A Priori Advantages of Meta-Induction and the No Free Lunch Theorem: A Contradiction?

Gerhard Schurz(✉) and Paul Thorn

DCLPS, Heinrich Heine University Düesseldorf, Düesseldorf, Germany
schurz@phil.hhu.de

Abstract. Recently a new account to the problem of induction has been developed [1], based on a priori advantages of regret-weighted meta-induction (RW) in online learning [2]. The claimed a priori advantages seem to contradict the no free lunch (NFL) theorem, which asserts that relative to a state-uniform prior distribution (SUPD) over possible worlds all (non-clairvoyant) prediction methods have the same expected predictive success. In this paper we propose a solution to this problem based on four novel results:

– RW enjoys free lunches, i.e., its predictive long-run success dominates that of other prediction strategies.
– Yet the NFL theorem applies to online prediction tasks provided the prior distribution is a SUPD.
– The SUPD is maximally induction-hostile and assigns a probability of zero to all possible worlds in which RW enjoys free lunches. This dissolves the apparent conflict with the NFL.
– The a priori advantages of RW can be demonstrated even under the assumption of a SUPD. Further advantages become apparent when a frequency-uniform distribution is considered.

Keywords: Problem of induction · No free lunch theorem · Online prediction under expert advice · Regret-weighted meta-induction

1 Introduction: The NFL Theorem and Hume's Problem of Induction

How can inductive inferences be rationally justified, in the sense of being reliable or at least preferable to non-inductive inferences? This is the problem of induction raised by the philosopher David Hume 250 years ago. Hume showed that all standard methods of justification fail when applied to the task of justifying induction. He concluded that induction has no rational justification at all.

The no free lunch theorem (NFL) expresses a deepening of Hume's inductive skepticism. In this paper we consider the NFL theorem in application to

This work was supported by the DFG (Deutsche Forschungsgemeinschaft), SPP 1516. For valuable help we are indebted to Ronald Ortner.

G. Kern-Isberner et al. (Eds.): KI 2017, LNAI 10505, pp. 236–248, 2017.
DOI: 10.1007/978-3-319-67190-1_18

online prediction tasks. A number of variants of the NFL theorem have been formulated (cf. [3–7]); the most general formulation is found in [8]. Wolpert's NFL theorem comes in a weak and a strong version. Since the strong version rests on unrealistic assumptions about the loss function, we focus in this paper on the weak NFL theorem. It says that the probabilistically expected success of any (non-clairvoyant) prediction method is equal to the expected success of random guessing or any other prediction method, provided one assumes (a) a *state-uniform* prior probability distribution (abbreviated SUPD) i.e., one that is uniform over all possible event sequences, and (b) a *weakly homogeneous* loss function (see below).

Does the NFL theorem undermine the project of learning theory? A standard defense of learning theorists against the NFL challenge maintains that one should not compute the expected success of learning strategies by means of a SUPD. Rather one should compute expected success using the (conjectured) *actual* distribution of the possible states of our environment, and 'according to our evidence' the latter distribution is clearly not uniform.[1] We argue that this line of defense against the NFL challenge does not work, because our beliefs about the actual distribution of possible states of our environment are themselves based on an inductive inference. Thus, this argument commits the *fallacy of circularity*. A general argument demonstrating the unacceptability of circular justifications runs as follows: If we accept the inductive justification of induction ("inductions were successful in the past, whence, by induction, they will be successful in the future"), then – on pain of inconsistency – we must also accept the anti-inductive justification of anti-induction ("anti-inductions were not successful in the past, whence by anti-induction they will be successful in the future").

For a robust defense of inductive learning methods against the NFL challenge a better argument is needed; one that does not presuppose what must be proved. Recently, a non-circular response to the problem of induction has been proposed, based on a priori advantages of regret-based meta-induction (in short: RW) in online learning. In Sects. 2 and 3 these results are presented and confronted with a version of the weak NFL theorem that applies to iterated prediction tasks in online learning. Thereafter the apparent contradiction is analyzed and dissolved, from the long-run (Sect. 4) and short-run perspectives (Sect. 5). Our analysis leads to four novel results that are summarized in the conclusion (Sect. 6).

