



Reichenbach's best alternative account to the problem of induction

Gerhard Schurz¹

Received: 14 November 2020 / Accepted: 15 June 2021 / Published online: 6 July 2021
© The Author(s) 2021

Abstract

In this paper Reichenbach's best alternative account (BAA) to induction is examined. In the first section, three versions of the BAA are distinguished that have been discussed in the literature. The major objections against all three versions are presented. In the second section it is shown by a text analysis that Reichenbach (The theory of probability, University of California Press, California, 1949) argues for all three versions of the BAA and does not sufficiently distinguish between them. In the third section it is explained how Reichenbach's third version of the BAA can be transformed into a provable optimality theorem within the account of meta-induction.

Keywords Problem of induction · Best alternative account · Prediction of frequency limits · Optimality of meta-induction

1 Introduction: three versions of the best alternative account

According to David Hume's skeptical insight, it is impossible to demonstrate that inductions will be successful without making inductive assumptions about the world, which in effect means reasoning in a viciously circular way. However, there exists a different approach to Hume's problem for which we can uphold the hope that it can succeed if it is adequately developed. This approach goes back to Reichenbach's *best alternative account* (BAA). The best alternative account does not try to demonstrate that induction must be successful in the sense of having a high success rate. This is impossible given Hume's insight. Rather, it argues that the expected

This article belongs to the topical collection "All Things Reichenbach", edited by Erik Curiel and Flavia Padovani.

✉ Gerhard Schurz
schurz@hhu.de
<http://www.philosophie.hhu.de/schurz>

¹ Department of Philosophy, Heinrich Heine University Duesseldorf, 40225 Duesseldorf, Germany

success of induction will be maximal among *all* competing methods. In simpler words, Reichenbach tried to show that if any method of prediction works, then the inductive method does. Reichenbach (1938, p. 363) draws the picture of a *fisherman* sailing to the sea and not knowing whether he will find fish (= find regularities). Yet it is clearly more reasonable for him to carry his fishing net (= induction) with him than not, because by doing so he can only win (catch fish) but not lose anything.

Reichenbach's best alternative account has been formulated in three versions. Against all three versions serious objections have been raised.

Version 1 of Reichenbach's BAA: This version attempts to argue that induction is the best that we can do in regard to predictive success. The version has been reconstructed by Salmon (1974, p. 83) by way of the utility matrix in Fig. 1.

According to this matrix, induction is optimal (in every possible world at least as successful as of any other competitor) and even dominant (in some possible words better than any other competitor). The key objection against this first version of the argument is that it simply *fails* to hold in regard to the goal of predictive success. It is easy to conceive worlds in which the success of inductive methods is not better than that of blind guessing, but where there exists a God-guided clairvoyant who successfully predicts the future. A perfect future-teller may have a 100% success in predicting random tossings of a coin, while the scientific inductivist can only have a predictive success of 0.5 in this case. The question mark "[?]" indicates this weak spot in Reichenbach's argument: a clairvoyant may have perfect success in this cell of the matrix.

One may object against this reconstruction that in Reichenbach philosophy of probability and induction, the central goal is the estimation of relative frequencies and not the prediction of events.¹ However, Reichenbach emphasizes that in all applications of probability we are interested in the *prediction* of *finite* frequencies (see Sect. 2). Since finite sequences (or 'samples') of a given fixed length are themselves complex events, the above objection apply likewise to the goal of predicting finite frequencies. A God-guided clairvoyant would be much more successful in the predictions of finite frequencies than an empirical inductivist, whose average error in the prediction of sample frequencies is given by the standard error of the sampling distribution.

Version 2 of Reichenbach's BAA: To avoid the above difficulties, Reichenbach switches to a second version of his argument. In this second version that has also been reconstructed by Salmon (1974, p. 87), the goal of induction is the approximation of the *frequency-limit* of a given property in an infinite sequence of events, by repetitive predictions of this limit in the long run based on the inductive 'straight rule', which projects the frequency observed until present to the conjectured limit. A "uniform world state" is now defined as one in which the events to be predicted possess a frequency limit in the given sequence, while a "non-uniform world state" is understood as a sequence whose events do not possess a limit. With respect to this interpretation of the decision-matrix in Fig. 1, Reichenbach's argument becomes true, but also more or less *trivial*. If the event has a frequency limit p , then the

¹ This objection was raised by a referee.

