Journal Pre-proof

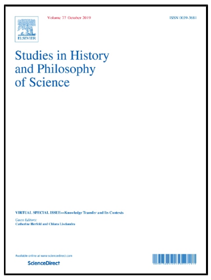

The Material Theory of Object-Induction and the Universal Optimality of Meta-Induction: Two Complementary Accounts

Gerhard Schurz, Paul Thorn

PII: S0039-3681(19)30128-1
DOI: https://doi.org/10.1016/j.shpsa.2019.11.001
Reference: SHPS 1481

To appear in: *Studies in History and Philosophy of Science*

Received Date: 27 April 2019
Accepted Date: 17 November 2019

Please cite this article as: Gerhard Schurz, Paul Thorn, The Material Theory of Object-Induction and the Universal Optimality of Meta-Induction: Two Complementary Accounts, *Studies in History and Philosophy of Science* (2019), https://doi.org/10.1016/j.shpsa.2019.11.001

Title:

The Material Theory of Object-Induction and the Universal Optimality of Meta-Induction: Two Complementary Accounts

Authors:

Gerhard Schurz and Paul Thorn

E-Mail address of corresponding author: schurz@phil.uni-duesseldorf.de

Internet: http://www.philosophie.hhu.de/schurz

Postal address of both authors:

Heinrich Heine University Duesseldorf

Department of Philosophy, DCLPS

Universitaetsstrasse 1, building no. 24.52,

40225 Duesseldorf, Germany

Abstract:

This paper brings together two accounts of induction that appear to be in opposition: John Norton's material account of induction (2003, 2010, manuscript) and Schurz' account of the universal optimality of meta-induction (2008, 2017, 2019). According to the material account of induction, all reliable rules of 'induction' are local and context-dependent. Here "induction" is understood in the sense of *object-induction*, i.e., induction applied at the object-level of events. In contrast, Schurz' account proceeds from the demonstration that there are universally optimal rules of *meta-induction*, i.e., rules of induction applied at the level of competing methods of prediction, including methods of object-induction. The two accounts are not in opposition; on the contrary, they agree on most questions related to the problem of induction. Beyond this agreement the two accounts are complementary: the material account suffers from a justificational circularity or regress problem that the meta-induction account can solve. On the other hand, the meta-inductive account abstracts from domain-specific aspects of object-induction that are supplied by the material account.

Keywords: object-induction, meta-induction, material account of induction, problem of circularity, universal optimality, no free lunch theorem

The Material Theory of Object-Induction and the Universal Optimality of Meta-Induction: Two Complementary Accounts

Gerhard Schurz and Paul Thorn

Abstract:

This paper brings together two accounts of induction that appear to be in opposition: John Norton's material account of induction (2003, 2010, manuscript) and Schurz' account of the universal optimality of meta-induction (2008b, 2017, 2019). According to the material account of induction, all reliable rules of 'induction' are local and context-dependent. Here 'induction' is understood in the sense of *object-induction*, i.e., induction applied at the object-level of events. In contrast, Schurz' account proceeds from the demonstration that there are universally optimal rules of *meta-induction*, i.e., rules of induction applied at the level of competing methods of prediction, including methods of object-induction. The two accounts are not in opposition; on the contrary, they agree on most questions related to the problem of induction. Beyond this agreement the two accounts are complementary: the material account suffers from a justificational circularity or regress problem that the meta-induction account can solve. On the other hand, the meta-inductive account abstracts from domain-specific aspects of object-induction that are supplied by the material account.

Keywords: object-induction, meta-induction, material account of induction, problem of circularity, universal optimality, no free lunch theorem

1. Introduction

This paper brings together two accounts of induction that appear to be in opposition: John Norton's material account of induction (2003, 2010, manuscript) and Schurz' account of the universal optimality of meta-induction (2008b, 2017, 2019). According to the material account of induction, all reliable rules of 'induction' are local and context-dependent. Here 'induction' is understood in the sense of *object-induction*, i.e., induction applied at the object-level of events. In contrast, Schurz' account proceeds from the demonstration that there are universally optimal rules of *meta-induction*, i.e., rules of induction applied at the level of competing methods of prediction, including methods of object-induction. The two accounts are not in opposition; on the contrary, they agree on most questions related to the problem of induction.

First of all, the two accounts agree in their criticisms of various attempts of establishing universally valid rules of object-induction, ranging from the simple straight rule to contemporary Bayesian accounts to inference to the best explanation (Schurz 2019, ch. 3-5; Norton manuscript, ch. 4-12). Second, the two accounts agree in their diagnosis that David Hume was basically right that there is no generally reliable method of object-induction. Rather, the reliability of a method of object-induction depends on the induction-friendliness of the domain to which it is applied (Norton manuscript, ch. 2, sec. 5), and every attempt to 'prove' a method's reliability without inductive uniformity assumptions ends in a circle or an infinite regress.

