

Metainduction over Unboundedly Many Prediction Methods: A Reply to Arnold and Sterkenburg

Gerhard Schurz*†

The universal optimality theorem for metainduction works for epistemic agents faced with a choice among finitely many prediction methods. Eckhart Arnold and Tom Sterkenburg objected that it breaks down for infinite or unboundedly growing sets of methods. In this article the metainductive approach is defended against this challenge by extending the optimality theorem (i) to unboundedly growing sets of methods whose number grows less than exponentially in time, (ii) to sequences of methods with an application to Goodman's problem, and (iii) to infinite sets of methods whose number of predictive equivalence classes grows less than linearly in time.

1. Introduction. In Schurz (2008) a new account to the problem of induction has been developed that is based on the optimality of metainduction. The account agrees with Hume's skeptical insight that it is impossible to demonstrate in a noncircular way that induction is *reliable*, in the sense of being predictively more successful than random guessing. A modern version of Hume's insight is the (in)famous no-free-lunch theorem (Wolpert 1996) demonstrating that under a uniform prior distribution over possible worlds all nonclairvoyant methods of prediction have the same expected success (Schurz 2017).

Reichenbach (1949, sec. 91) argued that it is at least possible to demonstrate a priori that induction is *optimal*, that is, is the best that we can do

Received August 2019; revised January 2020.

*To contact the author, please write to: Department of Philosophy, Heinrich Heine University Duesseldorf, Universitaetsstrasse 1, Geb. 24.52, 40225 Duesseldorf, Germany; e-mail: schurz@phil.hhu.de.

†For valuable help, I am indebted to Paul Thorn, Christian Feldbacher-Escamilla, Tom Sterkenburg, Ronald Ortner, Igor Douven, Stathis Psillos, and Tomoji Shogenji.

Philosophy of Science, 88 (April 2021) pp. 320–340. 0031-8248/2021/8802-0007\$10.00
Copyright 2021 by the Philosophy of Science Association. All rights reserved.

for the purpose of predictive success. Results in formal learning show, however, that it is not possible to demonstrate universal optimality at the level of *object induction*, that is, of induction applied to the task of predicting events (Kelly 1996, 263). In contrast, what the account of *metainduction* attempts to show is that induction is optimal if it is applied at the metalevel of competing *prediction methods* (which comprise functions mapping observed input information into a prediction of the next event). A metainductive strategy tracks the success rate of all prediction methods whose predictions are accessible to the given epistemic agent and predicts an optimal weighted average of the predictions of those methods that were most successful so far. Using results in the theory of regret-based online learning (Cesa-Bianchi and Lugosi 2006), Schurz (2008; and more extensively in [2019], chaps. 6–7) proves that there exists a particular weighting method, called *attractivity weighting*, which grants the metainductivist a predictive long-run success rate at least as high as that of every other accessible prediction method. Schurz and Thorn (2016) call this kind of optimality *access optimality*. Within the short run, the metainductivist may suffer a loss compared to the best method because her choice of methods for future predictions is based on the methods' past success rates; thus, there is an unavoidable delay. However, attractivity-weighted metainduction is designed in such a way that this short-run loss has tight upper bounds that quickly approximate zero if the number of rounds of the prediction game exceeds the number of competing methods.

Note that by itself, this justification of metainduction does not entail anything about the rationality of object-level induction. Which prediction method is metainductively evaluated as optimal is relative to the given empirical track record of the accessible methods and, thus, an a posteriori matter. It may well be that we live in a world in which a method different from object induction is predictively superior. However, the a priori justification of metainduction gives us the following a posteriori justification of object induction: to the extent that so far object-inductive prediction methods were more successful than all accessible noninductive (object-level) prediction methods, we are metainductively justified to continue favoring object-inductive prediction methods in the future. This argument is no longer circular because a noncircular justification of metainduction has been independently established. In conclusion, the 'full package' of Schurz' account has two parts: (i) the analytic (or a priori) justification of metainduction and (ii) the empirical (a posteriori) justification of object induction based on i. This is nicely worked out by Sterkenburg (2019), although in some passages he seems to consider this fact as a deficit, while I see it as a virtue. Since it is impossible, by Hume's skeptical insight, to demonstrate the superiority of object induction (or of any other method) in an a priori manner, such a demonstration can only be given in an a posteriori manner, based on the a priori justification of metainduction.

Remarkably, the access optimality of meta-induction holds in all possible worlds, even in ‘chaotic’ worlds in which the event frequencies and success rates do not converge against limits but are permanently changing in irregular ways or in ‘paranormal’ worlds that host perfect clairvoyants or demonic deceivers. In this sense, the optimality result is universal. Its only restriction is the assumption that the class of accessible prediction methods is finite. Precisely this restriction is the point of the critique of Arnold (2010) and, more recently, Sterkenburg (2019). Schurz (2008) justifies this restriction by the fact that real agents have only finitely bounded cognitive resources; thus, they can have simultaneous access only to a finite set of prediction methods, the so-called pool of candidate methods. Neither for Arnold nor for Sterkenburg is this argument satisfactory.

Arnold emphasizes that since from a logical viewpoint there are infinitely many methods, meta-induction cannot establish absolute optimality but only optimality relative to an empirically given pool of methods. This is surely right and agrees with what is said in Schurz (2008): meta-induction provides a noncircular justification of the preference for one method against a finite set of competitors. Without a solution to this basic problem, we would not even be able to justify induction against a single competitor, such as counter-induction, random guessing, or religious prophecy.

Sterkenburg respects the argument from cognitive finiteness, but he objects that there are infinitely many methods that could be selected for meta-inductive evaluation, so the question remains to which finite pool of candidate methods meta-induction should be applied? Sterkenburg suggests that the answer to this problem should be the pool of methods that have been actually proposed for prediction. However, as Sterkenburg points out, new methods may be invented and proposed in the midst of a prediction competition; thus, a *dynamically expanding* candidate pool is needed, to which my account of meta-induction does not apply. In this article I propose a solution to this problem.

In the remainder of this article the meta-inductive approach is defended against the challenges of Arnold and Sterkenburg as follows. In section 2, the fundamental theorems about meta-induction in the framework of prediction games are briefly recapitulated. It is argued that the optimality result for finitely many methods solves the most important part of the problem of induction. In section 3 it is shown that the optimality theorem can be extended to an unboundedly growing set of prediction methods, provided their number grows less than exponentially in time. The optimality result even extends to all piecewise combinations of prediction methods, which provides an elegant solution to a variant of Goodman’s problem (sec. 4). Section 5 shows how for cognitively infinite beings the meta-inductive optimality result can even be extended to infinite sets of methods, provided the number of predictive equivalence classes grows less than linearly in time.

2. Metainduction over Finite Pools of Methods and Cognitive Finiteness. The account of metainduction is formally developed in the framework of sequential prediction games. A *prediction game* is a pair $((e), \Pi)$ consisting of:

1. An infinite sequence $(e) := (e_1, e_2, \dots)$ of events e_n coded by real numbers between 0 and 1, for example, a sequence of daily weather conditions or stock values. Each time n corresponds to one round of the game. The space of possible event *values* $v_1, v_2, \dots \in \text{Val}$ is denoted by $\text{Val} \subseteq [0, 1]$, and e_n denotes the *true* event value at time n (thus, ‘ $e_n = v_i$ ’ says ‘the true event value at time n was v_i ’).
2. A finite set of prediction methods or players $\Pi = \{P_1, \dots, P_m, \text{MI}\}$. 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. The metainductivist is signified by ‘MI’, and the other players are the ‘non-MI players’ or ‘candidate methods’. They may be real-life experts, computational algorithms, or even ‘clairvoyants’ who can see the future in ‘paranormal’ possible worlds. Moreover, they may include independent methods (based on individual learning) as well as other metamethods (basing their predictions on those of independent methods). MI has *simultaneous access* to all non-MI methods, which means that MI can monitor and score their predictions.