	World-state:	Uniform	Non-uniform
Method applied:			
Inductive		Success	Failure
Other		Success or Failure	Failure [?]

Fig. 1 Salmon’s reconstruction of the BAA. “[?]” indicates the weak spot

simplest inductive generalization rule, the *straight* rule *must* approximate this limit in the long run, because *by definition* the existence of a limit p means that with $n \rightarrow \infty$ the finite frequencies $freq_n$ of the given event converge to p , $\lim_{n \rightarrow \infty} freq_n = p$. On the other hand, if the event frequencies $freq_n$ do not converge to a frequency limit, then for trivial reasons *no* method can find a limit.

The major objection against the second version of Reichenbach’s BAA is its triviality and its practical irrelevance. Lenz (1974) and Rescher (1980) have provided devastating arguments of this sort. They argue that the method of ‘finding’ provided by the straight rule is not really a method of finding, because after any finite number of observations, however large, our observed frequency could still be maximally distant from the true limit. Thus we never know when we have approximated the limit within non-trivial approximation bounds. *Second* and more importantly, conjecturing an approximate frequency limit is practically insignificant and not the primary goal of inductive inferences (cf. Kading, 1960). Rather, the primary goal is the *prediction* of future events or of the finite frequencies of future events. In this respect, Reichenbach’s account fails, for even if the inductivist performs equally well as a clairvoyant in conjecturing the frequency limit of an event, the clairvoyant could still be overwhelmingly more successful in regard to these predictions.

Version 3 of Reichenbach’s BAA: Reichenbach was well aware of the problems of the previous two versions of his argument. Faced with the triviality of version 2, he reacted by arguing that conjectures about the frequency limit should be interpreted as predictions about frequencies in finite time spans that can be tested after some time (see Sect. 2); thus in the effect, he switched back to version 1. Faced with the difficulties of version 1, Reichenbach made a crucial observation that lead him to the third version of the BAA. He observed that if a successful prophet existed, then this fact would already constitute some uniformity which the inductivist could recognize by applying induction to the success of prediction methods (1949, p. 476). Should it be observed that the prophet is successful, the inductivist would “consult him further”.

At this point, Reichenbach indicates precisely the direction in which the optimal approach to *meta-induction* (induction applied at the level of methods) has been developed later on, which will be briefly described in Sect. 3. Reichenbach himself stops his consideration at this point; he neither shows nor attempts to show that by this observation the inductivist could have an equally high predictive success as the future-teller. This difficulty has been highlighted by Skyrms’ objection against the third version of Reichenbach’s argument. Skyrms (1975, ch. III.4) distinguishes

between the object-level and the meta-level of methods. He argues that Reichenbach has indeed shown that if there exists a successful clairvoyance-method at the object-level, the meta-inductivist can find that out and construct an inductive argument why this clairvoyance-method is successful at the meta-level. However, what Reichenbach did not show, Skyrms objects, is that the meta-inductivist can produce equally successful inductive predictions at the object-level.

Which of the three versions of the BAA account was most close to Reichenbach's own version? This will be the question of the next section. By a text analysis of central quotes we will see that Reichenbach argued simultaneously for all three versions. Faced with their difficulties he switched between these versions without sufficiently distinguishing between the three versions or making the switches explicit.

2 Details and twists in Reichenbach's BBA

In this section we analyze what in our view are the most important passages in Reichenbach's BAA in Reichenbach (1949). In his chapter on the "problem of application" Reichenbach emphasizes that the central goal is the prediction of finite frequencies:

In fact, we are interested only in finite sequences because they will exhaust all the possible observations of a human lifetime or the lifetime of the human race. We wish to find sequences ... converging sufficiently well within that length (348).

... in all practical applications we wish to know the value of the limit before the sequence is completely produced; indeed, all practical use of probability statements consists in the fact that they are applied for the prediction of relative frequencies (351).

In the final chapter "The justification of induction" Reichenbach shifts the emphasis from finite frequencies to (infinite) frequency limits as follows:

Scientific method pursues the aim of predicting the future; in order to construct a precise formulation for this aim we interpret it as meaning that scientific method is intended to find limits of frequency. Classical induction and prediction of individual events are included in the general formulation as the special case that the relative frequency is = 1. (Reichenbach, 1949, p. 474f.).