It is precisely for the preceding reason that the account of meta-induction does not strive for a universal reliability justification, but for a universal optimality justification. Optimality is an epistemologically weaker epistemic goal than reliability; thus again, a conflict between the two accounts does not arise. The crucial feature which makes optimality accounts feasible is that the optimality claim at the meta-level is not raised in regard to all 'possible' methods of prediction – following from results in computational learning theory (Kelly 1996), such a claim would be irredeemable – but rather to the finite set of all methods of prediction (or object-induction) that are c*ognitively accessible* to the given epistemic agent. Theorems in computational learning theory demonstrate that universal long-run optimality in regard to all accessible prediction methods can indeed be achieved by certain 'clever' strategies of meta-induction. Thus meta-induction does indeed have an *a priori* justification. By itself this justification does not entail anything about the rationality of object-induction: it may well be that we live in a world in which methods other than object-induction – for example clairvoyance – are predictively superior. However, the a priori justification of meta-induction gives us, at least potentially, the following *a posteriori justification of object-induction:* to the extent that object-inductive methods have been much more successful than all accessible non-inductive methods in the past, we are justified by meta-induction to continue favoring object-inductive prediction methods in the future. This argument is no longer circular, because a non-circular justification of meta-induction has been established independently.

In conclusion, the account of meta-induction arrives at the same diagnoses concerning methods of object-induction as the material account of (object-) induction: the reliability of methods of object-induction are inescapably domain-specific and context-dependent. In the following sections, we work out more details of the potential synergy of the two accounts. We also argue that, over and above its attractive features, there is a deficiency in Norton's account of material induction that concerns its justification: the justification of material rules of induction is basically unsolved, and if we take this account literally it leads unavoidably into a circle or infinite regress. At this point, the optimality account of meta-induction can help, because it can terminate the regress in a way that the material account cannot.

2. The Material Account of Induction

As mentioned in the introduction, we agree with Norton's critique of contemporary accounts that seek to establish the universal reliability of inductive rules or methods. For instance, we agree that so-called 'inference to the best explanation' (IBE) is not a uniquely defined method of inference, because "there is no clearly defined relation of explanation that confers special inductive support on some hypotheses or theories"

(Norton manuscript, ch. 9, contents). Moreover, the criteria for 'good' explanations are ambiguous and may be in mutual tension (for example, simplicity versus strength). As argued for in Schurz (2008a), abduction is best understood as a *family* of inference patterns that divide into rather separate kinds (for example, non-creative versus creative abductions) that bear at most a family resemblance to each other.

We also agree with Norton that "any particular inductive inference can fail reliably if we try it in a universe hostile to it. That the universe is hospitable to the inference is a contingent, factual matter" (ibid., ch. 2, contents). Most importantly, we agree that

> "while probabilistic analysis of inductive inference can be very successful in certain domains, it must fail as the universal logic of inductive inference. ... Proofs of the necessity of probabilistic accounts fail since they require assumptions as strong as the result they seek to establish" (ibid., ch. 10, contents).

Several formal results support Norton's arguments. One example is the *no free lunch* (NFL) theorem. Several versions of this theorem have been demonstrated in the area of machine learning (cf. Wolpert 1996). The version most important to the philosophical problem of induction is a generalization of a result of Carnap (1950, pp. 564-566). The theorem is a consequence of the intuition that in the absence of knowledge, every possible world (event sequence) should have the same prior probability and can be expressed as follows (cf. Schurz 2017, theorems 4 and 5):

(1) *No free lunch theorem:* Let P be a uniform prior probability (density) distribution over the set of all infinite sequences of binary events (Carnap's measure $c^\dagger$). A non-clairvoyant prediction method is any function assigning, to a sequence of past observations $(e_1,\ldots,e_n)$ (with $e_i \in \{1,0\}$), a probability that the next event occurs, $P(e_{n+1}=1)$. Then: The P-expected success of every non-clairvoyant prediction method is equal to the expected success of random guessing and of any other non-clairvoyant prediction method, namely 0.5.

A distribution that is uniform over all possible worlds is called a *state-uniform* prior. (1) implies that a state-uniform prior makes all sorts of Bayesian long-run convergence results about posteriors impossible. These convergence results are central to subjective Bayesianism, since they establish a form of intersubjectivity, independently of the particular choice of prior distributions (Earman 1992, 141ff). However, these results presuppose that the prior distribution is continuous (and thus non-dogmatic) over the possible values the frequency limits, while a state-uniform distribution is provably non-continuous, but assigns a probability of 1 to the class of sequences with limiting frequency of 1/2 (Schurz 2017, theorem 5).