2 Regret-Based Meta-Induction

In the area of *regret-based learning*, theoretical results concerning the vanishing long-run regrets of certain meta-strategies of prediction have been developed that hold universally, i.e., for strictly all possible event sequences, independently from any assumed probability distribution [2]. Although labeled as "online learning under expert advice" these results characterize the performance of strategies of *meta-learning*, inasmuch as a forecaster which we call the "meta-inductivist"

[1] Cf. [6, Sect. 4] and [7, Sect. 3], citing statements from a 1994 e-mail discussion.

tracks the past success rates of accessible prediction methods ("experts") and utilizes that information in constructing an improved prediction strategy. Since the meta-inductivist predicts future events based on past success rates, short-run regrets (compared to the best method) are unavoidable. However, in the long run the regret-weighted meta-inductivist is guaranteed to predict at least as accurately as the best accessible prediction method, even in circumstances of non-convergent success rates of the independent methods. A standard label for this property is "Hannan-consistency" [2, p. 70]. Schurz and Thorn [9] argue that it is preferable to call this property *access-optimality*, because

– it expresses a long-run optimality result restricted to *accessible* methods, and
– this label is in line with standard game-theoretical terminology of "optimality" and "dominance"; results concerning *access-dominance* are stated below.

The proposed solution to the problem of induction developed in [1,10] works as follows: The meta-strategy RW has an 'a priori' justification, because in the long run it is recommendable in every possible environment to apply this meta-strategy on top of all prediction methods accessible to the epistemic agent. Following [1] we explicate this result within the framework of prediction games.

Definition 1 (Prediction game). *A prediction game is a pair* $((e), \Pi)$ *consisting of:*

(1) An infinite sequence $(e) := (e_1, e_2, \ldots)$ *of events* e_n *coded by real numbers between 0 and 1, possibly rounded according to a finite accuracy. In what follows* $\mathcal{V} \subseteq [0,1]$ *denotes the value space of possible events* $e_n \in \mathcal{V}$*. Each time* n *corresponds to one round of the game.*
(2) A finite set of prediction methods (or 'players') $\Pi = \{O_1, \ldots, O_m, M_1, \ldots, M_k\}$ *whose task, in each round* n*, is to predict the next event* e_{n+1} *of the event sequence. Methods are of two sorts, independent 'object-level' methods* $O_1, \ldots, O_m$ *(algorithms or experts) who base their predictions on the observed events, and dependent 'meta-level' methods* $M_1, \ldots, M_k$ *who base their predictions on those of the independent methods in dependence on their success (this is meant by the* O_i*'s 'being accessible' to the* M_j*'s).*

An example of (e) could be a sequence of daily weather conditions. In what follows the variable 'X' ranges over arbitrary prediction methods. We use the following notions:

– $p_n(X)$ is the prediction of method X *for* time n delivered *at* time $n-1$.
– The distance of the prediction p_n from the event e_n is measured by a normalized loss function, $\ell(p_n, e_n) \in [0,1]$.
– The *natural* loss-function is defined as the absolute distance between prediction and event, $|p_n - e_n|$. The theoretical results below apply to a much larger class, namely to all loss functions that are *convex* in the argument p_n.
– $s(p_n, e_n) := 1 - \ell(p_n, e_n)$ is the *score* obtained by prediction p_n of event e_n.
– $abs_n(X) := \sum_{i=1}^{n} s(p_i(X), e_i)$ is the *absolute success* achieved by method X until time n.

- $suc_n(X) := abs_n(X)/n$ is the *success rate* of method X at time n.
- $maxsuc_n$ is the maximal success rate of the independent methods at time n.

The simplest meta-inductive strategy is *Imitate-the-best*, abbreviated ITB, which, in each round n, imitates the prediction of the independent method with maximal success at time n. ITB fails to be universally access-optimal: Its success rate breaks down when it imitates *adversarial* methods, who return inaccurate predictions as soon as their predictions are imitated by ITB [1, Sect. 4].

The strategy of *regret-weighted meta-induction* comes in several versions. Its simplest version is abbreviated as RW and defined as follows (where $O_1, \ldots, O_m$ are the independent methods of the prediction game):

Definition 2 (Regret-weighted meta-induction)

(i) The absolute regret of RW with respect to independent method O_i at time n is defined as $Reg_n(O_i) := abs_n(O_i) - abs_n(RW)$ and the relative regret as $reg_n(O_i) := Reg_n(O_i)/n$.

(ii) Where $w_n(O_i) := max(Reg_n(O_i), 0)$, the predictions of RW are defined as

$$p_{n+1}(RW) := \frac{\sum_{i=1}^{m} w_n(O_i) \cdot p_{n+1}(O_i)}{\sum_{i=1}^{m} w_n(O_i)}$$

as long as $n > 0$ and the denominator is positive; else $p_{n+1}(RW) = 0.5$

RW is identical with the polynomially weighted forecaster $\mathrm{F_p}$ described in [2, p. 12] with parameter p set to 2.

