Minds & Machines (2016) 26:31–59
DOI 10.1007/s11023-015-9369-7

The Revenge of Ecological Rationality: Strategy-Selection by Meta-Induction Within Changing Environments

Gerhard Schurz[1] · **Paul D. Thorn**[1]

Received: 24 May 2013 / Accepted: 16 April 2015 / Published online: 8 May 2015

Abstract According to the paradigm of adaptive rationality, successful inference and prediction methods tend to be local and frugal. As a complement to work within this paradigm, we investigate the problem of selecting an optimal combination of prediction methods from a given *toolbox* of such local methods, in the context of *changing* environments. These selection methods are called *meta-inductive* (MI) strategies, if they are based on the success-records of the toolbox-methods. No absolutely optimal MI strategy exists—a fact that we call the "revenge of ecological rationality". Nevertheless one can show that a certain MI strategy exists, called "AW", which is *universally* long-run optimal, with provably small short-run losses, in comparison to any set of prediction methods that it can use as input. We call this property universal *access-optimality*. Local and short-run improvements over AW are possible, but only at the cost of forfeiting universal access-optimality. The last part of the paper includes an empirical study of MI strategies in application to an 8-year-long data set from the Monash University Footy Tipping Competition.

Keywords Prediction task · Adaptive rationality · Strategy selection · Meta-induction · Online learning

✉ Gerhard Schurz
schurz@phil.uni-duesseldorf.de

Paul D. Thorn
thorn@phil-fak.uni-duesseldorf.de

[1] Department of Philosophy, Heinrich Heine University Duesseldorf, 40225 Düsseldorf, Germany

Introduction: Prediction Tasks, Strategy Selection, and Paradigms of Rationality

Prediction (or inference) methods generate predictions about unobserved events or objects, utilizing available information about the given environment. We assume that the *rationality* of a prediction method should be evaluated according to its *objective success rate* in the given environment, which may either be understood as its predictive success rate (truth frequency), or in terms of its epistemic payoff as measured by cognitive costs and predictive gains.

According to the paradigm of *universal rationality*, good prediction strategies should be as general as possible, being applicable to all, or almost all, cognitive purposes and environments. In philosophy, this paradigm was promoted in its *deductive* variant by critical rationalists (Popper 1935) and more recently by formal learning theorists (Kelly 1996), and in its *inductive* variant by logical empiricists (Carnap 1950) and more recently by Bayesian philosophers of science (Howson and Urbach 1996). In *psychology*, the logical variant of this paradigm was represented by the Turing-model of cognition (cf. Wells 2005), and its inductive variant by universal learning theories based on behavioral conditioning or reinforcement (cf. Shanks et al. 1996).

The paradigm of universal rationality has been subjected to serious criticism by the younger paradigm of locally *adaptive* (or *ecological*) *rationality*. Advocates of the latter paradigm argue that good prediction methods are, and should be, adapted to the structure of *local* environments, being tailored to the *specific* tasks for which they provide highly efficient solutions. In philosophy, this paradigm is rather new. In psychology this paradigm was pioneered by Simon (1982) and has been developed by Gigerenzer, Todd, and the ABC research group.[1] This research program operates on the assumption that all successful cognitive methods used by humans are more-or-less local, and that simple heuristics are frequently more successful than computationally costly general reasoning mechanisms, following the slogan "less can be more".

The success of any locally adapted method depends on its being applied in the 'right' *environment*. However, biological organisms and especially humans frequently face *changing* environments. The present paper focuses on the investigation of adaptive rationality in prediction tasks conducted within changing environments. Under such conditions, one needs strategies that *select*, for each relevant environment, a method or a combination of methods that performs as well as possible in that environment. Following Rieskamp and Otto (2006, p. 207), we call this the problem of *strategy selection*. We thereby understand selection strategies as *meta-level* methods which apply their selection strategies to a given *toolbox* of locally adapted methods in the sense of Todd and Gigerenzer (2012, p. 10f). This paper focuses on *meta-induction*, a family of meta-level methods that base their selection strategy on the observed success records of the candidate methods in the toolbox, and on that basis, attempt to select an optimal prediction method, or an optimal combination of such methods.

[1] Cf. Gigerenzer et al. (1999), Todd and Gigerenzer (2012), and Hertwig et al. (2013).

Are There Optimal Selection Strategies? A Challenge for Ecological Rationality

Meta-level selection strategies can only have a cognitive benefit if their success is highly general, applying to a large class of environments and tasks. Indeed, if there were not meta-level methods whose applicability was significantly more general than the object-level methods in the adaptive toolbox, then there would be no point in distinguishing between these two levels at all. In other words, the strategy selection account only makes sense if the meta-level strategies are highly general without being too complex. This raises the question: *Do such methods exist?* In this paper we try to give an answer to this question.

Researchers within the adaptive rationality program acknowledge the importance of the strategy selection problem. For Todd and Gigerenzer (2012, p. 15), the study of ecological rationality centers around the question of which heuristics are successful in which kinds of environments. They propose a list of simple rules that are intended to indicate, for each of their studied heuristics, the kinds of environment in which it may be successfully applied, and in which it may not (ibid, Table 1.1). On closer inspection, however, their rules are not fully adequate. For example, consider the heuristic *take-the-best*, abbreviated as TTB, a well-known meta-inductive rule that will be studied below. TTB monitors predictively relevant cues, and bases its prediction on the cue with the highest validity that discriminates. According to Todd and Gigerenzer (ibid., p. 9), TTB is ecologically rational in environments with high redundancy (positive correlations between cue values) and a high dispersion of cue validities. Similarly, Rieskamp and Dieckmann (2012) report that high cue redundancy favors TTB while low cue redundancy favors weighting methods. However, these generalizations are too good to be generally true. The connection between high cue redundancy and TTB's optimality can be violated in both directions, as demonstrated in Schurz and Thorn (2014), among other sources.[2]

The preceding observations do not diminish the great success of the adaptive rationality program in discovering many surprising *less is more* effects. They rather point towards an underdeveloped area of this program, namely the selection-of-methods problem. They also indicate a major *challenge* (if not a dilemma) for the program of ecological rationality. Indeed, *if there were* simple rules of the form "In environment of type E_i, method M_i is optimal" (for all E_i in a partition of environment types $\{E_1, \ldots, E_n\}$), then the combined strategy "For all $i \in \{1, \ldots, n\}$: apply method M_i in environment E_i" would be a universally optimal (combined) strategy. The existence of such a strategy would, thereby, *re-install* universal rationality, and undermine the very program of adaptive rationality.