These and other results (cf. Schurz 2019, ch. 4) demonstrate, in support of Norton's position, that all probabilistic methods of induction presuppose that the assumed prior distribution satisfies certain principles of inductive uniformity. In reaction to this situation, Norton develops his material account of induction. It starts with the diagnosis that the reliability of any inductive method depends on the contingent uniformity properties of the domain to which induction is applied. Norton calls these uniformity properties "facts", but as we shall see soon, these *facts of uniformity* are general facts stretching into the indefinite future.

The argument that the reliability of inductive inferences relies on a certain fact, namely on the uniformity of nature, has prominently been proposed by John Stuart Mill (1865, III.3.1) and later by Russell (1912). However, in the writings of Mill and Russell, induction was conceived of as a *general* method, being justified by the *universal* fact of nature's uniformity. In contrast, Norton's uniformity account is decidedly *local*, dependent on the domain of application. According to Norton (2003, p. 649), inductive reasoning is not governed by formal and general rules, such as

(2) Some observed As are Bs; therefore all As are Bs,

but rather by local material inferences, such as

(3) Some samples of bismuth melt at 271°C; therefore all samples of bismuth melt at 271°C.

Norton compares the material induction (3) with the analogous induction for a sample of wax: "Some samples of wax melt at 91°C; therefore all of them do". In the latter case, we 'know' (according to Norton) that wax is not a chemical element or a definite chemical compound, but rather a mixture of different compounds with variable melting points. In contrast, bismuth is a chemical element, and we 'know' the following chemical uniformity-fact (ibid,. p. 650):

(4) Samples of the same element agree in their physical and chemical properties.

According to Norton, the reliability of induction over chemical elements is justified by fact (4). At this point, two clarifications are important.

First, we take it that Norton does not assume an externalist account of justification, according to which 'the mere fact of uniformity', independently of whether we will ever know this, justifies the inductive inference (3). In past work, Schurz criticized externalist accounts of justification, arguing that external 'justifications' that are cognitively inaccessible to human beings are of no practical significance (Schurz 2018, sec. 3.2). We do not want to discuss this topic here, but we assume that Norton understands justification in an internalist sense, implying that one must somehow be able to provide evidence and arguments for one's belief in local uniformity claims of the type exemplified by (4).

Second, (4) is obviously not a fact that can be observed, but a *general* fact that stretches into the indefinite future. All that we can know by pure observational knowledge is that *so far* all observed samples of the same chemical elements have agreed in their observed properties. Whether this will continue to hold in the indefinite future, or for properties that until now have not been observed, is a conjecture.

It is presumably true, as Norton asserts, that more or less all object-inductions in science are dependent on background uniformity assumptions that are assumed in some domains but not in others. For example, if for six days the Dow Jones index goes up or stays high, then no reasonable person would inductively infer that the Dow Jones will continue to go up forever or to stay high forever, because we 'know' by experience that the stock market is not uniform in this sense. So far so good. There is, however, an obvious challenge to Norton's approach – the Humean challenge: *how can* we know these facts of uniformity that stretch into the indefinite future? Without presupposing general scientific background 'knowledge', (4) could be false in many ways, not only in fanciful philosophical ways, but even in plausible ways. For instance, why should substances not slowly change their melting points with the evolution of our solar system? Stones and minerals have several properties that slowly change, for example, their isotope ratio; why not their melting points? How do we know that this is not the case? The answer is obvious: the only way to know this is by way of an inductive inference, in which we reason from past experience to the indefinite future.

In conclusion, the fundamental objection to the uniformity justification of induction is that it leads to a circle or an infinite regress, since the general uniformity facts that license inductive inferences must themselves be justified by inductive inferences, which in turn must be justified by other facts of uniformity, and so on. A similar circularity or regress objection has been directed against John Stuart Mill's uniformity account induction. Norton acknowledges this problem; however, he argues that the regress problem is only unsolvable for the formal (or general) account of induction, while for the material (or local) account the regress is neither demonstrably infinite nor demonstrably harmful (Norton 2003, sec. 6). We must admit that this is the part of Norton's account that we find least convincing. First of all, if the justification of every local induction depends on a uniformity-fact, and the justification of every uniformity-fact requires a material induction, then this is a 'perfect' situation for a circularity or an infinite regress, and we cannot see why this problem should be 'a jot' less severe for a material than for a formal account of induction. We are not alone in this view; for similar arguments cf. Kelly (2010) and Worrall (2010).

Moreover, a closer look at Norton's account shows that his claim that all uniformity assumptions are equally local is not tenable. The reason for this is that the uniformity assumptions that justify material inductive inferences become unavoidably *more and more general*. For example, the above inference (3) is justified by the uniformity assumption (4) "Samples of the same element agree in their physical properties". The inductive inference that justifies the uniformity-fact (4) is the following:

(5) So far all observed samples of the same element agreed in their physical properties. Therefore this is generally so.