Theorem 1 (Universal access-optimality of RW). (Cesa-Bianchi and Lugosi 2006, Corollary 2.1)

For every prediction game $((e), \Pi)$ with $RW \in \Pi$ the following holds:
(1.1) (Short run:) $(\forall n \geq 1)$ $suc_n(RW) \geq maxsuc_n - \sqrt{\frac{m}{n}}$.
(1.2) (Long-run:) $limsup_{n \to \infty}(maxsuc_n - suc_n(RW)) = 0$.

In the short run, RW may suffer from a possible regret. According to Theorem 1, RW's relative regret is upper-bounded by $\sqrt{\frac{m}{n}}$ and converges to zero when n grows large, or it oscillates endlessly but with a limsup converging to zero.

An improvement of RW is possible with help of so-called exponential weights. The weights of exponential regret-based meta-induction, abbreviated ERW, are defined as: $w_n(X) := e^{\sqrt{(8 \cdot ln(m)/n)} \cdot Reg_n(X)}$. If ERW's predictions are defined as in Definition 2(ii) but with help of exponential weights, then one can prove that ERW's short-run regret is upper-bounded by $1.77 \cdot \sqrt{ln(m)/n}$ [2, Theorem 2.3]. This is a significant improvement, but in regard to the NFL theorem the difference between RW and ERW is negligible: their long-run advantage is identical and their performance difference in the simulations presented in Sect. 5 turned out to be minor. Therefore we concentrate our investigation on RW.

Even if the events are binary, RW's predictions are real-valued, because proper weighted averages of 0 s and 1 s are real-valued. Thus the predictions are assumed to be elements of a value space $\mathcal{V}_p \subseteq [0,1]$ that may extend the space of event values: $\mathcal{V}$ ($\mathcal{V} \subseteq \mathcal{V}_p$).

What stands in apparent conflict with the NFL theorem is not the access-optimality of RW but rather its *access-dominance*, that is, the fact that RW performs at least as well and sometimes better than other accessible methods. By definition, a meta-method M *dominates* another method X (in the long run) iff (i) there is no prediction game $((e), \Pi)$ with $\{X, M\} \subseteq \Pi$ and $limsup_{n\to\infty}(suc_n(X) - suc_n(M)) > 0$, but there is a prediction game $((e)', \Pi')$ with $\{X, M\} \subseteq \Pi'$ and $limsup_{n\to\infty}(suc_n(M) - suc_n(X)) > 0$; this implies that X is not access-optimal. Theorem 1 asserts the access-optimality but not the dominance of regret-based meta-induction. Since there are other methods, different from RW, that are likewise long-run optimal (such as ERW mentioned above), RW cannot be universally access-dominant. However, the following restricted dominance result for RW can be derived from Theorem 1.

Theorem 2 (access-dominance for RW)

(2.1) RW dominates every accessible prediction method X (in the long run) that is not universally access-optimal.

(2.2) Not universally access-optimal in the long run are (a) all independent (non-clairvoyant) methods, and (b) among meta-strategies, for example, (b1) all one-favorite methods (who at each time point imitate the prediction of one independent method) and (b2) success-weighting, which identifies weights with success rates (also called "Franklin's rule" [11, p. 83]).

Proof. Theorem (2.1) is an immediate consequence of Theorem 1 and the definitions of "access-optimality" and "-dominance".

Proof of Theorem (2.2)(a): Let O be an independent method and (e') an O-adversarial event sequence defined as follows: $e'_1 = 0.5$, and $e'_{n+1} = 1$ if $pred_{n+1}(O) < 0.5$; else $e'_{n+1} = 0$. The predictions of the perfect (e')-forecaster O′ are identified with the so-defined sequence, i.e., $p_n(O') = e'_n$. In the prediction game $((e'), \{O, O', RW\})$ the success rate of O can never exceed 1/2, that of O′ is always 1 and that of RW converges to 1 (by Theorem 1). So O is not universally access-optimal.

The proof of Theorem (2.2)(b1) is found in [1, Sect. 4] and that of (2.2)(b2) in [9, Sect. 7]. □

Theorem (2.1 + 2) entails that in the long run there are "free lunches" for regret-based meta-induction in the sense that there are prediction methods X and event sequences (e) for which RW's long run success is strictly greater than that of X without there being any 'compensating' event sequences (e') in which RW's long-run success is smaller than that of X. This apparent conflict with the NFL folklore is investigated in the next sections.

3 NFL Theorems for Prediction Games

It is not straightforward to apply the NFL theorems to regret-based online learning. First of all, the RW account is more general than the NFL framework as the results of the former account hold even if clairvoyant methods are admitted – these are prediction functions that may have future events as input. However, regret-based meta-induction should not only be attractive for those who consider paranormal worlds as possible. Thus in what follows we take the *non-clairvoyance* assumption of the NFL theorems [8, p. 1380] as granted.