In *real-valued* games, the events need not be real-valued but may also be binary ($\text{Val} = \{0, 1\}$); a characteristic of these games is that players are allowed to predict real numbers, in the form of weighted averages of events in Val , for example, $0.4 \cdot v_1 + 0.6 \cdot v_2$. This is important for Bayesian prediction games in which predictions are probability distributions over Val (Schurz 2019, sec. 7.1; Sterkenburg 2019, sec. 2.1). Thus, formally, the space of possible predictions Val_{pred} may extend the event space: $\text{Val} \subseteq \text{Val}_{\text{pred}} \subseteq [0, 1]$. In contrast, in *discrete* prediction games this is forbidden; here, mixtures of events are not possible, and $\text{Val}_{\text{pred}} = \text{Val}$. Finally, note that all results about prediction games transfer to *action games*, by equating choices of actions with predictions about the action’s payoffs (Schurz 2019, sec. 7.5).

The *predictive success* of a method P is evaluated in terms of the following notions: (i) $\text{pred}_n(P)$ is the prediction of *player* P for time n delivered at time $n - 1$, (ii) the deviation of the prediction pred_n from the event e_n is measured by a normalized loss function, $\text{loss}(\text{pred}_n, e_n) \in [0, 1]$, (iii) $\text{score}(\text{pred}_n, e_n) =_{\text{def}} 1 - \text{loss}(\text{pred}_n, e_n)$ is the *score* obtained by prediction pred_n of event e_n , (iv) $\text{Suc}_n(P) =_{\text{def}} \sum_{1 \leq i \leq n} \text{score}(\text{pred}_i(P), e_i)$ is the *absolute success* achieved by player P until time n , (v) $\text{suc}_n(P) =_{\text{def}} \text{Suc}_n(P)/n$ is the *success rate* of player P at time n , and (vi) $\text{maxsuc}_n =_{\text{def}}$ the maximal success rate of all players at time n .

The *natural* (or absolute) loss function is defined as $|\text{pred}_n - e_n|$. The optimality theorem for real-valued games 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 including all linear, polynomial, or exponential functions of the natural loss function. The optimality theorem for discrete prediction games with randomized or collective metainduction holds even for all loss functions, though on some additional costs.

The simplest metainductive strategy is called ‘imitate the best’ and predicts what the currently best non-MI player predicts. It is easy to see that this metainductive method cannot be 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 metainductivist (Schurz 2008, sec. 4). Nevertheless, there exists a metainductive strategy that is provably universally optimal. This strategy is called *attractivity-weighted metainduction*, abbreviated as wMI, and is defined as follows:

Definition 1. The predictions of attractivity-weighted metainduction, abbreviated as wMI, are defined as

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

where the weight of method P at time n , $w_n(P)$, is a positive monotone function of the difference between P 's success and wMI's success at time n . This difference is called P 's *attractivity* for wMI at time n , defined as follows: $\text{At}_n(P) = \text{Suc}_n(P) - \text{Suc}_n(\text{wMI})$ is P 's absolute attractivity (for wMI) at time n , and $\text{at}_n(P) = \text{suc}_n(P) - \text{suc}_n(\text{wMI})$ is P 's attractivity rate at time n .

P 's attractivity for wMI is also called wMI's *regret* in comparison with P , and attractivity-weighted metainduction is a case of regret-based learning. Attractivity-based weights can be defined in different ways. Schurz (2008, sec. 7) used a simple definition that identified the methods' weights with the positive part of attractivity rates: $w_n(P) = \max(\text{at}_n(P), 0)$. This method grants long-run optimality but does not yield the best short-run bounds one can achieve. This article uses the exponential version of attractivity-weighted metainduction, abbreviated as eMI, whose weights are defined as exponentials of absolute attractivities:

Definition 2. $w_n(P) =_{\text{def}} e^{\eta \cdot \text{At}_n(P)}$, with $\eta = \sqrt{8 \cdot \ln(m)/(n+1)}$.

With these weights, the following optimality theorem can be proved:¹

Theorem 1 (short-run regret and universal access optimality of eMI). For every real-valued prediction game $((e)\{P_1, \dots, P_m, eMI\})$ with convex loss function,

- 1.1. Short run: $\max \text{suc}_n - \text{suc}_n(\text{eMI}) \leq \sqrt{2 \cdot \ln(m)/n} + \sqrt{\ln(m)/8 \cdot n^2} \leq \sqrt{2 \cdot \ln(m)/n} \cdot [1 + (1/4 \cdot \sqrt{n})] \leq 1.78 \cdot \sqrt{\ln(m)/n}$.
- 1.2. Long run: $\lim \sup_{n \rightarrow \infty} (\max \text{suc}_n - \text{suc}_n(\text{eMI})) \leq 0$.

Compared to the regret bound $\sqrt{m/n}$ obtained with the simple weight definition used in Schurz (2008, sec. 7), the bound of theorem 1 has $\log(m)$ in place of m , which is a significant improvement. Different from the simple weight definition, eMI does not disregard players with negative attractivities, but since exponential weights with negative attractivities decrease exponentially with increasing n , eMI ‘gradually forgets’ nonattractive non-MI players.

By using the logarithmic loss function that is preferred by Sterkenburg (2019, sec. 2.1), one obtains the even tighter bound of the metainductive regret of $\ln(m)/n$ (Cesa-Bianchi and Lugosi 2006, theorem 3.2 and proposition 3.1). However, the logarithmic loss function has rather special features: it tends to ignore small deviations between prediction and event but grows over all bounds if the absolute deviation converges to 1. For this reason we do not confine the central theorems to this loss function but state theorems holding for classes of loss functions as broad as possible.

Theorem 1 does not directly apply to discrete prediction games in which $\text{Val}_{\text{pred}} = \text{Val} = \{v_1, \dots, v_q\}$ must hold (thus, if the events are binary, the predictions must be binary, too). There are two methods by which the optimality result of theorem 1 can be transferred to discrete prediction games.

Method 1: Randomization, rMI (Cesa-Bianchi and Lugosi 2006, sec. 4.1).—Here rMI predicts each possible event value $v_r \in \text{Val}$ with a probability P equaling the weight sum of those non-MI players who predict this value:

Definition 3.

$$P(\text{pred}_{n+1}(\text{rMI}) = v_r) = \frac{\sum \{w_n(P_j) : \text{pred}_{n+1}(P_j) = v_r, 1 \leq j \leq m\}}{\sum \{w_n(P_j) : 1 \leq j \leq m\}}$$

1. See Schurz (2019, app. 12.24); the proof is based on theorems 2.2 and 2.3 of Cesa-Bianchi and Lugosi (2006) and holds for arbitrary prediction horizons. A simpler proof of the bound $\sqrt{2 \cdot \ln(m)/h}$ with a fixed prediction horizon h on which the weight definition depends is found in Shalev-Shwartz and Ben-David (2014, 253–54, theorem 21.11).