Now, the inductive uniformity that justifies the reliability of (5) is already entirely domain-unspecific and, thus, general. We propose to express this as the following principle of spatio-temporal invariance:

(6) Physically identical entities that differ only in their location in space and time behave in an identical way.

Of course, principle (6) could be expressed in different ways, but the point is that this principle is domain-unspecific, global and general. The inductive inference that justifies the uniformity-fact (6) would be

(7) So far physically identical entities differing only in their location in space and time behaved in an identical way. Therefore this is generally so.

The uniformity-fact that justifies (7) *would again be (6)*; thus Norton's regress terminates after a few steps in a perfect circle.

A similar analysis applies to Norton's case study of induction over the crystallographic structure of crystalline substances (manuscript, ch. 1; more about this case study is said in sec. 4). According to Norton, Marie Curie inductively inferred that all crystals bariumbromide and radiumbromide are crystallographically isomorphic, because "Curie already knew of the closeness of the chemical properties of barium and radium" (ibid., ch. 1, p. 23). Norton argues that Curie's knowledge of this general fact justified her inductive inference, but of course the general fact can itself only be justified by an inductive inference of the form "So far all samples of crystalline substances agreeing in their chemical properties belong to the same crystallographic system; therefore this is generally so". And obviously, this latter inductive inference can only be warranted by a fact of the highest generality level, namely fact (6), which terminates Norton's regress.

Norton appears to appreciate these difficulties. As a kind of side step, he offers another argument to overcome the circularity problem, the argument of the theory-ladenness of observation. Norton argues that even observation reports about the past, articulated in terms of qualitative properties, involve inductive assumptions (Norton 2003, e.g., p. 668, fn. 9). According to Norton, the singular observation sentence "This ball is red" involves the universal proposition "This ball has the same color as all balls in an infinite class of balls". We do not assume that Norton really wants to assert with his claim that when we report an observation such as "yesterday it was raining" we implicitly predict something about the weather in the future. Thus, in a sense, this point seems to be more a stopgap than a strong argument. Moreover, from a strict semantic viewpoint, the argument is unconvincing. The observation statement "This object is red" does not imply that this object has the same color as infinitely many other objects, but only that this object is perceived to have a certain memorized color-quality that is similar to the color of various other objects *observed in the past*.

This is merely a report about a present experience and its relation to finitely many past experiences, but not an inductive generalization.

In conclusion, the gap in Norton's account is the problem of how facts of uniformity can be inductively justified without entering a vicious circle. This is the point where the account of meta-induction offers help, and is the topic of the next section.

3. A Priori Justification of Meta-Induction: Universal Optimality

In other papers (cf. Schurz 2008b, 2009), Schurz developed a new type of higher order justification for inductive inferences that he calls an *optimality justification*. Optimality justifications do not attempt to 'prove' that a cognitive method (here induction) is reliable − something that, by Hume's arguments, cannot be done − but, rather, that it is *optimal*, i.e., that it is the best that we can do in order to achieve our epistemic goal, which in the case of induction is predictive success.

Reichenbach (1949, sec. 91) was the first philosopher who suggested something like an optimality account: he attempted to demonstrate that induction is the best that we can do for the purpose of predictive success. Reichenbach's attempt failed, because – as pointed out by Skyrms (1975, ch. III.4) – nothing in Reichenbach's best alternative account can exclude the possibility of a clairvoyant that is better in predicting random sequences than an empirical inductivist. More generally, results in formal learning theory show that no prediction method can be universally optimal at the object level, that is, optimal at the task of predicting events in all possible worlds (Kelly 1996, p. 263). In contrast, Schurz' account is focused on the concept of *meta-induction*, i.e., induction applied at the meta-level to a finite set of competing prediction methods.

Meta-induction tracks the success rate of all prediction methods whose predictions are *accessible* and predicts a weighted average of the predictions of those methods that were most successful so far. What Schurz' account attempts to show is that there is a meta-inductive strategy that is predictively optimal among all prediction methods that are (simultaneously) accessible to the epistemic agent. Since the restriction to accessible methods is crucial for the optimality theorem, Schurz and Thorn (2016) call this kind of optimality *access-optimality*. Remarkably, the access-optimality of meta-induction holds in *all* possible worlds, even in radically 'non-uniform' worlds or in 'paranormal' worlds that host perfect clairvoyants.

Technically the account of meta-induction is based on the notion of a prediction game:

(8) A *prediction game* is a pair ((e),Π) consisting of:

(1.) An infinite sequence (e) := (e_1, e_2,…) of events $e_n \in [0,1]$ coded by real numbers ranging between 0 and 1, possibly rounded according to a finite accuracy. For example, (e) may be a sequence of daily weather conditions, football game results or stock values. Each time *n* corresponds to one round of the game.