Two further possible hindrances of applying the NFL framework to regret-based online learning are treated as follows:

- Regret-based learning is defined for meta-strategies, while the NFL framework applies to arbitrary prediction methods (defined as computable functions from past event sequences into the next event). But every finite combination of a *fixed* set of independent prediction methods is itself a defined prediction method. Thus the NFL framework equally applies to prediction meta-strategies, given that they are applied to an (arbitrary but) fixed set of independent methods. This assumption will be made in the following.
- Online learning consists of a (possibly infinite) *iteration* of one-shot learning tasks in which the test item of round n is added to the training set of round $n+1$. For this reason the NFL theorems are only applicable if one assumes a SUPD (see below).

The strong version of Wolpert's NFL theorem presupposes that the loss function is *homogeneous* [8, p. 1349], which means by definition that for every possible loss value c, the number of possible event values $e \in \mathcal{V}$ for which a given prediction leads to a loss of c is the same for all possible predictions. This requirement is overly strong; it is satisfied for prediction games with binary events and the zero-one loss function $loss_{1-0}$, which has only two possible loss values: $loss_{1-0}(p, e) = 0$ if $p = e$ and $loss_{1-0}(p, e) = 1$ if $p \neq e$. As soon as real-valued predictions are allowed, a reasonable loss function will assign a loss different from 0 or 1 to predictions different from 0 or 1. Such a loss function is no longer homogeneous. So the strong NFL theorem does not apply to RW or any other real-valued prediction method. Note that real-valued predictions not only make sense in application to real-valued events but also to binary or discrete events, by predicting their conjectured probabilities. Only a weak version of the NFL theorem holds for prediction games with binary events and real-valued predictions, provided the loss function is weakly homogeneous:

Definition 3 (Weakly homogeneous loss function). [2] *A loss function is weakly homogeneous iff for each possible prediction the sum of losses over all possible events is the same, or formally, iff* $\forall p \in \mathcal{V}_p$*:* $\sum_{e \in \mathcal{V}} \ell(p, e) = c^\star$ *(where* $c^\star$ *is a constant).*

For binary games with real-valued predictions and natural loss function the condition of Definition 3 is satisfied, since for every prediction $p \in [0, 1]$, $\ell(p, 1) + \ell(p, 0) = 1 - p + p = 1$. Under this assumption the following weak NFL theorem holds for the probabilistic expectation value (Exp_P) of the success rate of a

[2] [8] mentions the weak no free lunch theorem in a small paragraph on p. 1354; for our purpose this NFL theorem is the most important one.

prediction method X, where "(e_{1-n})" abbreviates "$(e_1, \ldots, e_n)$" and $\mathcal{V}(C) = \{\ell(p,e) : p \in \mathcal{V}_p, e \in \mathcal{V}\}$ is the set of possible loss values:

Theorem 3 (Weak NFL theorem for prediction games). *Given a state-uniform P-distribution over the space of event sequences with r possible event values and a weakly homogeneous loss function, the following holds for every (non-clairvoyant) prediction method X and* $n \geq 0$*:*

The expectation value of X's success rate after an arbitrary number of rounds is $1 - \frac{c^\star}{r}$*, or formally,* $Exp_P(suc_n(X)) := \sum_{c \in \mathcal{V}(C)} c \cdot P(suc_n(X) = c) = 1 - \frac{c^\star}{r}$.

Proof. First we prove the following.

Lemma: For every prediction method X, the expectation value of X's loss in the prediction of the 'next' event equals c*/r, conditional on every possible sequence of 'past' events, or formally:

$$Exp_P(\ell(p_{n+1}, e_{n+1}) \,|\, (e_{1-n})) := \sum_{c \in \mathcal{V}(C)} c \cdot P(\ell(p_{n+1}, e_{n+1}) = c \,|\, (e_{1-n})) = c^\star/r.$$

Proof of lemma: As in [8] we allow that prediction methods are probabilistic, i.e., deliver predictions conditional on past events with certain probabilities $P(p_{n+1} \,|\, (e_{1-n}))$. First we compute the conditional probability of a particular loss value c. By probability theory it holds for all n $\geq$ 0:

$P(\ell(p_{n+1}, e_{n+1}) = c \,|\, (e_{1-n})) = \sum_{p_{n+1} \in \mathcal{V}_p} \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c) \cdot P(p_{n+1}, e_{n+1}) \,|\, (e_{1-n}))$, where "$\delta$" is the Kronecker symbol. By probability theory we obtain $= \sum_{p_{n+1} \in \mathcal{V}_p} \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c) \cdot P(p_{n+1} \,|\, (e_{1-n}), e_{n+1}) \cdot P(e_{n+1}|(e_{1-n}))$, which gives us by non-clairvoyance $= \sum_{p_{n+1} \in \mathcal{V}_p} \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c) \cdot P(p_{n+1} \,|\, (e_{1-n})) \cdot P(e_{n+1}|(e_{1-n}))$, and by rearranging terms $= \sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1} \,|\, (e_{1-n})) \cdot \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c) \cdot P(e_{n+1}|(e_{1-n}))$, and finally by the state-uniformity of P (*) $= \sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1}|(e_{1-n})) \cdot (1/r) \cdot \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c)$. Next we compute the expectation value: $Exp_P(\ell(p_{n+1}, e_{n+1}) \,|\, (e_{1-n})) := \sum_{c \in \mathcal{V}(C)} c \cdot P(\ell(p_{n+1}, e_{n+1}) = c \,|\, (e_{1-n}))$ from which we get by the result in line (*) $= \sum_{c \in \mathcal{V}(C)} c \cdot \sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1} \,|\, (e_{1-n})) \cdot (1/r) \cdot \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c)$ and by rearranging terms $= \sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1} \,|\, (e_{1-n})) \cdot (1/r) \cdot \sum_{c \in \mathcal{V}(C)} c \cdot \sum_{e_{n+1} \in \mathcal{V}} \delta (\ell(p_{n+1}, e_{n+1}), c)$.

Note that "$\sum_{c \in \mathcal{V}(C)} c \cdot \sum_{e_{n+1} \in \mathcal{V}} \delta(\ell(p_{n+1}, e_{n+1}), c)$" is nothing but the sum of p_{n+1}'s loss values for all possible events, i.e., $\sum_{e_{n+1} \in \mathcal{V}} \ell(p_{n+1}, e_{n+1})$. So, by the weak homogeneity of the loss function, we continue as follows: $= \sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1} \,|\, (e_{1-n})) \cdot (1/r) \cdot c^\star = c^\star/r$ (since $\sum_{p_{n+1} \in \mathcal{V}_p} P(p_{n+1} \,|\, (e_{1-n})) = 1$). *(End of proof of lemma.)*

The expectation value of X's success rate is the expectation value of the sum of X's scores divided by n. Since the result of the lemma holds for every round n, the additivity of expectation values $(Exp_P(X_1 + X_2) = Exp_P(X_1) + Exp_P(X_2))$ entails that $Exp_P(suc_n(X)) = n \cdot (1 - (c^\star/r))/n = 1 - (c^\star/r)$. $\square$

The SUPD is a *necessary* condition of the application of the NFL theorem to prediction games, because its proof presupposes that the P-distribution over $\mathcal{V}$ is uniform *conditional* on every possible past sequence. There are generalizations of NFL theorems for one-shot learning procedures to certain non-uniform P-distributions [6], but they are not valid for prediction games.

For prediction games with real-valued events, convex loss functions are not even weakly homogeneous, although certain restricted NFL theorems can be demonstrated [12]. However, in this paper we focus on prediction games with *binary* events, to which the weak NFL theorem applies, because here the apparent conflict of this theorem with RW's access-dominance is most vivid.

4 Meta-Induction and NFL: The Long-Run Perspective

Is there a contradiction between the weak NFL theorem and the existence of free lunches for RW meta-induction? In regard to the *long run* perspective our answer can be summarized as follows: No, the contradiction is only apparent. According to Theorem 2 (Sect. 2) there are RW-accessible methods whose long-run success rate is strictly smaller than that of RW in some world states and never greater than that of RW in any world state. Let us call these methods X_{inf} (for "inferior"). Nevertheless the state-uniform expectation values of the success rates of RW and X_{inf} are equal, because the state-uniform distribution that Wolpert assumes assigns a probability of zero to all worlds in which RW dominates X_{inf}; so these worlds do not affect the probabilistic expectation value.

Wolpert seems to assume that the state-uniform prior distribution is epistemically privileged. Reasonable doubts can be raised here, inasmuch as a well-known result in probability theory tells us that the state-uniform distribution is the most induction-hostile prior distribution one can imagine:

Theorem 4 (Induction-hostile uniformity). [13, pp. 564–566], [14, pp. 64–66]: *Assume the probability density distribution D_P is uniform over the space of all infinite binary event sequences $\{0,1\}^\omega$. Then $P(e_{n+1} = 1 \,|\, (e_{1-n})) = 1/2$, for every possible next event e_{n+1} and sequence of past events (e_{1-n}). Thus P satisfies the properties of a random IID-distribution over $\{0,1\}$, whence inductive learning from experience is impossible.*

Theorem 4 implies that a proponent of a state-uniform distribution believes with probability 1 that the event sequence to be predicted is an IID random sequence, i.e., (a) it consists of mutually independent events with a limiting frequency of 0.5, and (b) it is non-computable. Condition (a) follows from Theorem 4 and condition (b) from the fact that there are uncountably many sequences, but only countably many computable ones. However, the sequences for which a non-clairvoyant prediction method can be better than random guessing are precisely those that do not fall into the intersection of classes (a) or (b). Summarizing, according to a state-uniform prior distribution we are a priori certain that the world is completely irregular so that no inductive or other intelligent prediction method can have more than random success.