[2] The lack of connection between a high cue correlation and TTB's success is also reported in Czerlinski et al. (1999, p. 116f). The implication between high cue-validity dispersion and TTB's optimality holds only in "naive Bayes" environments (cf. fn. 13 and Katsikopoulos and Martignon 2006, Corollary 1). Gigerenzer and Brighton (2009, p. 143) and Brighton and Gigerenzer (2012, p. 55) describe an environment with zero cue validity dispersion (a so-called Guttman environment) in which TTB works particularly well.

Can universal rationality be re-installed in this simple way? The simple answer is: *No*. Following from well-known results in formal and computational learning theory (see next section), there cannot be a prediction method (be it an object-level or meta-level method) that is *absolutely* optimal, i.e., optimal in *all* environments among *all* possible prediction methods. This fact is frequently mentioned in work on adaptive rationality (cf. Martignon and Hoffrage 1999, p. 128; Todd and Gigerenzer 2012, p. 5). Therefore, there cannot be exhaustive and fully general meta-rules which specify for each task and environment a locally optimal method. In what follows, we call this fact the *revenge of ecological rationality*.

While there is no absolutely optimal selection strategy, the ecological rationality program presupposes selection rules whose success is at least very or sufficiently general. If such rules did not exist, one could not explain why humans are so successful in selecting the 'right' method for their given environment, in spite of the fact that their environment constantly changes. What makes it difficult to find general rules for selecting methods is that the success-relevant features of the environment are frequently cognitively inaccessible. Similarly, *changes* in the environment are often unrecognizable and unforeseeable: Consider the transitions between expansion and recession phases in the market economy. To deal with changing environments of this sort, one needs strategies for *learning* which object-level methods perform best in which environment, and in which temporal phases of the environment. This brings us to the account of strategy selection by *meta-inductive learning*. While Rieskamp and Otto (2006) recommend *reinforcement* as the learning method for strategy selection, *meta-induction* demarcates a more general family of selection strategies, which includes reinforcement as a special case.

The account of meta-induction was developed within the domain of epistemology as a means of addressing Hume's problem of induction (Schurz 2008, 2009; Arnold 2010; Vickers 2010, §6.3), utilizing results from the area of machine learning (Cesa-Bianchi and Lugosi 2006). Meta-inductive methods have also been applied in the area of social epistemology (Schurz 2012; Thorn and Schurz 2012; Feldbacher 2012). The simplest meta-inductive strategy is *Imitate-the-best* (ITB), and its relative *Take the best* (TTB), which imitate the predictions of the so far best available prediction methods. More elaborate meta-inductive strategies predict a weighted average of the predictions of the methods that have been successful so far, using different weighting methods (see below).

The method TTB has been investigated in a variety of studies with the well-known result that under certain conditions TTB is predictively more successful than more complex prediction methods with higher information demands.[3] In almost all of these studies, TTB was investigated as a prediction method based on (a) binary-valued cues, under the assumption that (b) the cue validities are either known or else estimated by random sampling from an environment whose probability distribution

[3] See (among others) Gigerenzer et al. (1999), ch. III, Brighton and Gigerenzer (2012), Rieskamp and Dieckmann (2012), and Katsikopoulos et al. (2010).

does not change in time.[4] In contrast, our investigation involves two major shifts in the problem setting:

1. Instead of assuming that cue validities are known or estimated by random sampling, we assume a situation of *online learning* conducted within a possibly changing environment, and
2. We pass from the object-level perspective to the *meta-level perspective*, applying meta-inductive strategies such as TTB not only to the selection of cues, but more generally to the selection of (object-level) prediction methods of any sort.

From the object-level perspective, one compares prediction methods that are based on a set of cues C1, …, Cn. For example, the method TTB uses the best cue, while the method SW (also called "Franklin's rule") uses a success-dependent weighting of cues.[5] The cues are themselves *predictive indicators* of a criterion variable whose values or value-relations must be predicted. Each cue has a given *probability* of predicting correctly, conditional on its delivering a prediction at all. This conditional probability is called the cue's (ecological) *validity* (Gigerenzer et al. 1999, p. 84). In one well known experiment that illustrates this perspective, the task was to predict which of two German cities has a higher population, based on binary cues such as (C1) whether it is a national or state capital, (C2) whether it has a first division soccer team, etc. In experiments of this sort, a cue (Ci) delivers a prediction if it 'discriminates' between the two compared cities: For example, if the cue has the value 1 for city A and 0 for city B, then it predicts that city A has a larger population than city B. We will see in the "Take-the-Best (TTB)" section that this *comparative* prediction format can be generalized to the format of arbitrary non-comparative prediction tasks (e.g., weather forecasts) without incurring any loss of intelligible results.

When we pass from the object-level perspective to the meta-level perspective, we (1) *re-interpret cues as locally adapted prediction methods*, which may either be simple cues, real-life experts, or computational methods, and (2) re-interpret TTB and related meta-induction methods as meta-level strategies for selecting optimal prediction methods.[6] Prima facie, this is only a change in perspective, consisting of a redescription of cues as prediction methods. Does it bring anything new?

The major innovation deriving from this change in perspective consists in raising new questions, concerning the success of meta-level strategies in comparison to the locally best object-level methods or cues. Within the adaptive rationality research program, these questions have not been asked, so far, presumably because cues were

[4] Exceptions to (a) are Hoffrage et al. (2000), Hogarth and Karelaia (2005) and Katsikopoulos et al. (2010), who study prediction tasks based on continuous-valued cues. Exceptions to (b) are Dieckmann and Todd (2012), and Rieskamp and Otto (2006), who study prediction tasks in the course of online learning.

[5] Other frequently studied methods are Dawes' rule (equal weights), regression (optimal weights), and "naive Bayes" (Gigerenzer et al. 1999, part III).

[6] A related idea is anticipated in Katsikopoulos and Martignon (2006, p. 491), who interpret a cue as a juror voting for one of two options in a social choice task.

not regarded as prediction methods. To understand the significance of such questions, recall the negative result that there is no method which is absolutely optimal. However, only a (finite) fraction of all possible prediction methods are *cognitively accessible* to any human-like agent. This raises the following new question: Is there a meta-inductive strategy which predicts optimally in comparison to all candidate prediction methods which are *accessible* to it, no matter what these methods are, and in what environment one happens to be? In what follows we call this property *access-optimality*, in distinction to *absolute* optimality, which is not restricted to the accessible methods. The philosophical importance of access-optimality lies in the context of the problem of induction. Before we turn to this problem, we make some notions, that have been informally introduced, logically precise.