The optimality theorem for rMI is defined in terms of rMI's cumulative expected success rate, $\overline{\text{suc}}_n(\text{rMI})$, which is the sum of the probabilistically expected scores of rMI until time n , divided by n . Or formally,

Definition 4. $\overline{\text{suc}}_n(\text{rMI}) =_{\text{def}} (1/n) \cdot \sum_{1 \leq i \leq n} \text{Exp}(\text{score}_i(\text{rMI}))$, with $\text{Exp}(\text{score}_i(\text{rMI})) =_{\text{def}} \sum_{1 \leq r \leq q} \text{P}(\text{pred}_i(\text{rMI}) = v_r) \cdot \text{loss}(v_r, e_i)$.

Let b_n^{eMI} be the regret bound for eMI in stated in theorem 1. Then we obtain (see Schurz 2019, app. 12.25):

Theorem 2 (universal access optimality of rMI). For every discrete prediction game $((e), \{P_1, \dots, P_m, \text{rMI}\})$ with a randomized metainductivist rMI who satisfies the following probabilistic independence assumption

$$\begin{aligned} & \text{(PIA)} \\ & \text{P}(\text{pred}_{n+1}(\text{rMI}) | e_{n+1}, \text{Input}_n(\text{rMI})) = \text{P}(\text{pred}_{n+1}(\text{rMI}) | \text{Input}_n(\text{rMI})), \\ & \text{where } \text{Input}_n(\text{rMI}) = \{(\text{pred}_{n+1}(P_j), w_n(P_j)) : 1 \leq j \leq m\}, \end{aligned}$$

the following statements hold:

- 2.1. $\text{maxsuc}_n - \overline{\text{suc}}_n(\text{rMI}) \leq b_n^{\text{eMI}}$.
- 2.2. $\limsup_{n \rightarrow \infty} (\text{maxsuc}_n - \overline{\text{suc}}_n(\text{rMI})) \leq 0$.

Theorem 2 applies to all discrete prediction games with an arbitrary loss function over $\text{Val} \times \text{Val}$, as the values of Val need not have any numerical structure. The proof of theorem 2 rests on the observation that even if the loss function of a discrete prediction game is arbitrary, the expected loss of rMI's prediction probabilities, $\overline{\text{loss}}(\text{P}(\text{pred}_n), e_n)$, is linear and thus convex in the argument $\text{P}(\text{pred}_n)$. Therefore, the proof of theorem 1 can be transferred from eMI to rMI.

A substantial restriction of theorem 2 is the probabilistic independence assumption (PIA). It requires that rMI's random choice of prediction is independent of the predicted event conditional on the input information of rMI's algorithm. This excludes the possibility that the environment reacts 'demonically' to eMI's choice of prediction, in the sense of displaying the opposite event of rMI's predictions (Cesa-Bianchi and Lugosi 2006, 68). Since for a solution to Hume's problem this restriction is unacceptable, Schurz (2008, sec. 8) introduced an alternative solution for discrete prediction games that works without this restriction and is truly access universal, namely, collective metainduction.

Method 2: Collective metainduction, cMI.—Here a collective of meta-inductivists approximates the predictive probabilities of rMI by the frequencies of their discrete predictions as close as possible, with the result that

their average predictive success rate, denoted as $\widehat{\text{suc}}_n(\text{cMI})$, is guaranteed to be approximately access optimal. Assume there are k collective metainductivists, abbreviated as $\text{cMI}_1, \dots, \text{cMI}_k$. With a slight variation of the rounding method described in section 6.7.2 of Schurz (2019), the following optimality theorem can be proved:²

Theorem 3 (universal access optimality of cMI). For every discrete prediction game with k collective metainductivists $((e_n), \{P_1, \dots, P_m, \text{cMI}_1, \dots, \text{cMI}_k\})$ and q discrete event values:

- 3.1. $\text{maxsuc}_n - \text{suc}_n(\text{cMI}) \leq b_n^{\text{cMI}} + (q - 1)/(2 \cdot k)$.
- 3.2. $\limsup_{n \rightarrow \infty} (\text{maxsuc}_n - \text{suc}_n(\text{cMI})) \leq (q - 1)/(2 \cdot k)$.

One might object that theorem 3 grants access optimality only for the average success of the metainductive collective but not to every individual metainductivist. However, if the cMI's form a cooperative collective and share their success equally, then every individual cMI is guaranteed to earn cMI's mean success and, thus, to be access optimal. In this way the game-theoretic requirement of cooperation achieves fundamental importance for a problem of epistemological justification.

Let us finally point out that theorems 1–3 assert the (access) optimality of eMI (or rMI or cMI) but not its (access) dominance in the sense that eMI beats every alternative method M^* in at least one prediction game. There exist different variants of attractivity-based metainduction that are likewise optimal. Thus, no particular version of metainduction can be universally dominant. However, a variety of restricted dominance theorems can be proved. For example, eMI dominates all independent methods and all metamethods that are not access optimal.³ These dominance results provide a partial solution to Wolpert's no-free-lunch theorem (Schurz 2017; Schurz and Thorn 2017).

This article, however, is focused on the optimality of eMI and its defense against Arnold's and Sterkenburg's challenge of infinitely many methods. To model such a situation, we abbreviate the set of non-MI methods as \mathcal{P} ; thus, $\Pi = \mathcal{P} \cup \{\text{MI}\}$, where \mathcal{P} may be infinite and MI is some kind of metainductivist. It is easy to see that no kind of metastrategy, however it is defined, can be universally optimal if it is confronted with an at least countably infinite pool of candidate methods. Assuming a finite value space

2. In the proof on p. 154 of Schurz (2019), we simply have to replace “1” by “0.5” in lines 8, 11, 12, and 17. In the continuation of this proof on p. 334, we replace “one” (resp. “1”) by “0.5” in lines 5 and 6, “ ≤ 1 ” by “ ≤ 0.5 ” in line 7, and “ $1/k$ ” by “ $1/2 \cdot k$ ” in lines 15 and 17.

3. Many well-known metamethods fail to be access optimal, e.g., take-the-best, linear regression or success-based weighting (Schurz 2019, sec. 8.3.1–2).

$\text{Val} = \{v_1, \dots, v_q\}$ with $v_1 = 0 < v_2 < \dots < v_q = 1$ and natural loss, we can demonstrate this by a slight extension of the counterexample of Arnold (2010, 588–89). In the first round of the prediction game, \mathcal{P} is split into q disjoint subsets $\mathcal{P}_{1,i}$ ($1 \leq i \leq q$) of infinite cardinality, where all players in $\mathcal{P}_{1,i}$ predict event value v_i . In each of the following rounds $n > 1$, the infinite subset of players that predicted correctly is again split into q infinite disjoint subsets $\mathcal{P}_{n,i}$ ($1 \leq i \leq q$), where all players in $\mathcal{P}_{n,i}$ predict event value v_i . It follows that at every time n there remains an infinite subset of players that have a perfect success rate of 1, independently of the chosen event sequence. Yet, by choosing an MI-‘demonic’ event sequence satisfying $e_n = 1$ if $\text{pred}_n(\text{MI}) < 0.5$ (else $e_n = 0$), the metainductivist’s success rate is driven down to 0.5 or lower. A slight modification of this proof applies to an infinite value space.⁴

In conclusion, a metastrategy that is universally access optimal among infinitely many prediction methods is impossible. Let us now investigate the philosophical relevance of this problem. Its relevance may be doubted in the light of the following fact:

Fact of cognitive finiteness (FCF). All real epistemic agents (whether human or nonhuman) have finitely bounded computational powers. Thus, they can have simultaneous access only to finite candidate sets whose success evaluation remains within a given complexity bound.