(2.) A finite set of prediction methods or 'players' Π = {P_1,…,P_m,MI}. In what follows we identify 'methods' with 'players'. In each round, it is the task of each player to predict the next event of the event sequence. 'MI' signifies the meta-inductivist, and the other players are the 'non-MI players' or 'candidate methods'. They may be real-life experts, virtual players implemented by computational algorithms, or even 'clairvoyants' who can see the future in 'para-normal' possible worlds. It is assumed that the predictions of the non-MI players are accessible to the meta-inductivist.

Each prediction game constitutes a *possible world*, or in cognitive science terminology a possible environment. Apart from the above definition, we make no further assumptions about these possible worlds. The sequence of events (e) can be arbitrary: a deterministic sequence, a random sequence or Markov chain, or a 'chaotic' sequence whose finite frequencies don't converge to limits. We also do not assume a fixed list of players − the list of players may vary from world to world, except that it always contains MI, and some *fallback* strategy of MI in situations in which there are no other accessible players. The only restriction concerning the set of non-MI players is that it is *finite*; this restriction will be discussed and relaxed at the end of this section.

The *predictive success rate* of a method P is defined by means of the following chain of definitions:

− $pred_n(P)$ is the prediction of *player* P *for* time *n* which is delivered *at* time *n*−1 (like the events, predictions are coded by real numbers between 0 and 1).
− The deviation of the prediction $pred_n$ from the event e_n is measured by a normalized loss function $loss(pred_n,e_n)$ ranging between 0 and 1.
− The *natural* loss-function is defined as the absolute (linear) distance between prediction and event, $|pred_n - e_n|$; however, results concerning the access-optimality of meta-induction apply for a much larger class of loss functions (see below).
− $score(pred_n,e_n) =_{def} 1- loss(pred_n,e_n)$ is the *score* obtained by prediction $pred_n$ of event e_n (ranging between 0 and 1).
− $abs_n(P) =_{def} \Sigma_{1\leq i\leq n}\, score(pred_i(P),e_i)$ is the *absolute* success achieved by player P until time n (ranging between 0 and n).
− $suc_n(P) =_{def} abs_n(P)/n$ is the *success rate* of player P at time n (ranging between 0 and 1).

The optimality theorem (10) below holds for all *convex* loss functions, which means that the loss of a weighted average of two predictions is not greater than the weighted average of the losses of two predictions. Convex loss functions comprise a large variety of loss functions including all linear, polynomial, and exponential functions of the natural loss function. The variants of the optimality theorem (10) in terms of expected or average success hold for all possible loss functions.

The simplest meta-inductive strategy is called *Imitate-the-best*, which predicts what the presently best non-MI player predicts. It is easy to see that this meta-inductive strategy is not universally access-optimal: its success rate breaks down when it plays against non-MI methods that are *deceivers*, which means that they lower their success rate as soon as their predictions are imitated by the meta-inductivist (cf. Schurz 2008b, sec. 4). A realistic example is the prediction of stock values in a 'bubble economy': Here the prediction that a given stock will yield a high rate of return leads many investors to put their money on this stock and by doing so they cause it to crash.

Nevertheless there is a meta-inductive strategy that is provably universally optimal. This strategy is called *attractivity-weighted meta-induction*; it predicts a weighted average of the predictions of the non-MI players, using their so-called 'attractivities' as weights. Attractivities of non-MI players are always monotonically increasing functions of their success difference compared to that of MI. From the viewpoint of MI, this attractivity is called *regret*, and attractivity-based meta-induction is a variant of regret-based learning. The most efficient definitions in terms of MI's worst case regret are exponential attractivities defined as follows (cf. Cesa-Bianchi and Lugosi 2006, pp. 16f):

(9) Predictions of eMI (short for exponential attractivity-weighted meta-induction):

$$pred_{n+1}(eMI) =_{def} \frac{\sum_{1\le i\le m} at_n(P_i)\cdot pred_{n+1}(P_i)}{\sum_{1\le i\le m} at_n(P_i)} \text{, where}$$

– $at_n(P_i)$ is the attractivity of a player P_i for eMI at time n, defined as

$at_n(P_i) =_{def} e^{\eta\cdot abs_n(P))}$, with $\eta = \sqrt{8\cdot \ln(m)/(n+1)}$.

Because the attractivities of non-MI players are exponential functions of their absolute successes, non-MI players having a lower relative success rate than eMI are gradually forgotten by eMI, because they achieve an exponentially smaller weight than those non-MI players whose success rate is comparable to that of eMI. This 'forgetting feature' is a necessary condition for eMI's access-optimality; it guarantees that eMI's success approximates the success rate of the best non-MI player, even if the best non-MI player is permanently changing. Let '$maxsuc_n$' denote the non-MI-players' maximal success rate at time *n*. Then the following universal optimality theorem for eMI has been proved:[1]

(10) *Theorem:* (universal access-optimality of eMI):
For every prediction game ((e){$P_1,\ldots,P_m$,eMI}) with a convex loss function:
(i) $maxsuc_n - suc_n(eMI) \leq 1.78\cdot\sqrt{\ln(m)/n}$
(ii) eMI is long-run access-optimal: $\limsup_{n\to\infty} (maxsuc_n - suc_n(eMI)) \leq 0$.