Dynamic Online Learning, Prediction Games, and Access-Optimality

In online learning, one must simultaneously predict future events and learn from the results of past predictions. In this paper, we study online learning conducted under possibly changing environments, henceforth called *dynamic online learning*. In Schurz (2008), the study of dynamic online learning was developed on the basis of prediction games. A *prediction game* ((e), Π) consists of:

1. An infinite sequence (e) := (e_1, e_2, …) of events $e_n \in [0,1]$, which are coded by real numbers between 0 and 1. For example, (e) may be a sequence of daily weather conditions, football game results, or stock values. Each time n corresponds to a round of the game.
2. A finite set of prediction methods or 'players' $\Pi = \{P_1, \ldots, P_m, MIx\}$, whose task, at each time n, is to predict the next event of the event sequence. "MIx" signifies a meta-inductivistic player of a certain "type x". The other players are called the "non-MI-players"; they form the 'toolbox' of MIx's candidate methods or cues, and may include real-life experts, 'virtual players' implemented by computational algorithms, and 'para-normal' players, e.g., clairvoyants who may be successful in para-normal worlds.

A prediction game corresponds to an *environment*, or *possible world*, where the event-sequence constitutes the *natural* part of the environment, and the player set the *social* part. We identify each method with one player; this is possible because prediction methods are evaluated according to their success, and not according to the number of their adherents. While all possible event-sequences are allowed, we assume that the player set is *finite*, because finite cognitive beings can only compare finitely many different methods.

Binary prediction games are a subcase of *real-valued* games in which the predictions and events take the value 0 or 1. We also distinguish between prediction games with *persistent* and *intermittent* players. Persistent players deliver a prediction for each event, while intermittent players do not. This may be the case because a player refrains from predicting in the given round, because she predicts but her prediction is inaccessible to the meta-inductivist, or because the prediction task is comparative in nature, being based on a pair of cue values (as in the studies mentioned in the previous section), and the cue doesn't discriminate.

Further notation that we use later includes:

- $\text{pred}_n(P) \in [0,1]$ is the prediction of *player* P *for* time n, which P delivers *at* time $n - 1$. In binary games, $\text{pred}_n(P) \in \{0,1\}$.
- $\text{loss}_n(P)$ is the *loss* of P for time n. Most natural is the *linear* loss-function, defined as $|\text{pred}_n(P) - e_n|$. The linear loss of a binary prediction is 0 if it is correct, and 1 if it is incorrect. (The theorems presented below do not presuppose that the respective loss functions are linear.)
- $\text{score}_n(P) =_{\text{def}} 1 - \text{loss}_n(P)$ is the score which player P earns for $\text{pred}_n(P)$.
- $\text{abs}_n(P)$ is the absolute success achieved by player P at time n, defined as the sum of P's scores for predictions *delivered* until time n.
- $\text{predfreq}_n(P) =_{\text{def}}$ the relative frequency with which player P delivered predictions until time n.
- $\text{val}_n(P) =_{\text{def}} \text{abs}_n(P)/(\text{predfreq}_n(P)\cdot n)$ is the (ecological) validity of player P at time n, i.e., P's success rate conditional on times at which P delivered a prediction.
- $\text{suc}_n(P) =_{\text{def}} \text{predfreq}_n(P)\cdot\text{val}_n(P) + (1 - \text{predfreq}_n(P))\cdot\text{ran}$. This is P's unconditional success rate at time n, where P's score for non-predictions is identified with the score of a *random guess*, denoted as "ran". The same convention is adopted by Martignon and Hoffrage (1999, p. 131).

The validity of a persistent player coincides with her success rate. In contrast, the validity of an intermittent player is greater than her success rate, provided her predictions are better than random guesses, because her non-predictions are scored as random-guesses. The latter convention is important for making a *fair comparison* of the success rates of persistent and intermittent players. Otherwise, a player could improve her success rate by refusing to predict when her uncertainty is high. Given the proposed convention, a player cannot improve her success rate by selectively refraining from prediction, but only her validity.

- limsuc(P) is P's limit success rate $=_{\text{def}} \lim_{n\to\infty} \text{suc}_n(P)$, provided the success frequencies converge. Limiting (event, success) frequencies, if they exist, are identified with (event, success) *probabilities*, written as "p(−)".
- maxsuc_n is the maximal success rate of the non-MI-players at time n, and maxlimsuc is their maximal limit success, provided the success frequencies converge.

In the setting of dynamic online learning, event frequencies, validities, and success rates may change over time. They may even fluctuate permanently, so that their relative frequencies *do not* converge to probabilities. Therefore validities and success rates are always relativized to certain *time* points of a prediction game. This is significantly different from learning by random sampling where sample frequencies will certainly converge with increasing sample size.

We define our basic notion of optimality in the manner that is usual in decision and game theory (Weibull 1995, p. 13). These notions are independent of any assumptions about probability distributions, but they are relativized to a class $\mathcal{E}$ of prediction-games or environments.

Definitions of optimality and dominance

1. A method M* is called optimal with respect to (w.r.t.) a class $\mathcal{E}$ of environments iff for all ((e), Π) in $\mathcal{E}$: M* $\in \Pi$ and for alternative methods M $\in \Pi$ (i.e., M $\neq$ M*), M* is *at least as good* as M in ((e), Π).—"M* is at least as good as M" can be understood in two ways:

   1.1. In the *long-run*, meaning that $\lim_{n\to\infty}(\mathrm{suc}_n(M^*) - \mathrm{suc}_n(M)) \geq 0$; we speak here of "long-run optimality".

   1.2. In the *short-run*, meaning that for all n, $\mathrm{suc}_n(M^*) \geq \mathrm{suc}_n(M) - L(n)$, where L(n) is a 'sufficiently small' short-run loss that converges 'sufficiently fast' to zero for $n \to \infty$. We speak here of "approximate short-run optimality". This is a vague notion. Only "strict short-run optimality" (defined by L(n) = 0 for all n) is a sharp notion, but this is *too good* to be achievable by meta-inductive methods.

   Note: The optimality of M* w.r.t. $\mathcal{E}$ does not exclude that there exists some other method M' that is 'equally optimal' w.r.t $\mathcal{E}$. This is excluded within the stronger notion of *dominance.*

2. A method M* is called *dominant* w.r.t. $\mathcal{E}$ iff M* but no other method M is optimal w.r.t. $\mathcal{E}$.—*Note:* (a) This implies that for every M $\neq$ M* there exists an environment ((e), Π) $\in \mathcal{E}$ such that either M $\notin \Pi$, or M* is better than M in ((e), Π). (b) If definition 2. does not hold for every M, but only for every M in a given class $\mathcal{M}$ of methods, we say that M* is called *dominant* w.r.t. $\mathcal{E}$and $\mathcal{M}$.