Mathematicians often argue that humans can represent infinite sets. But if they do this, they represent infinities by finite strings of symbols that do not exceed a certain complexity bound. Let us denote this complexity bound as ‘ b ’. FCF entails that real agents can have simultaneous access only to a finite class \mathcal{P} of candidate methods, whose success evaluation for a given data stream (i.e., computing $\text{Suc}_n(\mathcal{P})$ for all $P \in \mathcal{P}$) does not exceed their complexity bound b .⁵ Of course, the size of b is an empirical fact, depending on the stage of technical development.

4. If Val is countably infinite, the same proof applies, since ω can be partitioned into infinitely many subsets of infinite cardinality (e.g., into $\{\{m^n : n > 1\} : m \in \omega\}$). If Val is an interval, $[a, b] \subseteq [0, 1]$, then we partition it (for a small $\varepsilon > 0$) into $q =_{\text{def}} (b - a)/\varepsilon$ subintervals of ‘finest accuracy’ ε . Then (assuming a proper loss function) there will be in every round one subset of players with a success rate of $\geq 1 - \varepsilon$ while MI’s success rate is ≤ 0.5 .

5. The notion of complexity of a method M , $K(M)$, is understood in the sense of the algorithmic or Kolmogorov complexity (K): $K(M)$ is the length of the shortest program in the language of a universal computer that produces a description of M (Cover and Thomas 1991, chap. 7). The complexity of a computation performed by M can be defined as $K(M)$ plus the maximal length of temporarily stored descriptions during the computation.

It seems to follow that for a real agent the a priori justification of meta-induction remains intact: the best that she can do is to apply attractivity weighting to the finite pool of candidate methods that are simultaneously accessible to her. However, there arises a successor problem: that of selecting one's candidate set. The problem stems from the fact that there are many more methods that a b -bounded cognitive agent can access individually than she can access and evaluate simultaneously. Thus, a selection is needed, as in the Olympic games in which not all but only a limited number of top athletes may participate. I call this the problem of selecting one's candidate set. The problem achieves significant importance at the intersubjective level, when meta-inductive methods are used as a rational means of resolving disagreements between adherents of radically different 'epistemic cultures'. Meta-induction will fail to achieve this goal if members of different cultures choose different candidate sets.

One may respond that in order to arrive at an intersubjective consensus one should simply apply meta-induction to the union of all candidate sets of the disagreeing persons or parties. In most practical cases, this strategy will be feasible. Yet the situation may arise that more methods are proposed than can be simultaneously compared within a prediction game. In this situation a selection is needed. The theoretical need of selection mechanisms is obvious from the fact that one may define not only reasonable but also arbitrarily bizarre methods, including the infamous *Goodman-type* methods that will be treated in section 4. Moreover, as Sterkenburg (2019) points out, the development of methods is a historical process: new methods may be invented and proposed 'on the fly', in the midst of a prediction competition. Thus, we need an extension of the meta-inductive approach to prediction games with dynamically expanding candidate sets.

If we had an extension of the meta-inductive account to an expanding candidate set of proposed methods, we could use it for a solution to the problem of an unmanageably large candidate set by the following strategy: we ignore the class of methods of lowest attractivity whose weight sum is ε . This will change the meta-inductive predictions and success rate only marginally, by a small additional loss term whose worst case size depends on ε and the loss function. Since typically most methods in a hyperlarge candidate set \mathcal{P} are unsuccessful, this will reduce the size of the hitherto candidate set significantly and give room for adding new so far unexplored methods to the candidate set. Since the meta-inductive account avoids inductive a priori assumptions, the historical backlog of unsuccessful methods should not be ignored forever; periodically these methods should be given a new chance by adding them again to the candidate pool. In this way, an extension of meta-induction to an expanding candidate set would allow an elegant handling of the problem of selecting a candidate set. In the next section I propose such an extension.

3. Metainduction over Unboundedly Growing Sets of Players. Prediction games with unboundedly growing numbers of players have the form $((e), \mathcal{P}(n) \cup \{\text{MI}\})$, where $\mathcal{P}(n) =_{\text{def}} \{P_1, \dots, P_{m(n)}\}$ is a set of non-MI players increasing with time and $m(n)$ is the number of non-MI players at time n . The growth function $m(n)$ never decreases and increases unboundedly with $n \rightarrow \infty$, although it need not increase at every time step. We assume that each non-MI player P predicts persistently after the first time at which P was added to $\mathcal{P}(n)$; we call this time P 's *entrance time* and denote it as $t(P)$.

We cannot directly apply the results of theorems 1–3 because these theorems presuppose that the success of all non-MI players—including those added at later times of the game—is evaluated right from the start. Therefore, all plausible methods of embedding new players into a prediction game's success evaluation have to assume some way of attributing a *default success* to the new players before they were participating in the game, that is, default values $\text{Suc}_{t(P)}(P)$ and $\text{suc}_{t(P)}(P) = \text{Suc}_{t(P)}(P)/t(P)$ for the given event sequence $e_1, \dots, e_{t(P)}$. Without an attribution of past default successes, it is impossible to establish access optimality, because the future may give the new players an unproportional advantage. As an example, suppose the present time is 1,100 and at time $n = 1,000$ a couple of new players have been added, and assume that for all players predictive success was much harder to attain in the first 1,000 rounds than in the last 100 rounds. In that case, the 'lucky' new players have the advantage of only being evaluated in the easy part of the game, and eMI's success as evaluated from $t = 0$ will be much lower than the success of lucky new players. As time goes on, the 'lucky advantage' of the new players will diminish, but with unboundedly growing player sets this situation may recur at any time when a new player enters. Thus, metainduction over expanding sets of players cannot be universally access optimal without some method of attributing default past success.

Let $\text{Suc}_{n|t(P)} =_{\text{def}} \sum_{t(P) < i \leq n} \text{score}(P, e_i)$ be the actual absolute success player P acquired since P was participating in the game, and likewise for $\text{suc}_{n|t(P)}(P) = \text{Suc}_{n|t(P)}(P)/(n - t(P))$. Then P 's success evaluation uses the success measure $\text{Suc}_n(P) = \text{Suc}_{t(P)}(P) + \text{Suc}_{n|t(P)}(P)$ and $\text{suc}_n(P) = \text{Suc}_n(P)/n$. The question is, which default success should we attribute to a new player? One possibility would be to attribute to a new player P the average success of a random guess. But since one's candidate methods are typically much better than random guesses, this would amount to a strong negative prejudice against newly added players: by default their hypothetical success, had they participated from the start, is conjectured to be low. We need a better and unbiased method of attributing a default success to new players. A better method, first implemented by Chernov and Vovk (2009), consists in attributing to a new player by default the same track record as the metainductivist. We call this the method of *self-completion*.