If every sequence $(e) \in \{0,1\}^\omega$ is represented by a real number $r \in [0,1]$ in binary representation, then the state-uniform *density* distribution D_P is uniform over the interval $[0,1]$. Yet, if the same density is distributed over the space of statistical hypotheses H_r, asserting that the limiting frequency of 1s in (e) is r (for $r \in [0,1]$), it becomes maximally dogmatic, being concentrated over the point $r = 1/2$: $D(H_q) = 0$ for $q \neq r$ and $D(H_q) = \infty$ for $q = r$.

According to a second well-known result in probability theory, a prior distribution that is not state-uniform but *frequency-uniform*, i.e., uniform over all possible frequency limits $r \in [0,1]$ of binary sequences, is highly induction-friendly. Such a distribution validates Laplace's rule of induction, $P(e_{n+1} = 1 \,|\, freq_n(1) = \frac{k}{n}) = \frac{k+1}{n+2}$. Solomonoff [15, Sect. 4.1] has proved that a distribution is frequency-uniform iff the probability it assigns to sequences decreases exponentially with their algorithmic complexity.

Which prior distributions are more 'natural', state-uniform ones or frequency-uniform ones? In our eyes, this question has no objective answer. It is a great advantage of the optimality of meta-induction that it holds regardless of any assumed prior probability distribution. For a frequency-uniform prior distribution the probability of world-states in which meta-induction dominates random guessing in the long run is one. For a state-uniform prior the probability of world-states in which meta-induction dominates random guessing in the long run is zero. Nevertheless there are (uncountably) many such world-states and we should certainly not exclude these induction-friendly world-states from the start by assigning a probability of zero to them. This consideration gives us the following minimal acceptance criterion for prior distributions: They should assign a positive (even if small) probability to those world-states in which access-dominant prediction methods enjoy their free lunches.

5 Meta-Induction and NFL: The Short-Run Perspective

For finite sequences, the strict dominance of RW fails since the advantage of RW meta-induction comes at a certain regret. Is meta-induction still advantageous over the space of all short-run sequences? This question is addressed in this section.

Table 1 presents the result of a simulation of all possible binary prediction games with a length of 20 rounds. The considered independent methods are

- majority-induction "M-I", predicting the event that so far has been in the majority, or formally, $p_{n+1}(\text{M-I}) = 1/0.5/0$ iff $freq_n(1) > / = / < 0.5$,
- majority anti-induction "M-AI", predicting the opposite of M-I, i.e., $p_{n+1}(\text{M-AI}) = 1/0.5/0$ iff $freq_n(1) < / = / > 0.5$, and
- averaging "Av", which always predicts the average of all possible event values, which is 0.5 in binary games.

The considered meta-inductive strategies are RW and, for sake of comparison, ITB.

Table 1 displays the frequencies of sequences for which the absolute success of a prediction method (based on the natural loss function) lies in a particular interval, as specified at the left margin. Success intervals are arranged symmetrically around the average value 10. As expected, the weak NFL theorem applies: in accordance with it one sees on the bottom line that the state-uniform average success is the same for all five methods. Nevertheless the frequencies of sequences for which these methods reach certain success levels are remarkably different.

The success frequencies of M-I and M-AI are different for different success intervals, because M-I has its highest success in very regular sequences (e.g., 1111...) with high frequencies of 1s or of 0s, in which ties of so far observed frequencies are rare, while M-AI has its highest success in oscillating sequences (e.g., 1010...) in which ties of so far observed frequencies are frequent. This brings a score of 0.5 more often to M-AI than M-I. As a result, M-I's success can climb higher than M-AI's success, though the frequency of such cases is small. In compensation, the number of sequences in which M-AI does only little better than average is higher than the corresponding number of sequences for M-I. Observe the mirror-symmetric distribution of sequences over M-I's and M-AI's success intervals, following from the fact that for any given sequence abs_{20}(M-I) $= 1 - \mathrm{abs}_{20}$(M-AI).