A brief classification of methods

A method is *independent* if its predictions depend only on the events, but not on the predictions of the other players. Non-independent methods are called *social* methods. A method is *object-inductive* or *meta-inductive*, respectively, if it uses *some* kind of inductive method to infer future from past events, or future from past successes, respectively. A method is *normal* if its predictions depend only on *past* events or successes. A method is *para-normal* or *clairvoyant* if it has 'privileged' access to future events. Note that the admission of clairvoyant methods is needed in the philosophical context of the problem of induction, but not in the naturalistic setting of cognitive science.

Each *normal* method (or player) P can be extensionally identified with a function $f_P: \cup_{n\in \mathbb{N}}\Omega^n \to \Omega$ which maps each length-n history of elements of an event space Ω into a prediction of the next event e_{n+1}. The elements of Ω are ordinary events if P is an individual method, as in formal learning theory (Kelly 1996, p. 260f), and they consist of events together with the other players' predictions if P is a social method.[7]

[7] A clairvoyant P 'sees' the future, and can be identified with a function $f_P: \mathbb{N} \times \Omega^\infty \to \Omega$.

We mentioned above that no normal method is *absolutely optimal*, i.e., optimal with respect to the class of all environments. In the framework of prediction games, this result is easily proved: Let f_P be the prediction function of a normal method P. For every such method f_P (and for every set of competing players, if P is a meta-method) one can define a "f_P-demonic" event-sequence (e*), which produces for each time n a worst-score event (defined as $e_n^* = 0$ if $pred_n(P) > 0.5$; else $e_n^* = 1$). Moreover, one can define a method f* which perfectly predicts the event-sequence (e*). So f_P is not absolutely optimal.

Observe that if P is a meta-method, the preceding proof only goes through if P does not have *access* to f*'s predictions. If f_P can imitate f*'s predictions, it is no longer generally possible to construct an event-sequence (e*) which deceives f_P and at the same time rewards f*. This simple but crucial fact underlies all the results concerning access-optimality.

Definitions of access-optimality

3.1. A method M is accessible to a (social) method M* iff M* can observe—or otherwise simulate—M's present predictions and past predictive successes.
3.2. A (social) method M* is *access-optimal* w.r.t. $\mathcal{E}$ iff M* is optimal w.r.t. the class of all environments ((e), Π) in $\mathcal{E}$ in which all methods in Π are accessible to M*.
3.3. M* is *universally access-optimal* iff M* is access-optimal w.r.t. the class of *all* environments.

We now return to our major question: Is there a meta-level strategy that is universally access-optimal in the long-run? The answer is trivially *Yes*, if the success rates (or validities) of the non-MI-methods or cues converge to probabilities and are *known* in advance. In this case, applying TTB at the meta-level selects, in each environment, the best method in the 'toolbox' of candidate methods, whence TTB's predictive success is guaranteed to be access-optimal. So under the condition of *known* success rates the question of access-optimality is not particularly interesting.

The situation is decisively different when the cue validities are not known but must be learned by the uncertain methods of *induction*. Since inductively estimated cue validities may diverge from the true ones, neither TTB nor any other meta-inductive strategy is guaranteed to always select the best method, or combination of methods. Now the question of the existence of a universally access-optimal method becomes entirely non-trivial. This question is of particular importance for the *problem of induction*, as we shall explain at the end of the section on "Attractivity-Weighted Meta-Induction".

There are, however, two significantly different ways by which inductive inferences may be applied. The first is the method of *random sampling* that is applied in much of the research on adaptive rationality (cf. fn. 3). This method is only applicable under two "induction-friendly" conditions: first, samples are drawn from a population whose frequencies (or frequency-limits) don't change within the time window of the inductive experiment, and second, all individuals in the population (or events in the event-sequence) have the same chance of appearing in the sample or 'training set'. This means probabilistically that the sample distribution is IID (independent and identically distributed). The parameters found in the training set are inductively projected to the remainder of the population, the test set. Following from the laws of IID random sampling, the success rates (or validities) estimated from the training set will deviate from the true success rates in the test set by a symmetrically distributed error variance. Because of this error variance, complex prediction strategies such as success-based weighting and linear regression tend to *overfit* the frequencies found in small samples, i.e., they fit to random variations instead of representative tendencies of samples. This explains why complex strategies, if based on small samples, frequently perform worse than TTB and other frugal prediction methods (cf. Brighton and Gigerenzer 2012, p. 36, 44f; Katsikopoulos et al. 2010).

The second way of applying inductive inference is investigated in this paper: dynamic online learning, as modeled by prediction games. For many real-life-situations, dynamic online learning is the more realistic learning situation compared to random sampling, because it is also applicable in the absence of induction-friendly conditions. It differs from random sampling in three respects. First of all, in dynamic online learning, *past* observations are inductively projected into the *future*: Since one can only observe past but not future events, not all members of the event sequence have the same chance of entering the observed sample. Secondly, there are no delineated training and test phases. Admittedly one may artificially distinguish between a training and test phase at each round, by considering the observations made so far as the training set, and the predicted event as the test set. But thirdly—this makes the crucial difference—the success rates (and validities) of the cues may systematically *change* in time, i.e., the future may be systematically different from the past. Since after the training phase has passed the success rates may have already changed, dynamic online learning requires constant updating of the inductively projected success rates.

In probabilistic terms, the 'sampling' procedure in dynamic online learning not only generates an error variance, as in the case of random sampling, but may also generate a systematic bias (cf. Brighton and Gigerenzer 2012, p. 46f). This systematic bias manifests itself in the form of a correlation between the (event or success) frequencies in the training phase and those in the test phase, which creates difficulties when the correlation is negative. As we shall see, the difficulties generated by such correlations affect simple prediction methods, such as TTB, as well as complex ones.

A correlation between past and future events means that the underlying event sequence is a Markov chain (whose elements are not IID). This possibility is admitted in the framework of dynamic online learning. It is even admitted that the event or success frequencies do not converge to a limit but oscillate forever. In that case, the sequence is not generated by a probabilistic source, and all that one can study are its finite frequencies.