Let eMI_g denote the version of strategy eMI applied to growing sets of methods (and likewise for rMI_g and cMI_g). The method of self-completion

is philosophically satisfying because it guarantees that the new player P has a nonnegligible weight at its entrance time $t(P)$ (since eMI_g 's success is typically close to the success of the best player). If P predicts better than the other players, eMI_g will quickly recognize this by increasing P 's weight. The price coming with this advantage is that if a new player P predicts badly, eMI_g needs some time to degrade P 's weight. But we think this price is worth its cost. The technical advantage of the method of self-completion is that by setting the default prediction and score of an absent player P equal to that of the metainductivist, it follows that the hypothetical inclusion of P 's default predictions in eMI_g 's weighted average prediction would not change eMI_g 's prediction, which is de facto only based on the actually present players. Summarizing, the weights of eMI_g are defined as follows:

Definition 5 (default evaluation and weighting method of eMI_g [rMI_g, cMI_g]).

5.1. For every player P :

- 5.1.1. For eMI_g , $\text{pred}_n(P) = \text{pred}_n(eMI_g)$ for all times $n \leq t(P)$; thus, $\text{score}_n(P) = \text{score}_n(eMI_g)$ for all $n \leq t(P)$, and $\text{Suc}_{t(P)}(P) = \text{Suc}_{t(P)}(eMI_g)$.
- 5.1.2. For rMI_g and cMI_g , $P(\text{pred}_n(P) = x) = P(\text{pred}_n(rMI_g) = x)$ for $n \leq t(P)$; thus, $\text{score}_n(P) = \overline{\text{score}}_n(P) = \overline{\text{score}}_n(rMI_g)$, and $\text{Suc}_{t(P)}(P) = \overline{\text{Suc}}_{t(P)}(P) = \overline{\text{Suc}}_{t(P)}(rMI_g)$ (i.e., absent players are default evaluated like randomizing players).

5.2. The metainductive algorithm is applied only to present players. The other notions are defined as before.

With these assumptions we can prove the following result:

Theorem 4 (access optimality of eMI_g, rMI_g , and cMI_g for growing player sets).⁶ Let xMI_g be one of eMI_g, rMI_g , and cMI_g (thus x varies over $\{e, r, c\}$) and $b_n^{xMI_g}$ be the regret bound of xMI established in theorems 1–3, but with ‘ m ’ replaced by the growth function ‘ $m(n)$ ’. Then:

- 4.1. For every prediction game $((e), \{P_1, \dots, P_{m(n)}, xMI_g\})$, (a) $\text{maxsuc}_n - \text{suc}_n(eMI_g) \leq b_n^{eMI_g}$, (b) $\text{maxsuc}_n - \overline{\text{suc}}_n(rMI_g) \leq b_n^{rMI_g}$, and (c) $\text{maxsuc}_n - \overline{\text{suc}}_n(cMI_g) \leq b_n^{cMI_g}$.
- 4.2. If $m(n)$ grows slower than exponentially with n , that is, $\lim_{n \rightarrow \infty} m(n)/e^n = 0$, then (a) $\limsup_{n \rightarrow \infty} (\text{maxsuc}_n - \text{suc}_n(eMI_g)) \leq 0$,

6. Mourtada and Maillard (2017, 7) develop a variant of this algorithm based on a possibly nonuniform prior distribution and using an exp-concave loss function.

$$(b) \limsup_{n \rightarrow \infty} (\maxsuc_n - \overline{\text{suc}}_n(\text{rMI}_g)) \leq 0, \text{ and } (c) \limsup_{n \rightarrow \infty} (\maxsuc_n - \widehat{\text{suc}}_n(\text{cMI}_g)) \leq (q - 1)/(2 \cdot k).$$

Proof. We first consider $e\text{MI}_g$. We choose a variable time horizon n and denote by $e\text{MI}_g^*$ the virtual metainductive method that applies the method $e\text{MI}_g$ to all players in $\mathcal{P}(n)$ from the start, attributing predictions to absent players by the self-completion method. The players present at time $t \leq n$ are $P_1, \dots, P_{m(t)}$, and $P_{t+1}, \dots, P_{m(n)}$ are the players absent at time t ; $w_t(P_i)$ is the weight of player i at time t (which for absent players, as we will see, need not be known). We abbreviate as follows: $W_{\text{pres},t} = \sum_{1 \leq i \leq m(t)} w_t(P_i)$ is the weight sum of players present at t , attributed by $e\text{MI}_g^*$; $W_{\text{abs},t} = \sum_{m(t) < i \leq m(n)} w_t(P_i)$ is the weight sum of players absent at t ; and $w_{\text{pres},t} = W_{\text{pres},t} / (W_{\text{pres},t} + W_{\text{abs},t})$ and $w_{\text{abs},t} = W_{\text{abs},t} / (W_{\text{pres},t} + W_{\text{abs},t})$ are the relative weights of the classes of present and absent players at time t , respectively.

We now prove that $\text{suc}_t(e\text{MI}_g^*) = \text{suc}_t(e\text{MI}_g)$, by showing that $\text{pred}_t(e\text{MI}_g^*) = \text{pred}_t(e\text{MI}_g)$ holds for all $t \leq n$. We have

$$\begin{aligned} & \text{pred}_{t+1}(e\text{MI}_g^*) \\ &= \frac{\sum_{1 \leq i \leq m(t)} w_t(P_i) \cdot \text{pred}_{t+1}(P_i) + \sum_{m(t) < i \leq m(n)} w_t(P_i) \cdot \text{pred}_{t+1}(P_i)}{W_{\text{pres},t} + W_{\text{abs},t}} \\ &= \frac{\sum_{1 \leq i \leq m(t)} w_t(P_i) \cdot \text{pred}_{t+1}(P_i)}{W_{\text{pres},t}} \cdot w_{\text{pres},t} + \text{pred}_{t+1}(e\text{MI}_g^*) \cdot w_{\text{abs},t} \\ &= \text{pred}_{t+1}(e\text{MI}_g) \cdot w_{\text{pres},t} + \text{pred}_{t+1}(e\text{MI}_g^*) \cdot w_{\text{abs},t}. \end{aligned} \tag{1}$$

Since $w_{\text{pres},t} + w_{\text{abs},t} = 1$, equation (1) has the mathematical form $y = x \cdot w + y \cdot (1 - w)$, whose only nontrivial solution is $y = x$; that is, $\text{pred}_t(e\text{MI}_g) = \text{pred}_t(e\text{MI}_g^*)$. Thus, $e\text{MI}_g$ predicts identically to $e\text{MI}_g^*$. So we can apply theorem 1 to $e\text{MI}_g^*$, which proves theorem 4 for $e\text{MI}_g$. To prove theorem 4 for rMI_g (and likewise for cMI_g), we identify the absent players' default weights with the probabilities of rMI_g 's predictions. Thus, instead of equation (1) we have, abbreviating ' $\text{pred}_{t+1}(\text{rMI}_g^*) = x$ ' as ' $\text{pred}_{t+1}^* = x$ ',

$$\begin{aligned} & P(\text{pred}_{t+1}^* = x) \\ &= \frac{\sum \{w_t(P_i) : 1 \leq i \leq m(t), \text{pred}_{t+1}(P_i) = x\} + W_{\text{abs},t} \cdot P(\text{pred}_{t+1}^* = x)}{W_{\text{pres},t} + W_{\text{abs},t}} \\ & \quad (\text{where } W_{\text{abs},t} = m(n) - m(t)) \\ &= P(\text{pred}_{t+1}(\text{rMI}_g) = x) \cdot w_{\text{pres},t} + P(\text{pred}_{t+1}^* = x) \cdot w_{\text{abs},t}, \end{aligned}$$

which implies $P(\text{pred}_{t+1}(\text{rMI}_g) = x) = P(\text{pred}_{t+1}^* = x)$, as above. QED