According to theorem (10)(ii), attractivity-weighted meta-induction is long-run optimal for *all* possible event sequences and finite sets of (simultaneously accessible) prediction methods. In the short run, attractivity-weighted meta-induction may suffer from a possible loss, compared to the leading player. This possible loss derives from the fact that eMI must base her prediction of the next event on the *past* success rates of the candidate methods, and the hitherto most attractive methods may perform badly in the prediction of the next event. Fortunately theorem (10)(i) states a worst-case upper bound for this loss, which is small if the number of competing methods (m) is not too large compared to the number of rounds (n), and converges to zero when n grows large.

In conclusion, theorem 10 and its variants establish the following a priori justification of meta-induction:

(11) *A priori justification of meta-induction:* In all possible worlds, it is reasonable for an epistemic subject X to apply the strategy eMI to all prediction methods accessible to X, since this can only improve but not worsen X's success in the long

[1] Proof: See Schurz (2019), appendix 12.24, based on theorem 2.2+3 of Cesa.-Bianchi and Lugosi 2006; see also theorem 21.11 in Shalev-Shwartz and Ben-David 2014, pp. 253f. There the tighter worst case bound of $\sqrt{2\cdot\ln(m)/n}+\sqrt{\ln(m)/8\cdot n^2}$ is proved, which implies the bound of (10).

run.

The justification of meta-induction given by (11) is a priori and analytic, since it does not rely upon any assumptions about contingent facts. Moreover, the justification is non-circular, because it does not rest on any inductive inference or assumption of inductive uniformity. Note that claim (11) should not be misunderstood as entailing that the application of eMI *in isolation* is the best epistemic strategy. Rather, it implies that eMI is optimal relative to the given set of accessible candidate methods. Besides this, it is always reasonable *in addition* to try to improve one's candidate set. But this does not constitute an objection against the universal recommendation of applying eMI on top of one's candidate set.

Theorem (10) proves the optimality, but not the *dominance* of attractivity-based meta-induction. As it turns out, there are other meta-inductive methods, different from eMI, that are also access-optimal. However, under most conditions, the exponential version eMI has the best short-run performance (cf. Cesa-Bianchi and Lugosi 2006; Schurz 2019, ch. 6+7; Thorn and Schurz 2019).

An important restriction of theorem (10) is the assumption that the number of competing prediction methods is finite. This restriction is a necessary condition for the proof of the *universal* access-optimality of meta-induction (without it only weaker results are provable). In Schurz (2008b; 2019, sec. 9.2.3) the finiteness restriction is justified by the following fact:

(12) *Fact of cognitive finiteness:* Real epistemic agents are finite beings who can simultaneously access (and compare) only finitely many methods of finite complexity. Therefore the optimality justification of meta-induction is not impaired by the finiteness restriction.

Although the argument from cognitive finiteness is rather strong, it is not the only answer to the 'challenge of infinitely many methods'. An important extension of meta-induction for cognitively finite beings concerns prediction games with finite but *unboundedly growing numbers* of prediction methods. Relevant to this situation, there is a beautiful extension of theorem 10 to games with unboundedly growing numbers players, which holds provided the number of players grows more slowly than an exponential function of the number of rounds (Schurz 2019, sec. 7.3, theorem 7.3).

In conclusion, the optimality justification of meta-induction provides us with a tenable solution to Hume's problem of induction. The *core* of this solution consists in the fact that meta-induction has an indefeasible learning ability: whenever the strategy is confronted with a so far better method, it will learn from it and reproduce its success. This is what makes it optimal – not among all possible methods, but among all accessible prediction methods.

4. A Posteriori Justification of Object-Induction: In Support of the Material Theory

What does theorem 10 and its variants imply for the rationality of object-induction? Without further assumptions nothing, since it is logically as well as metaphysically possible that we live in a world that hosts persons with 'super-natural' abilities – clairvoyants, God-guided fortune tellers or whatever – whose predictive success outperforms the success of ordinary empirical scientists. Of course, the optimality of eMI is not affected by this possibility: in these worlds, meta-inductivists would favor not the predictions of the scientists but those of the clairvoyants.

Nevertheless, conditional on past success rates, the optimality of meta-induction provides us with an a posteriori justification of object-induction, based on the following idea:

(13) As a matter of contingent fact, object-inductive prediction methods were so far much more successful than all accessible non-inductive (object-level) prediction methods. Therefore, it is justified, by meta-induction, to continue favoring object-

inductive prediction methods in the future.