The study of dynamic online learning of individual sequences is complementary to Bayesian learning. In addition to meta-inductive approaches, methods of dynamic online learning were developed within formal learning theory (Kelly 1996) and, under the rubric of "online learning under expert advice", within computational learning theory (Cesa-Bianchi and Lugosi 2006). Our notion of *access-optimality* is a generalization of the property known as "Hannan-consistency" (Cesa-Bianchi and Lugosi 2006, p. 70). The major advantage of the framework of dynamic online learning is that its key results are independent of particular assumptions about prior probability distributions (which does not exclude that these results can be extended by such assumptions, see below). In contrast, key Bayesian results depend on assumptions about prior probabilities: This point is substantiated at the end of the section on "Attractivity-Weighted Meta-Induction". Some Bayesians argue that our prior probabilities should be adapted to our local environment. However, prior to inductive inference from experience we have no clue which prior distribution fits with our (future) environment. In other words, prior probabilities are not empirically grounded but subjective in nature (cf. Schurz 2013, ch. 4.7).

In the following three sections, we present the major theoretical results concerning meta-induction. From this point on, we always assume that the non-MI-players of the prediction game are accessible to MIx.

Imitate-the-Best (ITB)

A simple yet surprisingly efficient meta-inductive method is "Imitate-the-best", ITB, which predicts what the non-MI-player with the presently highest predictive *success rate* predicts; this player is called ITB's present *favorite*. ITB changes her favorite only if another player becomes *strictly* better. If there are several best players, ITB chooses the first-best player according to an arbitrary ordering. If her favorite fails to deliver a prediction, ITB predicts according to a random guess. ITB's initial prediction (at time 0) is a random guess.

Theorem 1 tells us that ITB is (access-) optimal, not in all environments, but in the class of those environments where the leading method remains constant after some "winning time" w, i.e., where there exists a P_k such that for all $n \geq w$, $suc_n(P_k) > suc_n(P_i)$, for all $i \neq k$:

Theorem 1 *For each prediction game ((e), {P_1, …, P_m, ITB}) whose set of non-MI-players contains a best player P_k after winning time w, the following holds:*

1.1. *(Short-run) For all $n > w$: $suc_n(ITB) \geq maxsuc_n - (w/n)$.*
1.2. *(Long-run) ITB's success rate approximates the maximal success rate: $lim_{n\to\infty}(maxsuc_n - suc_n(ITB)) = 0$.*

Theorem 1 holds for all monotonic loss functions, where $loss_n(P)$ is a strictly positive monotonic function of $|pred_n(P) - e_n|$ (the same holds for Theorem 3 concerning TTB). Theorem 1.2 implies that ITB is access-optimal in the long run in all environments containing a best non-MI-player after some winning time w. Figure 1 shows the result of a computer simulation of a binary prediction game of the sort described by Theorem 1: ITB always imitates the best player, which changes from the object inductivist to Player 2.

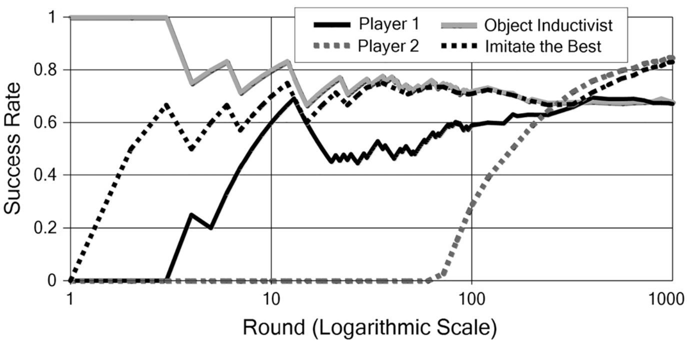

Fig. 1 ITB, an object inductivist, and two alternative success-convergent players

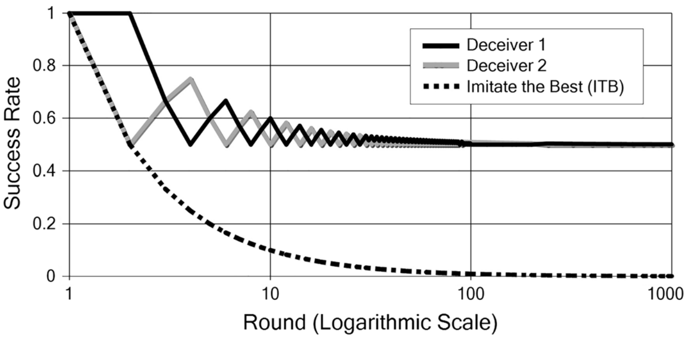

Fig. 2 Binary prediction games with ITB against two deceiving players with convergent success-oscillations

Whenever ITB switches her favorite, ITB has to buy a loss of (at most) one score-point compared to the best player. These losses may accumulate. The existence of a winning time w excludes the possibility that these losses may grow infinitely, because after time w, ITB will no longer switch her favorite. ITB's worst case short run loss is (w/n), which is only small if the winning time is small.

The optimality of ITB breaks down whenever the success rates of two or more leading non-MI-players oscillate endlessly around each other in negative correlation to their position of being ITB's favorite. This situation is programmed in Fig. 2 in a binary prediction game with two success-oscillating 'deceiving' players who each predict incorrectly exactly when they become ITB's favorite. As a result, ITB's success rate goes to zero while the success rate of each deceiver converges to 1/2.

The success oscillations in Fig. 2 are called *convergent* because the oscillation amplitudes are decreasing so that the oscillating success rates converge to the same limit. A simple modification of ITB yields a form of meta-induction that is resistant to convergent success oscillations: a conservative variant of ITB with a small switching threshold ε, abbreviated as εITB, which switches her favorite only if the success difference between her present favorite and the new better player exceeds ε.

It is demonstrable that εITB ε-approximates the maximal success rate in the long run, in all prediction games whose success rates converge to limiting frequencies (cf. Schurz 2008, Theorem 2).

Yet the access-optimality of εITB breaks down in games with *non-convergent* success-oscillations of the non-MI-players. The worst case involves *systematic deceivers* who are assumed to know when the meta-inductivist will choose them as favorite. Systematic deceivers deliver a false or minimal-score prediction whenever they are ITB's favorite, otherwise they predict as accurately as possible. In the presence of εITB, the success rates of systematic deceivers oscillate around each other with a non-diminishing amplitude greater than ε, with an exponentially increasing time period. Figure 3 represents a computer simulation of a prediction game pitting εITB (and AW, to be introduced below) against four systematic deceivers: While εITB's success rate converges to zero, the deceivers' success rates approximate 3/4 (the success curve of the method AW will be explained in the section on "Attractivity-Weighted Meta-Induction".

Theorem 2 informs us that the negative results for ITB and εITB (Figs. 2, 3) generalizes to all *one-favorite* meta-inductive methods, which (by definition) imitate, at each time point, the prediction of a single non-MI-player (or cue). More generally, one-favorite meta-inductivists cannot be long-run optimal if the success rates of the non-MI players are negatively correlated with their position of being MIx's favorite. A Proof of Theorem 2 is found in the Appendix.