Having switched the goal from predictive success to approximating the frequency limit (version 2 of the BBA) he draws the following preliminary conclusion:

It has been shown that if the aim of scientific method is attainable, it will be reached by the inductive method. ... if there is a limit of the frequency we

will find it. If there is none, we shall certainly not find one – but then all other methods will break down also.

The answer to Hume's question is thus found ... " (Reichenbach, 1949, p. 475).

Three major points of critique have to be raised against this part of Reichenbach's BAA from a contemporary viewpoint:

Critique 1: Reichenbach correctly identifies the "aim of predicting the future" as the central aim of induction. "To make this precise" he interprets this aim as the aim of finding the frequency limit. However, by this move he doesn't make this aim "precise", but he *changes* it. Success in predictions and success in finding the frequency limit are not generally correlated. For example, an inductivist may perfectly conjecture the right limiting frequency of a coin tossing sequence, but her success in event predictions has an average error of 0.5. Moreover, her success in predictions of the frequencies of length-10-subsequences has an average error of $(0.5^2/10)^{0.5} \approx 0,16$, provided the sequence is random; otherwise the error could be 0.5 even in this case. In contrast, a perfect clairvoyant would have zero error in all of these cases.

Moreover, an important part of the success in predicting events is finding the most specific *reference class* for the predictive target. In contrast, for the success in approximating the frequency limit of a given event type in a given reference sequence, finding most narrow reference classes for the events in that sequence is irrelevant; one just has to apply induction by enumeration. To give an example, for predicting whether or not it will rain the next day in London it is important to conditionalize on the preceding weather situation, but for estimating the long-run average frequency of rainy weather in London this is irrelevant.

Critique 2: Limit 'predictions' are *not* genuine predictions, because the 'frequency limit' is a *theoretical* concept that is not empirically observable, not even approximately within an arbitrary long but fixed time horizon. What we only know is that at *some* (arbitrarily late) time the conjectures provided by the straight rule will have approximated the limit, but we cannot know when this time is reached. There is no future time point at which we can be guaranteed to have approximated the true limit within a given accuracy interval $\pm \epsilon$.

In fact, this seems to be the core problem behind Reichenbach's switches of perspectives. For example, he writes.

The practical geometry of lines of small width and points of small extension can be identified with the ideal geometry to a certain degree of approximation. ... For the same reason we shall not use the concept of a practical limit in the following chapter, but shall carry through all technical analysis with respect to infinite sequences. We know that the results will hold approximately for practical limits (348).

But this is a crucial error. Different from geometrical approximation, there exists no general (sequence-independent) approximation relation between length- n frequencies (freq_n) and the frequency limit p : for arbitrary large n the difference between freq_n and p may still be maximal, i.e. 1.

Critique 3: The third critique is independent from the above problems. Reichenbach writes (in the above quote on p. 474 and the quote below on p. 476) that his method of predicting frequency limits would include the prediction of individual events as a limiting case, but this is not correct. Reichenbach is right that limit 'predictions' include predictions with *certainty*, $p = 1$, but this is not the same as predicting individual events. Also for statistical event sequences—e.g., sequences of binary weather events (rain versus not-rain) with a 70% limiting frequency of rain—one may forecast, instead of a frequency limit, the individual events, either rain or not-rain. An important scoring result is connected with this observation. The so-called *maximum* rule says that one should predict the event with maximal observed frequency. If binary events are coded with 1 (rain) and 0 (not rain), this means that the maximum rule does not predict the observed frequencies, but their *rounding* to 1 or 0. It is well-known that for IID (identically independently distributed) sequences and a *linear* scoring function, the maximum rule is optimal among all competing prediction rules of the form "predict 1 in $r\%$ and 0 in $1-r\%$ of cases". Reichenbach (1938, p. 310f) was aware of this result. Linear scoring scores the success of a prediction proportional to 1 minus its absolute distance from the truth value. In order to reward forecasters who predict their estimated probabilities or limiting frequencies, one has to apply nonlinear scoring methods, for example Brier's (1950) scoring based on a *quadratic* loss function, or other non-linear scoring rules that are subsumed under the term *proper* scoring rules.