In contrast, Av always predicts 0.5 and earns a sum-of-scores of 10 in all possible worlds. The meta-inductive methods ITB and RW reach the top successes that object-induction (M-I) achieves in highly regular worlds, although in a diminished way due to their short run regrets. The advantage of RW is that it manages to avoid low success rates: following from its near access-optimality RW's success is in every possible sequence close to the maximal success in this sequence; so RW cannot fall much behind Av's success rate which is 0.5 (10 of 20 points) in all sequences. In contrast, M-I has a poor performance, and ITB and M-AI an even worse performance, in some sequences.

Similar tendencies can also be observed in other settings. Increasing the number of rounds has the effect that the frequency of sequences with high or low success rates steadily decreases, as explained in the previous section.

Based on these results we obtain a justification of object-induction and meta-induction even *within* the induction-hostile perspective of a state-uniform prior distribution for *binary short-run* sequences. What counts are two things: (a) To reach high success in those environments which allow for non-accidentally high success by their intrinsic regularities. This is what independent inductive methods do. (b) To protect oneself against high losses (compared to average success) in induction-hostile environments. This is what cautious methods such as Av do. The advantage of RW meta-induction is that it combines *both*: reaching high success rates where it is possible and avoiding high losses. Thus RW achieves 'the best of both worlds'. This, however, comes at the cost of a certain short-run regret.

The preceding version of a justification of meta-induction works within the most induction-hostile prior distribution – a SUPD. If one switches to a frequency-uniform prior distribution, one thereby adopts an induction-friendly perspective. This result is displayed in Table 2 for a simulation of all possible

Table 1. Meta-induction and no-free-lunch for binary-event games. Cells show percentage of possible binary sequences with 20 rounds, for which the five methods M-I, M-AI, Av, ITB and RW have reached certain intervals of absolute successes (left margin).

		M-I	M-AI	Av	ITB	RW
Absolute success intervals	[0,1)	0	0.0002	0	0	0
	[1,2)	0	0.003	0	0.0004	0
	[2,3)	0	0.029	0	0.008	0
	[3,4)	0	0.159	0	0.077	0
	[4,5)	0	0.618	0	0.394	0
	[5,6)	0.537	1.824	0	1.412	0
	[6,7)	3.540	4.254	0	3.708	0
	[7,8)	9.579	8.035	0	7.555	0
	[8,9)	15.622	12.476	0	12.966	36.491
	[9,10)	18.346	16.065	0	18.238	23.472
	10	8.910	8.910	100	9.848	0
	(10,11]	16.065	18.346	0	17.642	14.835
	(11,12]	12.476	15.622	0	14.213	11.880
	(12,13]	8.035	9.579	0	8.155	7.469
	(13,14]	4.254	3.540	0	3.558	3.595
	(14,15]	1.824	0.537	0	1.486	1.513
	(15,16]	0.618	0	0	0.555	0.560
	(16,17]	0.159	0	0	0.152	0.153
	(17,18]	0.029	0	0	0.029	0.029
	(18,19]	0.003	0	0	0.003	0.003
	(19,20]	0.0002	0	0	0.0002	0.0002
State-uniform average success		10	10	10	10	10

sequences of length 20 applied to the methods of Table 1. In Table 2 the left margin displays intervals for the possible frequencies of 1s in the 20-round sequences and the cells display the achieved (average) absolute successes of the methods for sequences whose frequencies lie in these intervals.

Note that the more decentered a frequency interval, the lower is its corresponding entropy. M-I is most successful in all frequency intervals that are not close to the center. It is only in the interval [0.4,0.6] – which, of course, contains many more individual sequences than the other intervals – that Av performs better than M-I. However, Av's mean success in this interval is worse than the mean success of the anti-inductive method M-AI, which performs badly in the decentral intervals. Again, the meta-inductive methods combine both features: in the central interval they don't lose much compared to Av, while in the decentral intervals their mean success rate comes close to that of M-I and beats

Table 2. Meta-induction for binary events from the perspective of frequency-uniform distributions. Cells show (average) absolute successes of the five methods M-I, M-AI, Av, ITB and RW, for binary sequences with 20 rounds, in dependence of their event-frequencies (left margin).

		M-I	M-AI	Av	ITB	RW
Frequency intervals	[0,0.1]	17,5	2,5	10	17,12	17,14
	[0.1,0.2]	15,4898	4,5102	10	14,7254	14,7103
	[0.2,0.3]	13,4314	6,56857	10	12,2263	12,1212
	[0.3,0.4]	11,2313	8,76873	10	9,81723	9,90422
	[0.4,0.5]	8,82824	11,1718	10	9,80868	9,77248
	[0.5,0.6]	8,82824	11,1718	10	9,80868	9,77248
	[0.6,0.7]	11,2313	8,76873	10	9,81723	9,90422
	[0.7,0.8]	13,4314	6,56857	10	12,2263	12,1212
	[0.8,0.9]	15,4898	4,5102	10	14,7254	14,7103
	[0.9,1]	17,5	2,5	10	17,12	17,14
Average for frequency uniform distribution		13.30	6.70	10	12.74	12.73

that of Av and M-AI. As expected, the frequency-uniform expectation values of the absolute success is much higher for inductive than for non-inductive methods; M-I has the lead, closely followed by ITB and RW.