Theorem 2 *For every prediction game {(e), {P_1, …, P_m, MIx} in which MIx is a one-favorite meta-inductivist:*

2.1. *If for all $i \in \{1, \ldots, m\}$, P_i is a systematic deceiver, then (a) limsuc(MIx) = 0, but (b) limsuc(P_i) = $(m - 1)/m$.*

2.2. *If for every $i \in \{1, \ldots, m\}$, P_i's limiting success rate exists and is negatively correlated with P_i's position of being MIx's favorite, MIx is not long-run optimal, i.e., maxlimsuc − limsuc(MIx) > 0.*

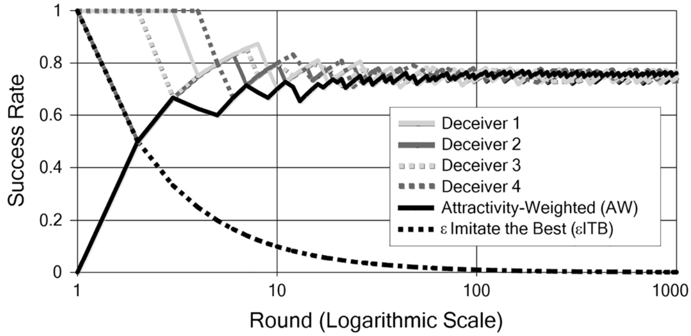

Fig. 3 Binary prediction game with εITB and AW against four systematic deceivers with non-convergent success-oscillation (ε = 0.05)

Theorems 1 and 2 also have important implications for prediction tournaments that are not based on online learning but on random sampling. Following from the laws of IID random sampling, probabilistic correlations between past events or predictions (training set frequencies) and future events (test set frequencies) are impossible in the long run. This implies that ITB is guaranteed to be access-optimal in the long run, if the sequences of events and scores are IID sequences.

It may be objected that long run results are irrelevant for humans making practical decisions. Similarly, ITB's short-run performance is determined by a 'winning time' which may occur arbitrarily late. But keep in mind that Theorem 1 reflects the worst possible case. Under more induction-friendly probabilistic assumptions, one can obtain better short-run results for ITB. Assume, for simplicity, a prediction game with two players P_1 and P_2 whose success-scores are generated by IID probabilities, and the difference between the success probabilities of P_1 and P_2 is 2δ. Let $p_{1>2}$ denote the probability that time point n is a winning time for P_1. Then $p_{1>2} = (1 - p_\delta)^2$, where P_δ is the probability that for some $m \geq n$, $|suc_m(P_i) - p_i| > \delta$ holds (with $p_i = limsuc(P_i)$). p_δ can be shown to be upper bounded by $(c/n^{0.5}) \cdot x^n/(1 - x)$, with $c = 2/\sqrt{2 \cdot \pi \cdot \delta}$ and $x = 1/e^{0.5 \cdot \delta^2}$.[8] For $n = 1000$ and $\delta = 0.1$ this gives $p_{0.1} \leq 0.1$ and $p_{1>2} \geq 0.9^2 = 0.81$. In the worst case ITB favors P_2 until time 1000, and by IID-laws $suc_{1000}(P_2) > limsuc(P_2) - 0.04$ is almost certain. So with $p > 81\ \%$ ITB's short run loss after n rounds is smaller than $0.24 \cdot (1000/n) = 240/n$.

Concerning the limitative results of Theorem 2 for ITB and other kinds of one-favorite meta-induction, one could argue that in real-life situations, adversarially fluctuating success rates almost never occur. However, they *do* occur in social environments in which predictions have a negative feed-back effect on the social structure to be predicted—which is one of the possible causes of biased sampling in real environments. An example which illustrates this possibility concerns the prediction of stock values in a so-called *bubble economy*: Here it is predicted that a given stock will yield a high rate of return, which leads many investors to put their money in this stock, and by doing so cause it to crash (since this stock lacks sufficient economic support). In such a situation it would be a bad recommendation to always put all of one's money into the stock that is presently most successful (which would be ITB's strategy), instead of distributing it over several stocks. The latter strategy corresponds to the weighting methods which are discussed in the sections on "Attractivity-Weighted Meta-Induction" and "Local Improvements".

[8] For fixed n we approximate $p(|suc_n - p| \geq \delta) \approx c/(n^{0.5} \cdot e^{0.5 \cdot n \cdot \delta^2})$ (see de Finetti 1974, sect. VII.5.4). p_δ is upper bounded by the infinite sum $c \cdot \Sigma_{n \leq i \leq \infty}(1/(i^{0.5} \cdot e^{0.5 \cdot i \cdot \delta^2}))$. This sum is lower-equal $(c/n^{0.5}) \cdot \Sigma_{n \leq i \leq \infty} x^i$), which is (by the sum-formula for a convergent geometric series) equal to $(c/n^{0.5}) \cdot x^n/(1-x)$.

Take-the-Best (TTB)

With minor modifications all results about "Imitate-the-Best (ITB)" generalize to the strategy Take-the-best (TTB) within games with intermittent players. In such games, TTB selects at every time n that player as its favorite whose *validity* is greatest among those players who delivered a prediction (with ties resolved according to an arbitrary ordering). If no player delivers a prediction, TTB predicts according to a random guess.

In other words, TTB is the intermittent version of ITB. In what follows, we generally write "iMIx" for the intermittent version of the meta-inductive strategy MIx. Thus, TTB = iTB. The intermittent strategy iMIx differs from the persistent strategy MIx only in games with intermittent players: Here iMIx bases its success evaluation on validities, while MIx bases its success evaluation on success rates, scoring non-predictions as random guesses. We speak here of *intermittent* versus *persistent* success evaluation.

Within research on adaptive rationality, prediction competitions have been frequently performed in the following comparative format: 'Players' must predict which of two randomly selected objects A, B has a higher value on a given criterion variable X (e.g., city size), based on various binary cues $Ci_{A/B} \in \{1,0\}$ for the two objects A, B (e.g., having a first division soccer team). The *cue-difference* $\Delta_i =_{def} Ci_A - Ci_B \in \{1,0,-1\}$ is taken as the prediction variable: The cue Ci predicts $X_A > X_B$ if $\Delta_i = 1$ and $X_B > X_A$ if $\Delta_i = -1$. If $\Delta_i = 0$, the cue C_i doesn't deliver a prediction (note that draws $X_A = X_B$ are assumed to be absent). Table 1 lists the correspondences between these formats.