Reichenbach was well aware of critiques 1 and 2 concerning the practical irrelevance of infinite limits. To avoid these objections, he proposes to replace the notion of a theoretical limit by the notion of the *practical limit* that he introduced earlier in the book (1949, 347, 447). A binary sequence has a practical limit p if after some chosen finite time Δ the frequency of the event lies within given small bounds $\pm \epsilon$ of p . It is not excluded that the frequencies may fall outside these bounds later, or that the practical limit after time Δ is significantly different from the true limit for $n \rightarrow \infty$. Nevertheless, as Reichenbach maintains, for purposes of application the notion of the practical limit is the only notion of relevance (1949, 448).

Reichenbach assumes that for his notion of a practical limit, the BAA argument would go through in the same way as for the theoretical limit. However, this is not the case, for two reasons:

First, even if both the inductive straight rule, I, finds the practical limit after the given time span Δ , it may be that an alternative method M finds this limit already much *earlier* in the sequence, which would constitute an obvious improvement of M's predictive success over that of I. This predictive advantage is acknowledged by Reichenbach in his quote on p. 351: "we wish to know the value of the limit before the sequence is completely produced".

Second, even if both methods I and M have predicted a wrong practical limit p after time Δ , it may nevertheless be that the sequence has the practical limit p after a somewhat longer time span $\Delta' > \Delta$. If method M finds the practical limit p for Δ' much earlier than I this would still be an advantage for M.

As explained, Reichenbach's preferred inductive method for predicting frequency limits was the above-mentioned inductive straight rule I. Reichenbach (1949,

p. 475f.) argued for the optimality of this method by distinguishing between two classes of 'alternative' methods, as follows:

1. *Provably convergent* methods (Reichenbach's "first class"): The predictions of these methods deviate from the so-far observed frequency freq_n (n = the present time) by an additive n -dependent factor c_n that goes to zero when n approaches ∞ (in the theoretical case) or Δ (in the practical case). Reichenbach argues that these methods are equally good as the induction rule I in reaching the limit, but they are more complicated than I; therefore he prefers I (1949, p. 476). What Reichenbach forgets here is that under certain conditions the alternative methods may approximate the limit faster than I.
2. *Nonprovably convergent* methods (Reichenbach's second class): These methods converge but we cannot prove their convergence. As an example of such a method, Reichenbach assumes a purported *clairvoyant* claiming that he is better than a scientific inductivist in ordinary predictions. Reichenbach writes that "such a method is usually presented in the form of a prediction of individual events. This is included in our theory as the case in which the probability, or the frequency limit, is $p = 1$ " (1949, p. 476). As we have explained above, this is wrong and the prediction of individual events is different from the prediction of certainties. Anyhow, Reichenbach goes on to make the assumption that the purported clairvoyant makes predictions about the (practical) frequency limit of the sequence. Reichenbach was aware of the problem that the possibility of a predictively superior clairvoyant cannot be excluded by logical proofs. This problem led Reichenbach to the third version of his BAA. He writes:

Assume that a clairvoyant asserts that he is able to predict only the probability of an event – to predict the limit of the frequency in a sequence. We shall not be willing to believe him until we have checked his abilities. ... But how can such methods be tested? (Reichenbach, 1949, p. 476).

And he continues:

Obviously, there is only one way – to test these methods by means of the rule of induction. We would ask the soothsayer to predict as much as he could, and see whether his predictions finally converged sufficiently with the frequency observed in the continuation of the sequence. Then we would count his success rate. If the latter were sufficiently high, we would infer by the rule of induction that it would remain so, and thus conclude that the man was an able prophet. If the success rate were low, we would refuse to consult him further. (Reichenbach, 1949, p. 476).

There are three major conclusions to draw from this quote:

First, Reichenbach assumes we can test the soothsayer's prediction of the frequency limit after some finite time; thus he assumes the notion of the practical limit. What Reichenbach suggests here (although he doesn't say it explicitly) is that

predictions of practical frequency-limits should be tested in a repetitive manner after certain time spans (say blocks of 100 or 1000 predictions), by comparing them with the frequencies obtained within these time spans. This implicit idea of Reichenbach will be made precise in the next section. It fits with Reichenbach's chapter 7 on predicting frequency sequences, which for Reichenbach "concerns problems important in practical statistics" and "plays a prominent role in practical applications" (262).