Argument (13) is no longer circular, because a non-circular justification of meta-induction has been established independently. Argument (13) presupposes a contingent premise about the past success rates of inductive compared to non-inductive prediction methods. Although this premise seems to be plausible, a closer look makes clear that the premise needs refinement. Two complications must be considered.

First, 'object-induction' is not just one method, but an unboundedly large *family* of – simple or increasingly refined – methods applied at the level of observed events. Many scientific debates concern the question of which inductive method (e.g., TTB heuristics, multilinear regression, simple or full Bayes estimation) is most appropriate for which domain. It is easy to see that different inductive methods, if applied to the same event sequence, may produce mutually inconsistent predictions. As an example, consider the binary event sequence (0011100111) and two binary prediction methods: M1 always predicts the rounding of the so-far observed frequency to 1 or 0 (the so-called 'maximum rule' of prediction). In contrast, M2 predicts repetitive patterns, in our example the repetitive pattern 00111. Thus while M1 predicts $e_{11} = 1$, M2 predicts $e_{11} = 0$ (for a similar example cf. Norton manuscript, ch. 2, sec. 6). In a situation of this sort, the application of meta-induction to the competing inductive prediction methods is the recommended choice, as it is guaranteed to select an optimal combination of the methods. Of course, sometimes the result of a meta-inductive combination of methods may lead to the conclusion that the odds are equal and one should simply remain agnostic.

Second, there are several domains in which object-inductive prediction methods are not more successful than random guessing, because of the chaotic dynamics or the chance-driven nature of the events in these domains.

In conclusion, our thesis that object-induction was so far predictively more successful than non-inductive methods should be explicated as follows:

(14) *Contingent premise of the a posteriori justification of object-induction:* Until the present time and according to the presently available evidence, object-inductive methods dominated non-inductive methods in the following sense: In many fields some object-inductive method was significantly more successful than every non-inductive method, though in no field was a non-inductive method significantly more successful than all object-inductive methods.

It should be kept in mind that the justification of object-induction based on premise (14) is always relative to the present time and available evidence. As already discussed in section 2, there are many domains in which the superiority of object-inductive prediction methods is not obvious. This latter point brings us back to Norton's material account of object-induction. It is only in domains that are regulated by strong uniformities that the superiority of object-inductive methods over non-inductive methods of prediction will be strong enough that it can convince even skeptical persons. The strength of Norton's material account of object-induction lies in the fact of illuminating the detailed structure of these local uniformities make the success of inductions in science possible. An example of this sort was already presented in sec. 2: the bismuth example of Norton (2003). Here the local uniformity is expressed as the strict (exceptionless) generality (4): All samples of the same element agree in their physical and chemical properties. However, as Norton emphasizes, in most cases the general facts are merely statistical generalization that may have exception. An example is Norton's case study of the crystallographic structures of minerals (manuscript, ch. 1). Inductions generalizing the crystallographic structures of samples (e.g., cubic, octahedral, etc.) have the form

(15) This sample of salt A belongs to crystallographic system B, therefore all samples to salt A belong to crystallographic system B (ibid., 7).

According to the simple strict analyses, the reliability of these inductions is based on the general fact-hypothesis

(16) Each crystalline substance has a single crystallographic form B (Haüy's principle, cf. ibid., p. 19).

However, this fact-hypothesis is false, because there are exceptions: some crystalline m may possess several different crystallographic forms. Norton concludes that the uniformity fact behind inductions over crystallographic systems substances is merely *weak* generality which he expresses as follows (ibid., p. 21):

(17) *Generally* each crystalline substance has a single characteristic crystallographic form.

Generalizations of this form are called *normic generalizations* in the literature, because they express normal-case hypotheses that have exceptions (cf. Schurz 2001); another strategy of expressing these weak generalizations in the literature is by means of *ceteris paribus* laws (Reutlinger et al. 2010, Schurz 2002). Compared to strictly general facts, weakly general facts imply two changes in the formal nature of the induction that are licensed by them. First, the reliability of these inductions is now no longer strict but merely probabilistic, which makes probabilistic considerations more important than it seems according to Norton's discussions of Bayesianism. Second, the inductive inference becomes defeasible by exceptional evidence, which now has to be explicitly excluded in the premises. Therefore the proper formal structure of inductive inferences can now no longer have the simple form that Norton attributes to them in (2003) and (manuscript, ch. 1):

(2) Some observed As are Bs; therefore all As are Bs.

Rather, the extent of the observed exceptions must now be made explicit in the premises, in one of the two following forms:

(2') *All* observed As are Bs (i.e., so far no counterexamples have been observed); therefore all As are Bs,

or

(2") r percent of the observed As are Bs (i.e., so far 1-r% counterexamples have been observed); therefore approximately r% of all As are Bs.