From Table 1, we see that comparative prediction tasks can be subsumed under binary intermittent prediction tasks, where 1 is mapped to 1, −1 is mapped to 0, and 0 is mapped to "no prediction". For this reason, all formal results concerning binary prediction games apply to comparative ones. We therefore expect that the shift from the comparative to the binary format doesn't cause any significant changes in results. The shift from binary to real-valued prediction games is more significant (see the section on "Attractivity-Weighted Meta-Induction", and the last section).

Independent of the prediction format, TTB's success rate converges (provably) to a weighted average of the validities of the intermittent non-MI-players, if these

Table 1 Correspondences between different formats of prediction games

Format	Real-valued	Binary	Comparative
Events e:	Event-interval	Yes–No	Which has greater X-value?
	$e \in [0,1]$	$e \in \{0,1\}$	$e \in \{X_A > X_B, X_A < X_B\}$
Prediction of player/cue P:	$\text{pred} \in [0,1]$	$\text{pred} \in \{0,1\}$	$\text{pred} = X_A >/< X_B$ iff $\Delta P = 1/-1$
			where ΔP is P's cue-difference
Intermittent:	Sometimes no pred. is delivered by P		No pred. is delivered if $\Delta P = 0$
Score of pred.:	$1 - \lvert e - \text{pred}\rvert$	1 if pred. is correct, otherwise 0	
Success rate:	Average score	Frequency of correct predictions	

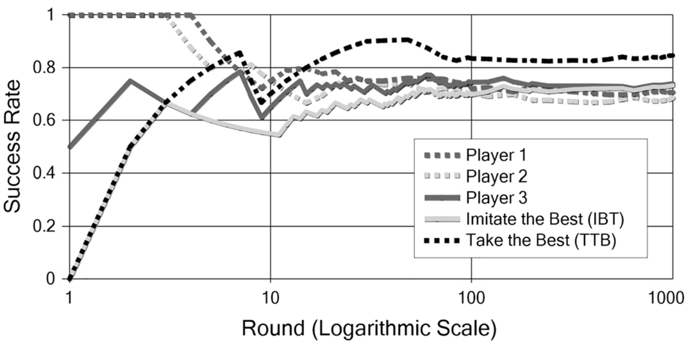

Fig. 4 TTB in a binary prediction game playing against three intermittent non-MI-players delivering a prediction in 60 % of all rounds

validities are convergent and conditionally independent from each other. Figure 4 illustrates the performance of TTB in a binary prediction game satisfying this condition. Since the non-MI-players' validities are greater than random successes, TTB's success rate is significantly greater than the (random-guess-updated success rates of the non-MI-players. TTB's success is also greater than the success rate of ITB, which imitates the 'random guess' of the so far best player, whenever this player doesn't deliver a prediction, rather than searching (as TTB) for the next-best player who delivers a prediction. So TTB enjoys an additional advantage over ITB in prediction games with intermittent players.

Theorem 3 generalizes the result of Fig. 4. According to this theorem, the meta-level optimality of TTB requires that after a certain convergence time the entire success-ordering of the non-MI-players becomes stable, and not only the top non-MI player (as in the case of ITB):

Theorem 3 *For each prediction game ((e), {P_1, ..., P_m,TTB}) with intermittent non-MI-players who converge to a stable validity ordering* $\mathrm{P}_{s_1}, \ldots, \mathrm{P}_{s_m} (1 \leq s_i \leq m)$ *after a stabilization time s, and whose conditional validities (see below) are better than the success of a random guess after s, the following holds:*

3.1. *(Short-run)*
(a) $\mathrm{suc_n(TTB)} \geq \Sigma_{1 \leq k \leq m} \mathrm{val_n(P_{s_k}|fav\text{-}P_{s_k})} \cdot \mathrm{freq_n(fav\text{-}P_{s_k})} - (s/n)$.
Here "$\mathrm{freq_n(fav\text{-}P_{s_k})}$" *is the relative frequency of times (until time n) for which player* P_{s_k} *was TTB's favorite, i.e., for which the kth best player but no better player delivered a prediction, and* "$\mathrm{val_n(P_{s_k}|fav\text{-}P_{s_k})}$" *is the success rate of* P_{s_k} *(at time n) conditional on those times.*
(b) $\mathrm{suc_n(TTB)} > \mathrm{maxsuc_n} - (s/n) \geq \mathrm{suc_n(ITB)} - (s/n)$.

3.2. *Long run* results are obtained from 3.1 by omitting "s/n" and replacing finite frequencies by their limits. For example, instead of (3.1)(b) we have: limsuc(TTB) > maxlimsuc ≥ limsuc(ITB).

A Proof of Theorem 3 is found in the Appendix. Note that, according to Theorem 3, what is relevant for TTB's predictive success is not the players' validities simpliciter but their validities conditional on the times for which no better player delivered a prediction. In the case where the players' validities are conditionally independent, one may replace the conditional validities in Theorem 3.1 by the unconditional validities, $val_n(P_{s_k})$ (cf. fn. 13).

The modified TTB-method which uses the conditional validities instead of the unconditional ones is called "greedy TTB" by Schmitt and Martignon (2006). They argue that greedy TTB is provably better that ordinary TTB. But this only holds if the conditional validities are known. Brighton and Gigerenzer (2012, p. 44) show that in situations of learning from not too large samples, greedy TTB is less successful than simple TTB. The reason for this is the already mentioned problem of *overfitting:* Greedy TTB frequently fits its estimated conditional validities to random accidentalities of the sample which disappear when the samples size grows larger. A similar effect arises in situations of online learning, since conditional validities converge more slowly than ordinary validities.

The short-run performance of TTB depends on the stabilization time for ordinary validities: If it occurs late, TTB's short-run performance will be bad. As in the case of ITB in Fig. 2, TTB's performance is particularly bad in prediction games with players whose success rates are negatively correlated with their being imitated by TTB. More precisely, Theorem 2 holds also for TTB, because TTB is a one-favorite meta-induction method.

Attractivity-Weighted Meta-Induction (AW) and the Problem of Induction

Are there meta-inductivist strategies which can handle systematic deceivers and, hence, are universally access-optimal? *Yes*, although only in the long run. Generally speaking, all *weighted* meta-inductive methods predict a weighted average of the predictions of the non-MI-players. The crucial property of universally access-optimal meta-inductive methods is that the weights of the non-MI-players are identified with their *attractivities*, which from MI's viewpoint are called *regrets* (cf. Cesa-Bianchi and Lugosi 2006, Sect. 2.1). Attractivity weights have so far not been studied in the ecological rationality literature. In what follows, we use "AW" to denote the attractivity-weighted meta-inductivist. The *attractivity* of a player P for AW, at a given time, is the surplus of her success rate compared with AW's own success rate at that time:

$$at_n(P) =_{def} suc_n(P) - suc_n(AW),\ \text{if this expression is positive; else}\ at_n(P) = 0.$$

AW's predictions are defined as follows:

$$pred_{n+1}(AW) = \frac{\sum_{1 \le i \le m} at_n(P_i) \cdot pred_{n+1}(P_i)}{\sum_{1 \le i \le m} at_n(P_i)}$$

When the denominator is zero, AW's prediction is identified with the prediction of the player with maximal success rate at time n[9] (which at $n = 1$ or if no player predicts is identified with a random guess).