Second, Reichenbach is aware that induction applied at the level of events need not necessarily be the best method of approximating the limit. However, if there are better methods, we can find them by applying induction at the level of methods. This is clearly stated in the following quote:

We thus come to the conclusion that the rule of induction can by no means maintained to be the best method of approximation. But with its help it is possible to find better methods of approximation. (Reichenbach, 1949, p. 447).

Moreover, in the sentence of the preceding quote "If the success rate were low, we would refuse to consult him further", Reichenbach suggests implicitly that as long as the soothsayer has a superior success rate we should "consult him", i.e., use his predictions for our predictions. However, as explained in Sect. 1, Reichenbach doesn't make anything out of this idea. He does not show nor even attempt to show that by performing induction over alternative methods the inductivist can have an optimal success rate. This is the point where the account of meta-induction comes into play, that is briefly described in the next section.

3 The optimality of meta-induction and its application to Reichenbach

By object-induction (abbreviated as OI) we mean methods of induction that are applied at the object-level of events, while by meta-induction (abbreviated as MI) we understand methods applying induction at the meta-level of competing prediction methods. The approach of meta-induction attempts to show what, according to Skyrms (1975), Reichenbach has failed to show: that by observing all accessible prediction methods, the inductivist can achieve an at least equally high predictive success than all competitors method and sometimes even better success. Generally speaking, the problem of Reichenbach's account lies in the fact that it is impossible to demonstrate that a method is optimal in regard to *all* possible prediction methods. This is a lesson of formal learning theory (Putnam, 1965): For every method M one can construct a possible 'demonic' world (event sequence) w in which this method systematically fails, and for every possible world w one can construct a method M^* which is maximally reliable in this world. However, this observation holds only true if the method M is not a meta-inductive method that has cognitive access to M^* , i.e., can observe M^* 's success and use M^* 's predictions. Thus by restricting the optimality claims to methods that are *accessible* to the meta-inductivist, the optimality account may possibly succeed. Even under this restriction there remains the difficulty that in the short run the meta-inductivist may suffer from unavoidable losses

because of the *delay* problem: the meta-inductivist bases his or her prediction of the *next* event on the *past* success rates of the accessible methods, and the hitherto best methods may perform badly in the prediction of the next event. Fortunately, as we shall see, these worst-case losses due to the delay problem may be held small.

The account of meta-induction was developed in Schurz (2008, 2019)² based on mathematical results in machine learning (Cesa-Bianchi and Lugosi 2006). Technically the account is based on the notion of a *prediction game*. This is a pair ((e), Π) consisting of.

- (1) An infinite sequence (e) := (e₁, e₂, ...) of events e_n ∈ [0, 1] coded by real numbers (possibly rounded) in the interval [0, 1], and
- (2) a finite set of accessible prediction methods (or 'players') Π = {M₁, ..., M_m, MI}.

In each round n it is the task of each method to predict the next event (e_{n+1}) of the event sequence. "MI" is the meta-inductivist and the other 'players' M_i are accessible to MI. Each prediction game constitutes a possible world. Apart from the above definition no further assumptions about possible worlds are made: the sequence of events (e) can be arbitrary; its finite frequencies can but need not converge to limits; even 'para-normal' worlds hosting 'clairvoyants' are admitted; the list of players may vary from world to world, except that it always contains MI.

The *predictive success* of a given method is evaluated by a normalized *loss* function that measures the deviation of the prediction pred_n(M) (of a method M) from the event e_n: loss(pred_n(M), e_n) ∈ [0, 1]. The natural (linear) loss-function is defined as the absolute distance |pred_n - e_n|. However, the optimality theorem holds for all *convex* loss functions (meaning that the loss of a weighted average of two predictions is not greater than the weighted average of the losses of two predictions). The corresponding score is defined as 1 minus loss. The *absolute* success achieved by a method M until a time n is defined as its sum-of-scores: Suc_n(M) =_{def} ∑_{1 ≤ i ≤ n} score(pred_i(M), e_i). Moreover, suc_n(M) =_{def} Suc_n(M)/n is the *success rate* of method M at time n, and maxsuc_n is the maximal success rate of all methods at time n.