As these considerations make clear, not only the reliability but also the formal structure of local inductions depends on the nature of the local uniformity facts licensing them (in fact, much more complicated forms of object-induction are possible and suited for specific environments; cf. Schurz 2019, ch. 5, ch. 5, sec. 8.3.2).

5. Conclusion

The two accounts of induction have been brought together: John Norton's material account of induction (2003, 2010, manuscript) and Schurz' account of the universal optimality of meta-induction (2008b, 2017, 2019). According to the first account, all reliable rules of object-induction are local and context-dependent. According to the second account, there are universally optimal rules of *meta-induction*. The two accounts are not in opposition but complementary. The material account suffers from a justificational circularity or regress problem that the meta-induction account can solve. On the other hand, the meta-induction account abstracts from domain-specific aspects of object-induction that are supplied by the material account.

References

Carnap, R. (1950). *Logical foundations of probability*. Chicago: Univ. of Chicago Press.

Cesa-Bianchi, N., and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge: Cambridge Univ. Press.

Earman, J. (1992). *Bayes or bust?* Cambridge/Mass.: MIT Press.

Kelly, T (2010). Hume, Norton, and induction without rules. *Philosophy of Science* 77, 754-764.

Mill, J. St. (1865). *System of logic*. London: Longmans, Green & Co.

Norton, J. (2003). A material theory of induction. *Philosophy of Science* 70, 647-670.

Norton, J. (2010). There are no universal rules for induction. *Philosophy of Science* 77, 765-777.

Norton, J. (manuscript). *The material theory of induction*. https://www.pitt.edu/~jdnorton/papers/material_theory/material.html

Reichenbach, H. (1949). *The theory of probability*, Berkeley: University of California Press.

Reutlinger, A., Schurz, G., and Hüttemann, A. (2011). Ceteris paribus laws. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), http://plato.stanford.edu/entries/ceteris-paribus.

Russell, Bertrand. (1912). The problem of induction. In B. Russell, *The problems of philosophy*, 1959 (ch. 6).Oxford: Oxford Univ. Press.

Schurz, G. (2001). What is 'normal'? An evolution-theoretic foundation of normic laws and their relation to statistical normality. *Philosophy of Science* 28, 476-97.

Schurz, G. (2002). Ceteris paribus laws: classification and deconstruction. In Earman, J., Glymour, C. and Mitchell, S. (eds.), *Ceteris paribus laws*, *Erkenntnis* 57/3 (special issue), 351-372.

Schurz, G. (2008a). Patterns of abduction. *Synthese* 164, 201-234.

Schurz, G. 2008b. The meta-inductivist's winning strategy in the prediction game: a new approach to Hume's problem. *Philosophy of Science* 75, 278-305.

Schurz, G. (2017) No free lunch theorem, inductive skepticism, and the optimality of meta-induction. *Philosophy of Science* 84, 825-839.

Schurz, G. (2018). Optimality justifications: new foundations for foundation-oriented epistemology. *Synthese* 195, 3877-3897.

Schurz, G. (2019). *Hume's problem solved: the optimality of meta-induction*. Cambridge/Mass.: MIT Press.

Schurz, G., and Thorn, P. (2016). The revenge of ecological rationality: strategy-selection by meta-Induction within changing environments. *Minds and Machines*

26(1), 31-59.

Shalev-Shwartz, S. and S. Ben-David. (2014. *Understanding machine learning. From theory to algorithms*. New York: Cambridge University Press.

Skyrms, B. (1975). *Choice and chance*. Encinco: Dickenson (4th ed. Wadsworth 2000).

Thorn, P., and Schurz, G. 2019. Meta-inductive prediction based on attractivity weighting: an empirical performance evaluation. *Journal of Mathematical Psychology* 89, 13–30.

Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation* 8/7, 1341-1390.

Worrall, J. (2010). For universal rules, against induction. *Philosophy of Science* 77, 740-753.

Title:

The Material Theory of Object-Induction and the Universal Optimality of Meta-Induction: Two Complementary Accounts

Highlights:

- This paper brings together two accounts of induction that appear to be in opposition: John Norton's material account of induction and Schurz' account of the universal optimality of meta-induction
- Norton's material account is about object-induction (induction applied at the level of events). Norton argues that all reliable rules of object-induction are local and context-dependent. In contrast, Schurz' account is about meta-induction (induction applied at the level of prediction methods). Schurz demonstrates that there are universally optimal rules of meta-induction.
- The two accounts are not in opposition but agree on most questions related to the problem of induction. Beyond this agreement the two accounts are complementary: The material account suffers from a justificational regress problem that the meta-induction account can solve. On the other hand, the meta-inductive account abstracts from domain-specific aspects of object-induction that are supplied by the material account.