. Hume's circularity objection strikes again.

And yet there is something of positive epistemological value to be extracted from the Hardin–Taylor theorem. It solves a weaker problem of induction than the hard problem, *modulo* the definability objection. This problem is that of persuading an inductive sceptic that there is a mostly correct inductive function given T , that is, given the structure of time.²⁹ Earlier, we called this the parametrised hard problem of induction: the hard problem of induction with T a shared parameter in the dialectic between sceptic and anti-sceptic. If it weren't for the problem of definability, the Hardin–Taylor theorem would solve this problem. If only our generally accepted set theory took a slightly different form (for example, if it incorporated the extra axiom $\mathbf{V} = \mathbf{L}$ or $\mathbf{V} = \mathbf{HOD}$), and if time were correctly modelled by some interval of the real numbers, as we generally think it is, the theorem would demonstrate that the parametrised problem of induction for the actual world is solvable. If set theory were to be extended in such a way that well-orderings of the reals are definable—certainly not something we can guarantee, yet neither something we can conclusively rule out—the parametrised problem has an a priori solution. To put it another

²⁸Consider two sets S_1 and S_2 , each of which is an exhaustive and mutually exclusive set of world states. Let f_1 represent the world history with world-states drawn from S_1 and f_2 the world history with world-states drawn from S_2 ; thus $f_1 \in {}^T S_1$ and $f_2 \in {}^T S_2$ for some set of times T . Suppose M_1 and M_2 are Hardin–Taylor functions derived for this respective choice of set of world states. M_1 and M_2 correctly predict f on all but all a well-ordered set of times, so they are compatible with each other on all but a well-ordered set of times, since they conflict at most on the union of two well-ordered sets of times, which is well-ordered.

²⁹As noted in §8, the sceptic can choose any S , though the finer-grained S the more informative the mostly correct predictions delivered by M will be.

way, it is consistent with present-day mathematics that the parametrised problem is a priori solvable.

Although it falls short of all we might have wished for in reply to the sceptic, this is still a surprising epistemological implication. Many of us would not have thought this much could be achieved.³⁰

REFERENCES

- [1] FEFERMAN, S. (1965) “Some applications of the notions of forcing and generic sets”, *Fundamenta Mathematicae* 56, pp. 325–345.
- [2] GEORGE, A. (2007), “A Proof of Induction?”, *Philosophers’ Imprint* 7, pp. 1–5.
- [3] HARDIN, C. & TAYLOR, A. (2008), “A Peculiar Connection Between the Axiom of Choice and Predicting the Future”, *American Mathematical Monthly* 115, pp. 91–96.
- [4] HOWSON, C. (2000), *Hume’s Problem: Induction and the Justification of Belief*, Oxford: Oxford University Press.
- [5] HUME, D. (1748), *An Enquiry Concerning Human Understanding*, eds. L. A. SELBY-BIGGE (1893) and P. H. NIDDITCH, (1975), Oxford: Oxford University Press.
- [6] JECH, T. (2003), *Set Theory*, Berlin: Springer.
- [7] LEVY, A. (1979), *Basic Set Theory*, Berlin: Springer-Verlag.
- [8] MADDY, P. (1993), “Does V Equal L?”, *Journal of Symbolic Logic* 58, pp. 15–41.
- [9] PASEAU, A. (2008) “Justifying Induction Mathematically: Strategies and Functions”, *Logique et Analyse* 203, pp. 263–269.
- [10] VICKERS, J. (2009), “The Problem of Induction”, *Stanford Encyclopedia of Philosophy*.

³⁰Thanks to Alan Taylor, Christopher Hardin, Greg Restall, Peter Millican and an anonymous *AJL* referee.

The *Australasian Journal of Logic* (ISSN 1448-5052) disseminates articles that significantly advance the study of logic, in its mathematical, philosophical or computational guises. The scope of the journal includes all areas of logic, both pure and applied to topics in philosophy, mathematics, computation, linguistics and the other sciences.

Articles appearing in the journal have been carefully and critically refereed under the responsibility of members of the Editorial Board. Only papers judged to be both significant and excellent are accepted for publication.

The journal is freely available at the journal website at

<http://www.philosophy.unimelb.edu.au/ajl/>.

All issues of the journal are archived electronically at the journal website.

SUBSCRIPTIONS Individuals may subscribe to the journal by sending an email, including a full name, an institutional affiliation and an email address to the managing editor at ajl-editors@unimelb.edu.au. Subscribers will receive email abstracts of accepted papers to an address of their choice. For institutional subscription, please email the managing editor at ajl-editors@unimelb.edu.au.

Complete published papers may be downloaded at the journal's website at <http://www.philosophy.unimelb.edu.au/ajl/>. The journal currently publishes in pdf format.

SUBMISSION The journal accepts submissions of papers electronically. To submit an article for publication, send the \LaTeX source of a submission to a member of the editorial board. For a current list of the editorial board, consult the website.

The copyright of each article remains with the author or authors of that article.