4. A Solution to Goodman's Problem for Prediction Methods. As mentioned in section 2, the 'hard side' of the selection problem consists in the fact that one may propose arbitrary bizarre 'methods', including the infamous Goodman-type methods. In Goodman's (1946) original example, the projection of the complex predicate 'grue' over emeralds' amounts to the inductive projection of their observed color green until a future time point k and the counterinductive prediction of their color being blue after time k .⁷ Applying this idea to methods, a Goodman method G_k with switch point k predicts by an object-inductive method M until time k , and afterward it predicts the opposite of M (i.e., $1 - \text{pred}(M)$). Generalizing this idea, the class of Goodman combinations over a given class of base methods $\{M_1, \dots, M_m\}$ consists in all arbitrary concatenations of these base methods along the discrete time: 'between time 1 and n_1 apply M_{i_1} , between time $n_1 + 1$ and n_2 apply M_{i_2} ', and so on. Since the number of Goodman methods increases exponentially with time and the number of admissible switches, it is impossible to include all of them in the candidate set, because by theorem 4 \mathcal{P} has to grow slower than exponentially in time. Nevertheless, there is an elegant solution to this problem, based on results in online learning about *player sequences*. I explain this solution for the metainductive strategy eMI in real-valued prediction games.

Let the variable $s_{n,k}$ range over player sequences of length n over a given set of *base* players (or methods, experts) $\{P_1, \dots, P_m\}$ with at most k ($\leq n$) switches between players. Thus, $s^n = (P_{i_0}, \dots, P_{i_{n-1}}) \in \{P_1, \dots, P_m\}^n$ such that $P_{i_t} \neq P_{i_{t+1}}$ holds for at most k distinct t 's. These player sequences are now treated as 'complex players'. The prediction and score of a player sequence $s_{n,k}$ for a time $t + 1 \leq n$ is identified with the prediction and score of its player at time t , P_{i_t} . The metainductivist's regret is defined with respect to the success rates of all player sequences of length n in a given set \mathcal{S}_n^k of such sequences. Since the number of all possible player sequences grows with n , metainduction applied to player sequences is a variant of metainduction with unboundedly growing player sets, with the advantage that all player sequences reach backward toward time 0 and are success evaluable from the start. The application of metainduction to a set of possible sequences \mathcal{S}_n^k of length n is straightforward: the regret bounds are obtained by inserting the number of these sequences $|\mathcal{S}_n^k|$ for the time-dependent number $m(n)$ in the regret bounds of theorem 4. The only problem is that the set of all possible player sequences grows exponentially with n and k and is, thus, unfeasible.

7. Goodman showed that by a linguistic redefinition of the primitive symbols of one's language one may translate a *positional* predicate that refers to a location in time into a simple *qualitative* predicate and vice versa. The search for a language-independent distinction between qualitative and positional predicates is a separate problem for which solutions are possible (Schurz 2019, sec. 4.2). Here we focus on the selection problem for candidate methods described in a language with a given distinction between qualitative and positional predicates.

Fortunately, there exists a version of the metainductive algorithm that only tracks the base players and is yet access optimal with respect to all player sequences in S_n^k : the *fixed share* variant of eMI, abbreviated as feMI. The feMI algorithm is defined as follows, where $H(x) = -x \cdot \ln(x) - (1 - x) \cdot \ln(1 - x)$ denotes x 's binary entropy:

Definition 6 (weighting method of feMI [fixed share eMI]). The unnormalized weight $w_n(P_i)$ of player P_i at time n is given as $w_n(P_i) = \alpha \cdot (W'_n/m) + (1 - \alpha) \cdot w'_n(P_i)$, with $w'_n(P_i) = e^{\eta \cdot A_n(P_i)}$, $W'_n = \sum_{1 \leq j \leq m} w'_n(P_j)$, $\alpha \in [0, 1]$, the constants set to

$$\eta = \sqrt{8 \cdot (k + 1) \cdot \frac{\ln(m)}{n} + (n - 1) \cdot H\left(\frac{k}{n - 1}\right)},$$

and for $k > 0$, $\alpha = k/(n - 1)$; $k = 0$, $\alpha = 0$.

The weights $w'_n(P_i)$ are eMI's standard weights; feMI adds to them a uniformly distributed α -share that guarantees the weight of all base players stays above some minimal level, and feMI can access the best player sequence. Remarkably, feMI has provably the same upper regret bound as eMI applied to the exponentially many player sequences in S_n^k (Cesa-Bianchi and Lugosi 2006, 103–4, theorem 5.1). Thus, feMI achieves a kind of *access superoptimality*: it tracks only base players but is access optimal for all sequences of them with a bounded number of switches.

When applying the strategy feMI to the problem of Goodman methods, we are interested in feMI's regret in regard to combinations of methods that *reach into the future*, that is, much later than the present time n . Let $f > n$ be an arbitrary future time. It is easy to see that S_n^k (the set of length- n sequences with at most k switches) equals $S_f^k \uparrow n$, that is, the set of initial subsequences of sequences in S_f^k of length n . Since the success rate of a sequence $s_{f,k}$ (reaching to future time f) at time n depends only on $s_{f,k}$'s initial subsequence until time n , feMI's regret bound also holds for all player sequences of length greater than n , provided the maximal number of their switches is still equal to k . In a final extension we allow also m (the number of base players) and k (the number of allowed switches) to grow with n ; that is, $m = m(n)$ and $k = k(n)$. The resulting strategy feMI_g (feMI over growing base sets) is characterized by the following theorem:

Theorem 5 (access superoptimality of feMI_g over sequences of methods).⁸ For every prediction game $((e), \{P_1, \dots, P_{m(n)}, \text{feMI}_g\})$, where $S_{\geq n}^{k(n)}$ denotes the set of all admissible player sequences of length $\geq n$ with at most $k(n)$ switches,

8. A similar result is obtained by Mourta and Maillard (2017, 5).

5.1. $\max \text{suc}_n(\mathcal{S}_{\geq n}^{k(n)}) - \text{suc}_n(\text{feMI}_g) \leq$

$$\sqrt{2 \cdot (k(n) + 1) \cdot \frac{\ln(m(n))}{n} + H\left(\frac{k(n)}{n-1}\right) \cdot \left(1 + \frac{1}{4 \cdot \sqrt{n}}\right)}.$$

5.2. feMI_g is access optimal over $\mathcal{S}_{\geq n}^{k(n)}$ if $k(n) \cdot \ln m(n)$ grows slower than linearly with n .

5.3. Results 5.1 and 5.2 apply to frMI_g (with ‘ $\widehat{\text{suc}}_n$ ’ replacing ‘ suc_n ’) and to crMI_g (with ‘ $\widehat{\text{suc}}_n$ ’ replacing ‘ suc_n ’ and adding the regret term $(q - 1)/(2 \cdot k)$).