The simplest meta-inductive strategy is *Imitate-the-best*. It predicts what the presently best non-MI player predicts. It is easy to see that this meta-inductive method cannot be universally access-optimal: Its success rate breaks down when it plays against 'deceiving' methods that lower their success rate as soon as their predictions are imitated (cf. Schurz, 2008, sec. 4). Nevertheless there exists a meta-inductive strategy that is provably universally optimal. This strategy is called *attractivity-weighted meta-induction*, abbreviated as aMI. It predicts a weighted average of the predictions of the accessible methods:

- (1) Predictions of aMI (attractivity-weighted meta-induction):

$$\text{pred}_{n+1}(\text{aMI}) =_{\text{def}} \frac{\sum_{1 \leq i \leq m} w_n(M_i) \cdot \text{pred}_{n+1}(M_i)}{\sum_{1 \leq i \leq m} w_n(M_i)}$$

² For further literature see Thorn and Schurz (2019), Henderson (2018, Sec. 7.3); Sterkenburg (2019), Feldbacher (forthcoming), Douven (forthcoming).

The weights of method M at time n , $w_n(M)$, reflect the attractiveness of M for aMI and can be defined in different ways (Schurz, 2019, Sec. 6.6–6.8). A way that minimizes short-run regrets bounds is $w_n(P) =_{\text{def}} e^{\eta \cdot n \cdot \text{suc}_n(P)}$, with $\eta = \sqrt{8 \cdot \ln(m)/(n+1)}$. Because of the exponentiation, the weights of methods whose success rates are constantly non-maximal are soon negligible for increasing n . The following universal optimality theorem for aMI has been proved (Schurz, 2019, appendix 1 2.24, based on Cesa-Bianchi & Lugosi, 2006):

(2) *Theorem:* (universal access-optimality of aMI):

For every prediction game $((e), \{M_1, \dots, M_m, \text{aMI}\})$ with convex loss function:

- (i) (Short run:) $(\forall n \geq 1:) \text{suc}_n(\text{aMI}) \geq \max \text{suc}_n - 1.78 \sqrt{\ln(m)/n}$.
- (ii) (Long-run:) $\text{suc}_n(\text{aMI})$ converges to $\max \text{suc}_n$ for $n \rightarrow \infty$.

According to theorem (2)(ii), attractivity-weighted meta-induction is universally long-run optimal for arbitrary finite sets of accessible prediction methods. Attractivity-weighted meta-induction cannot be deceived: even if the success rates of the candidate methods are oscillating in adversarial ways, aMI's long run success is guaranteed to be maximal. In the short run, weighted meta-induction may suffer from a possible loss compared to the leading method (a so-called 'regret'). However, theorem (2)(i) states a worst-case upper bound for this loss, which is small (provided n is not small and not smaller than m) and quickly converges to zero when n grows large. In prediction games in which one of the candidate methods dominates its competitors, the weights of aMI converge quickly to 1's and 0's, meaning that aMI imitates the best method.

Theorem (2) applies to *real-valued* prediction games. The events of these games need not be properly real-valued but may also be binary, coded by 0 and 1. However, the predictions are allowed to be proper weighted averages (mixtures) of events. In so-called *discrete* prediction games, such a 'mixing' is impossible or forbidden. Fortunately, there are possibilities of generalizing theorem (2) to discrete prediction games (either by randomization, or by collective meta-induction; see Schurz, 2019, sec. 6.7). A further restriction of theorem (2) is the fixation of a finite pool of candidate methods. In real-life situations new methods may be invented 'on the fly' and the set of candidate methods may grow (Sterkenburg, 2019). An extension of the optimality theorem to sets of candidate methods that grow unboundedly in time is possible under the mild assumption that their number grows less than exponential in time (Schurz, 2019, sec. 7.3).

We finally explain how Reichenbach's idea of predicting practical frequency limits in finite time spans can be implemented within the account of prediction games. We make the following assumptions:

- (i) We assume a binary event sequence $(e_i \in \{0, 1\})$ and divide it in consecutive blocks (time-spans) b_1, b_2, \dots , each having a fixed length k . The frequency of the event $e_i = 1$ within this block b_n is denoted as $\text{freq}(b_n)$ (i.e., $\text{freq}(b_n)$ is the number of 1's in block b_n divided by k). We identify the "new events" of the sequence to which theorem (2) is applied with these frequencies; thus the new

sequence is the frequency sequence $(\text{freq}(b_1), \text{freq}(b_2), \dots)$. This sequence is what Reichenbach calls a "frequency sequence counted in sections" (261).