6 Conclusion

In this paper we applied the (weak) no free lunch (NFL) theorem to regret-based meta-induction (RW) in the framework of prediction games. The challenge of the NFL theorem cannot be 'solved' by arguing that expected successes should be computed relative to the 'actual' (instead of some prior) distribution, because this idea is viciously circular. A more robust defense is possible based on an a priori result concerning the access-dominance of RW. Since this dominance result implies the existence of free lunches for RW it seems to contradict the NFL theorem. This conflict was dissolved based on four core result:

(1) A weak NFL theorem can be proved for prediction games (with binary events and natural loss function) under the assumption of a SUPD (a state-uniform probability distribution). A SUPD is maximally induction-hostile. In contrast, a frequency-uniform distribution is induction-friendly. Either sort of prior distribution is subjective and biased.

(2) Concerning success in the long run, the meta-inductive prediction strategy RW enjoys free lunches compared to all prediction methods that are not access-optimal (and most prediction methods aren't). However, the SUPD underlying the NFL theorem assigns a probability of zero to the class of all event sequences in which RW dominates other methods. This dissolves the apparent conflict with the NFL within the long-run perspective.

(3) Concerning success in the short run, the following short-run advantage of RW can be demonstrated even under the induction-hostile perspective of an SUPD: What counts is (a) to reach high success rates in regular (low-entropy) environments, which is what independent inductive methods do, and (b) to protect against high losses, compared to average success, in irregular (high-entropy) environments, which is what cautious "averaging" methods do. RW meta-induction combines both advantages, at the cost of a small short-run regret.
(4) If one assumes a frequency-uniform prior, then (meta-) inductive prediction strategies outperform non-inductive methods for all event sequences whose entropy is not close-to-maximal.

References

1. Schurz, G.: The meta-inductivist's winning strategy in the prediction game: a new approach to Hume's problem. Philos. Sci. **75**, 278–305 (2008)
2. Cesa-Bianchi, N., Lugosi, G.: Prediction, Learning, and Games. Cambridge University Press, Cambridge (2006)
3. Wolpert, D., Macready, W.: No free lunch theorems for search. Technical report SFI-TR-95-02-010, Santa Fe Institute (1995)
4. Wolpert, D.: On the connection between in-sample testing and generalization error. Complex Syst. **6**, 47–94 (1992)
5. Schaffer, C.: A conservation law for generalization performance. In: Machine Learning (Proceedings of ICML 1994), pp. 259–265. Morgan Kaufmann, Burlington (1994)
6. Rao, R., Gordon, D., Spears, W.: For every generalization action, is there really an equal and opposite reaction? In: Machine Learning (Proceedings of ICML 1995), pp. 471–479. Morgan Kaufmann, Burlington (1994)
7. Giraud-Carrier, C., Provost, F.: Toward a justification of meta-learning: is the no free lunch theorem a show-stopper? In: Proceedings of the ICML-2005 Workshop on Meta-learning, pp. 12–19 (2006)
8. Wolpert, D.: The lack of a priori distinctions between learning algorithms. Neural Comput. **8**, 1341–1390 (1996)
9. Schurz, G., Thorn, P.: The revenge of ecological rationality: strategy-selection by meta-induction. Minds Mach. **26**(1), 31–59 (2016)
10. Schurz, G.: Meta-induction is an optimal prediction strategy. In: Proceedings of the 18th Annual Belgian-Dutch Conference on Machine Learning, pp. 66–74. University of Tilburg (2009)
11. Gigerenzer, G., Todd, P.M., the ABC Research Group. (eds.): Simple Heuristics That Make Us Smart. Oxford University Press, Dordrecht (1999)
12. Schurz, G.: The Optimality of Meta-induction. A New Approach to Hume's Problem. Book manuscript (2017)
13. Carnap, R.: Logical Foundations of Probability, 2nd edn. University of Chicago Press, Chicago (1962)
14. Howson, C., Urbach, P.: Scientific Reasoning: The Bayesian Approach, 2nd edn. Open Court, Chicago (1996)
15. Solomonoff, R.: A formal theory of inductive inference. Inf. Control **7**(1–22), 224–254 (1964)