Proof. Result 5.1 follows from theorem 5.1 and corollary 5.1 of Cesa-Bianchi and Lugosi (2006, 105). The bound for feMI_g is obtained from the bound for the number M of sequences in \mathcal{S}_n^k , $M \leq m^{k+1} \cdot e^{(n-1)m \cdot H(k/(n-1))}$ (2006, 101), by replacing ‘ m ’ in theorem 1 by this bound and outfactoring $(1 + [1/(4 \cdot \sqrt{n})])$. Note that Cesa-Bianchi and Lugosi’s m is our $k(n)$, and their N our $m(n)$. While Cesa-Bianchi and Lugosi insert the upper M bound into their result for a variant of eMI with a fixed prediction horizon, we insert it into eMI’s bound stated in theorem 1, which gives result 1 of theorem 5; moreover, we state our bound in terms of regret rates by dividing through n . Result 5.2 follows, since if $\lim_{n \rightarrow \infty} k(n) \cdot \ln m(n)/n = 0$, the expression under the square root in 5.2 converges to zero for $n \rightarrow \infty$. Result 5.3 is an obvious consequence. QED

Theorem 5 offers a beautiful solution to Goodman’s problem for prediction methods. It entails that feMI_g is access optimal in regard to all Goodman-type combinations over a fixed set of m basic prediction methods whose switch number $k(n)$ grows sublinearly with n . If m is allowed to grow, the sublinear growth condition applies to the product $k(n) \cdot \ln(m)$.

5. Metainduction over Infinite Sets of Players. In this section we show how our results can be extended to prediction games with infinite sets of candidate methods, under the *counterfactual assumption* that the metainductivist has infinite cognitive powers allowing him to simultaneously access an infinite pool of methods \mathcal{P} . There are two ways of generalizing the notion of a weighted sum $\sum_{1 \leq i \leq m} w_n(P_i) \cdot \text{suc}_n(P_i)$ to an infinite class of players \mathcal{P} . One way is explained in section 4.2 of Sterkenburg (2019); this method is restricted to a countably infinite $\mathcal{P} = \{P_i : i \in \omega\}$ and assigns a countably additive prior weight $w_0(P_i)$ to all P_i ($i \in \omega$). Thus, $\sum_{i \in \omega} w_0(P_i) = 1$, and $w_n(P_i) = w_0(P_i) \cdot e^{\eta \cdot \text{Atn}(P)}$ (recall definition 2). The problem with this method, as explained in section 4.2 of Sterkenburg (2019), is that it fails to achieve access optimality, because on reasons of mathematical necessity $w_0(P_i)$ must

decrease rapidly with increasing n , implying that at any time n of the prediction game there will be k such that the weight sum of all players with index $\geq k$ will still be negligibly small, so that eMI's success rate cannot approximate the maximal success rate of all players (although it can approximate the success rate of every individual player, but not simultaneously).

In this section we propose a different method that generalizes eMI to arbitrary infinite classes \mathcal{P} that is based on splitting the infinite set \mathcal{P} at any given time point n into a finite partition $\mathcal{P}_1, \dots, \mathcal{P}_{m(n)}$ of *predictive equivalence classes*; thus, each player in \mathcal{P}_i ($1 \leq i \leq m(n)$) has delivered the same predictions and has the same success up to time n . The assumption of finitely many predictive equivalence classes is strictly justified only for finite predictive value spaces: if $|\text{Val}_{\text{pred}}| = c$, then in round n there are at most c^n predictive equivalence classes. For real-valued predictions Val_{pred} is not finite. However, we assume that every cognitive representation of a real number r has a maximal accuracy interval, α , by which r is approximated. Given α , there are only $c =_{\text{def}} 1/\alpha$ possible distinct events and predictions and, thus, at most c^n predictive equivalence classes in round n .

At any time n we identify every equivalence class \mathcal{P}_i with one virtual player and let the sum in the definition of $\text{pred}_{n+1}(\text{eMI})$ run over the elements of $\{\mathcal{P}_1, \dots, \mathcal{P}_{m(n)}\}$. We denote the version of eMI that applies the definition of attractivities and weights to the finitely many predictive equivalence classes of an infinite \mathcal{P} as eMI_{inf} . Since $m(n)$ is a nondecreasing function of n and in each new round new 'splittings' of equivalence classes may occur (i.e., hitherto equivalently predicting players may predict differently), eMI_{inf} looks like a version of eMI for growing player sets. However, we cannot apply the method of attributing a default success, because all new equivalence classes that arise from a splitting of the same old equivalence classes share their actual success history. We can only reduce an infinite eMI_{inf} -prediction game to a success-equivalent game with an ordinary eMI metainductivist, for a given time point n , if we assume that there are as many individual players already at the start as predictive equivalence classes of players have been formed until round n in the infinite eMI_{inf} game. In what follows, let $m(n)$ be a function of n that returns the number of equivalence classes (virtual players) $\mathcal{P}_1(n), \dots, \mathcal{P}_{m(n)}(n)$ that emerged by round n . Let $P_1, \dots, P_{m(n)}$ be the corresponding ordinary players in the corresponding eMI game. Each ordinary player P_i corresponds to one branch of the tree of branching equivalence classes until time n . Thus, in each earlier round $t < n$ with fewer equivalence classes $\mathcal{P}_1(t), \dots, \mathcal{P}_{m(t)}(t)$, $m(t) < m(n)$, some of the equivalence classes correspond to more than one player of the eMI game with $m(n)$ players. This implies that eMI will assign different weights to the predictions of the players corresponding to one equivalence class than the weights assigned by eMI_{inf} to these predictions, except when in every round each equivalence class corresponds to the same number of players in the eMI game. If this is

the case, we speak of a *symmetric* splitting process: in such a process, at any time point t at which a split occurs, every cell of the partition $\mathcal{P}_1(t - 1), \dots, \mathcal{P}_{m(t-1)}(t - 1)$ produced until time $t - 1$ splits into the same number of new cells. In a symmetric split process, the predictions of eMI_{inf} and that of eMI are identical in every round, since if $p(t)$ is the number of players that each equivalence class contains at time $t \leq n$, the following equation holds:

$$\begin{aligned} \text{pred}_{t+1}(eMI_{inf}) &= \frac{\sum_{1 \leq i \leq m(t)} w_t(\mathcal{P}_i) \cdot \text{pred}_{t+1}(\mathcal{P}_i)}{\sum_{1 \leq i \leq m(t)} w_t(\mathcal{P}_i)} \\ &= \frac{\sum_{1 \leq i \leq m(t)} p(t) \cdot w_t(\mathcal{P}_i) \cdot \text{pred}_{t+1}(\mathcal{P}_i)}{\sum_{1 \leq i \leq m(t)} p(t) \cdot w_t(\mathcal{P}_i)} = \text{pred}_{t+1}(eMI). \end{aligned} \tag{2}$$

Thus, for every time point n , an infinite eMI_{inf} game with a symmetric split process can be reduced to an ordinary eMI game with $m(n)$ players.

If the split process is asymmetric (as in Arnold’s counterexample sketched in sec. 2), this simple method of transferring our theorems about eMI to infinitely many players is impossible. However, we can obtain an analogous result by constructing the *symmetric completion* of an arbitrary splitting process as follows: we add to the game a sufficient number of pseudoequivalence classes that agree in their predictions and produce a symmetric split process. Thus, now the equivalence classes (virtual players) consist of proper equivalence classes (each having a different prediction history) and pseudoequivalence classes (agreeing in their prediction history with one of the proper equivalence classes).