- (ii) At the beginning of each block b_n , the event frequency $\text{freq}(b_n)$ within this block is predicted. This frequency is nothing but Reichenbach's 'practical limit' after a time span of k time units.
- (iii) At the end of each block b_n , the deviation of the prediction of $\text{freq}(b_n)$ from the true value of $\text{freq}(b_n)$ is scored, for each method of the prediction game.

A prediction game satisfying these three assumptions is called a *Reichenbach game*. Theorem (2) applies to such a game in exactly the same way and we obtain the result that the meta-inductivist's success rate in a Reichenbach game is universally optimal in regard to all accessible methods of predicting practical frequency limits.

Two variations of this version of a Reichenbach game are possible. A first variation would be what Reichenbach calls "a frequency sequence counted in overlapping segments" (261). In this case each block b_n of length k would start with event number n ; thus at every time point a prediction of the frequency of 1 s in the next k events is made. A second variation of a Reichenbach game would be obtained if the blocks are non-overlapping, but the frequency $\text{freq}(b_n)$ to be predicted is not the frequency of 1 s within block n , but the frequency of 1 s from the start until the end of block b_n , i.e., in the first $k \cdot n$ events. The meta-inductive optimality theorem also applies to these two variations of a Reichenbach game.

4 Conclusion

The examination of Reichenbach's best alternative account to induction has led to the following results. In the first section, three versions of the BAA have been distinguished. All three versions seem to be refuted by pertinent objections raised in the literature. In the second section it has been shown by a text analysis that Reichenbach (1949) simultaneously argues for all three versions of the BAA. He switches between the three versions and does not sufficiently distinguish between them. In the third section it is explained how Reichenbach's third version of the BAA can be transformed into a provable optimality theorem within the account of meta-induction.

Acknowledgements This work was supported by the DFG (Deutsche Forschungsgemeinschaft), priority program 1516, "New Frameworks for Rationality". For important discussion I am indebted to Miklós Rédei, Tom Sterkenburg, Clark Glymour and Erik Curiel.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the

material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Brier, G. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78, 1–3.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.
- Douven, I. (forthcoming). Explaining the success of induction. *British Journal for the Philosophy of Science*.
- Feldbacher-Escamilla, C. (forthcoming). *Epistemic engineering*. To appear with Rowman & Littlefield International, Lanham.
- Henderson, L. (2018). The problem of induction. In *Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), <https://plato.stanford.edu/archives/sum2018/entries/induction-problem/>.
- Kading, D. (1960). Concerning Mr. Feigl's "vindication of induction." *Philosophy of Science*, 27, 405–407.
- Lenz, J. (1974). Problems for the practicalist's justification of induction. In R. Swinburne (Ed.), *The justification of induction* (pp. 98–101). Oxford University Press.
- Putnam, H. (1965). Trial and error predicates and a solution to a problem of Mostowski. *Journal of Symbolic Logic*, 30, 49–57.
- Reichenbach, H. (1938). *Experience and prediction*. University of Chicago Press.
- Reichenbach, H. (1949). *The theory of probability*. California: University of California Press.
- Rescher, N. (1980). *Induction*. University of Pittsburgh Press.
- Salmon, W. (1974). The pragmatic justification of induction. In R. Swinburne (Ed.), *The justification of induction* (pp. 85–97). Oxford University Press.
- Schurz, G. (2008). 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. (2019). *Hume's problem solved: The optimality of meta-induction*. MIT Press.
- Skyrms, B. (1975). *Choice and chance*. Encinco: Dickenson (4th ed. Wadsworth 2000).
- Sterkenburg, T. (2019). Meta-inductive justification of induction. The Pool of Strategies. *Philosophy of Science*, 86, 981–992.
- Thorn, P., & Schurz, G. (2019). Meta-inductive prediction based on attractivity weighting: An empirical performance evaluation. *Journal of Mathematical Psychology*, 89, 13–30.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.