The equivalence classes are recursively defined as follows. Let $m(n)$ be the number of proper equivalence classes and $s(n) > m(n)$ the number of equivalence classes (including ‘pseudos’) constructed in the symmetric completion of this game until round n . (Initially, $m(0) = s(0) = 1$; i.e., before making a prediction all players are predictively equivalent.) Assume in round $n + 1$ that only one of the $s(n)$ equivalence classes $\mathcal{P}_1(n), \dots, \mathcal{P}_{s(n)}(n)$ of round n , say $\mathcal{P}_k(n)$, splits into Δ new proper equivalence classes, while the players in all other $s(n) - 1$ equivalence classes predict identically—this is the most asymmetric case. Then we add to each old equivalence class $\mathcal{P}_i(n)$ of round n ($i \neq k$) $\Delta - 1$ new pseudoequivalence classes that predict identically as $\mathcal{P}_i(n)$, which means that we multiply the number of equivalence classes of round n by Δ , to extend the asymmetric splitting in round $n + 1$ to a symmetric splitting.

Note that $\Delta \leq c$ (c being the number of possible predictions); thus, an equivalence class can split into at most c new classes. Therefore, if we add in each round $n + 1$ as many new pseudoequivalence classes as are needed such that number of equivalence classes $m(n)$ is multiplied with c , we are guaranteed to have added sufficiently many equivalence classes to obtain a symmetric extension of the game. Thus, if $m(n)$ is the growth function of proper equivalence classes in an asymmetric game, the corresponding growth function of equivalence classes that guarantees the construction of a symmetric extension is $s(n) = c^{m(n)}$. For this growth function every infinite eMI_{inf} game can be extended to an infinite game with a symmetric split process and identical success profile, which allows its reduction to an ordinary eMI game with $s(n)$ players at the start. This gives us the following theorem:

Theorem 6 (eMI_{inf} in prediction games with infinitely many players). For every infinite prediction game $((e), \mathcal{P} \cup \{eMI_{\text{inf}}\})$ in which eMI_{inf} has infinite cognitive resources, ‘ c ’ is the number of possible predictions and ‘ $m(n)$ ’ the number of predictive equivalence classes produced until time n :

- 6.1. $\text{maxsuc}_n - \text{suc}_n(eMI_{\text{inf}}) \leq \sqrt{2 \cdot (\ln c) \cdot m(n)/n} \cdot (1 + \frac{1}{4 \cdot \sqrt{n}})$.
- 6.2. If $m(n)$ grows slower than linearly with n , that is, $\lim_{n \rightarrow \infty} m(n)/n = 0$, then $\limsup_{n \rightarrow \infty} \text{maxsuc}_n - \text{suc}_n(eMI_{\text{inf}}) \leq 0$.
- 6.3. Results 6.1 and 6.2 apply to rMI_{inf} (with ‘ $\widehat{\text{suc}}_n$ ’ replacing ‘ suc_n ’) and to cMI_{inf} (with ‘ $\widehat{\text{suc}}_n$ ’ replacing ‘ suc_n ’ and adding the regret term $(q - 1)/(2 \cdot k)$).

Proof. By equation (2) and because the symmetric completion of the infinite game with prediction horizon n produces at most $c^{m(n)}$ equivalence classes, an infinite eMI_{inf} game is reducible to a finite eMI game with $c^{m(n)}$ players. This yields result 6.1 by applying theorem 1. Results 6.2 and 6.3 are obvious consequences. QED

Theorem 6 tells us that our infinite versions of metainduction are access optimal if the number of predictive equivalence classes increases less than linearly. This restriction is stronger than that for unboundedly growing numbers of players (sec. 3), which was less than exponential growth.

6. Conclusion and Outlook. In this article the optimality results for meta-induction (Schurz 2008) are developed further, in reply to the challenges of Arnold (2010) and Sterkenburg (2019). According to these challenges the universal optimality theorem is restricted to metainduction over finite sets of prediction methods but does not apply to infinite or unboundedly growing sets of methods. The metainductive approach is defended against this challenge in four steps:

1. Epistemologically, the optimality argument for finitely many methods solves the most important part of the problem of induction, since humans are cognitively finite beings (sec. 2). However, under slight restrictions optimality theorems can even be proved for meta-induction over unbounded sets of prediction methods. In this regard three further results are established:
2. It is shown that universal optimality can even be granted if the number of candidate methods grows unboundedly, provided it does not grow too fast, that is, slower than exponentially (sec. 3). This is a strongly encouraging result for meta-induction in view of the unlimited creativity of the human mind.
3. Moreover, the optimality result for meta-induction can be extended to arbitrary sequences of methods (sec. 4). This extension does not only give as a new solution to Goodman's problem for methods; it also has interesting application to realistic cases of epistemic agents who switch between radically different (e.g., scientific and esoteric) methods in the course of history.⁹
4. For an agent with infinite cognitive resources, meta-induction is even optimal over infinite classes of methods, provided the number of predictive equivalence classes increases less than linearly (sec. 5). This restriction is quite strong; still, the result is impressive. Together with the lesson of Arnold's counterexample, it gives us a sufficient and necessary condition for the infinite access optimality of eM_{inf} , since in Arnold's counterexample the number of predictive equivalence classes increases linearly (every round by one).

REFERENCES

- Arnold, Eckhart. 2010. "Can the Best-Alternative-Justification Solve Hume's Problem?" *Philosophy of Science* 77:584–93.
- Cesa-Bianchi, Nicolo, and Gabor Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge: Cambridge University Press.
- Chernov, Alexey, and Vladimir Vovk. 2009. "Prediction with Expert Evaluator's Advice." In *Algorithmic Learning Theory: 20th International Conference, ALT 2009, Porto, Portugal, October 3–5, 2009*, ed. Ricard Gavaldà, Sandra Zilles, Gábor Lugosi, and Thomas Zeugmann, 8–22. Berlin: Springer.
- Cover, Thomas M., and Joy A. Thomas. 1991. *Elements of Information Theory*. New York: Wiley.
- Goodman, Nelson. 1946. "A Query on Confirmation." *Journal of Philosophy* 44:383–85.
- Kelly, Kevin T. 1996. *The Logic of Reliable Inquiry*. New York: Oxford University Press.
- Mourtada, Jaouad, and Odalric-Ambrym Maillard. 2017. "Efficient Tracking of a Growing Number of Experts." *Journal of Machine Learning Research* 76:1–23.
- Reichenbach, Hans. 1949. *The Theory of Probability*. Berkeley: University of California Press.

9. This possibility was raised as a problem for meta-induction by J. Brian Pitts in the 2017 conference Quo Vadis Selective Scientific Realism? at Durham University.

- Schurz, Gerhard. 2008. "The Meta-inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem." *Philosophy of Science* 75:278–305.
- . 2017. "No Free Lunch Theorem, Inductive Skepticism, and the Optimality of Meta-induction." *Philosophy of Science* 84:825–39.
- . 2019. *Hume's Problem Solved: The Optimality of Meta-induction*. Cambridge, MA: MIT Press.
- Schurz, Gerhard, and Paul Thorn. 2016. "The Revenge of Ecological Rationality: Strategy-Selection by Meta-induction." *Minds and Machines* 26 (1): 31–59.
- . 2017. "A Priori Advantages of Meta-induction and the No Free Lunch Theorem: A Contradiction?" In *Advances in Artificial Intelligence*, 236–48. Lecture Notes in Computer Science 10505. Cham: Springer.
- Shalev-Shwartz, Shai, and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms*. New York: Cambridge University Press.
- Sterkenburg, Tom F. 2019. "The Meta-inductive Justification of Induction: The Pool of Strategies." *Philosophy of Science* 86:981–92.
- Wolpert, David H. 1996. "The Lack of A Priori Distinctions between Learning Algorithms." *Neural Computation* 8:1341–90.