Even if the success rates of adversarial players oscillate around each other, AW does not favor just one of them, but predicts an attractivity-weighted average of the correctly and incorrectly predicting adversaries, and so doesn't fall into the same trap as ITB. Figure 3 in the "Imitate-the-Best (ITB)" section illustrates this fact: In contrast to εITB, AW is not deceived but approximates the maximal success of the deceivers.

The important mathematical result concerning AW is stated in Theorem 4. It does not hold for all monotonic but only for all *convex* loss functions. A loss function is convex if the loss for a weighted average of two predictions is less than or equal to the weighted average of the losses for the two predictions. Convex loss functions cover a wide class, including the natural linear loss function $|pred_n - e_n|$ as well as loss functions which are polynomial or exponential in $|pred_n - e_n|$.

Theorem 4 *For every real-valued prediction game ((e), {P_1, …, P_m, AW}) with a convex loss-function, the following holds:*

4.1. *(Short run:)* $(\forall n \geq 1:)\ suc_n(AW) \geq maxsuc_n - \sqrt{m/n}$.
4.2. *(Long-run:)* $suc_n(AW)$ *approximates the non-MI-players' maximal success for* $n \to \infty$.

Theorem 4 tells us that AW is indeed a universally access-optimal meta-inductive strategy, in the long run. A proof of Theorem 4 is found in Schurz (2008, th. 4); it refers to a central result from Cesa-Bianchi and Lugosi (2006, ch. 2.1, Theorem 2.1, Corollary 2.1).

Theorem 4 does not directly apply to binary prediction games, because AW's predictions are real values between 0 and 1, which is forbidden in binary games. Theorem 4 can be transferred to the prediction of binary events by interpreting AW's real-valued prediction $r \in [0,1]$ as AW's probability of predicting 1.[10]

In prediction games with intermittent forecasters, the performance of AW can be improved by the intermittent version of AW, iAW. At each time, iAW ignores all predictors who didn't deliver a prediction. The 'intermittent attractivities' of the predicting players ("i-at(P)") are defined with the help of their validities, as follows:

$$\text{i-at}_n(P) =_{def} val_n(P) - suc_n(iAW) \text{ if this difference is positive; else } \text{i-at}_n(P) = 0.$$

Provided that these validities are better than the success rates of random guesses, one can prove that the success rate of iAW exceeds that of AW, analogously to

[9] This requirement guarantees that under the conditions of Theorem 3 the intermittent version of AW approximates TTB in the long run.

[10] Cf. Cesa-Bianchi and Lugosi (2006, ch. 4.2). The randomization method presupposes that the event sequence does not react adversarially to AW's predictions. For adversarial event sequences, Theorem 4 can be transferred by assuming a *collective* of binary meta-inductivists who approximate real-valued predictions by the mean value of their binary predictions (cf. Schurz 2008, Theorem 5).

Theorem 3. (We dispense with stating this as a formal theorem, since the basic facts are clear.)

We now turn to the relevance of Theorem 4 for the philosophical *problem of induction*. This problem goes back to David Hume and consists in the fact that it is apparently impossible to justify the reliability of inductive inference without reasoning in a vicious circle: By arguing that "we expect induction to be successful in the future, because it was successful in the past", we presuppose what we wanted to demonstrate, namely that induction is successful. In the area of computational learning theory, Wolpert's *no free lunch* theorem (1996) has deepened Hume's insights.

Hume's skeptical results are compatible with the possibility of an 'a priori' justification of a meta-induction, by appeal to its access-optimality. The demonstration that a method M is optimal does not entail that M is successful, since M may merely be the 'best of a bad lot'. However, if one can demonstrate by a priori arguments that meta-induction is universally access-optimal, then one can use this to derive a non-circular a posteriori justification of ordinary object-level induction, as follows: So far induction was most successful among all accessible prediction methods, therefore by meta-induction (for which we have an independent justification) it is rational to apply inductive methods in the future (cf. Schurz 2008, 2009).

Theorem 4 establishes an a priori justification of attractivity-based meta-induction: In all environments, it is always reasonable (in *addition* to searching for good candidate methods) to apply AW/iAW, as this can only improve but not worsen one's success in the long run. Given Theorem 4, we can not only infer that AW is universally access-optimal, but even more, that AW is *dominant* (in the long run) w.r.t. the class of all non-meta-inductive methods,[11] since for every non-meta-inductive method M' one can construct 'sufficiently normal' environments in which AW's long run success exceeds that of M', while AW is never worse than M'. Since this fact seems to contradict Wolpert's famous no free lunch theorem, an explanation is in order.

Wolpert (1996, p. 34) proved that the probabilistically expected success of any normal learning algorithm is equal to the expected success of random guessing or of any other learning algorithm, under the assumptions of a uniform prior probability distribution over all possible (completely specified) states of the world (Wolpert's "targets"). Wolpert's theorem is a far-reaching generalization of an earlier result in probability theory about the prediction of infinite binary sequences (cf. Carnap 1950, p. 565; Howson and Urbach 1996, p. 189). For this application, Wolpert's result says the following: However the learning function f (with $pred_{n+1} = f((e_1,\ldots,e_n)) \in \{0,1\}$) is defined, there are as many sequences of a given length $k > n$ that verify f's prediction as there are sequences of length k that falsify it. So by attaching an equal probability to every possible sequence the expected score of each learning algorithm

[11] This dominance-claim can be strengthened (cf. Schurz 2008, Sect. 9.2). But AW is not universally access-dominant, since there are variations of AW with a different short-run performance.

will be 1/2. This result applies equally to all meta-level selection strategies given that they are applied to a finite toolbox of normal prediction methods.

Wolpert's theorems make certain assumptions which do not hold for all prediction games, e.g., the homogeneity of the loss function (ibid., p. 1377), which holds for binary but not for real-valued events. But let us grant these assumptions, and turn to a more fundamental problem behind Wolpert's result: Wolpert's assumption that each particular world (binary sequence) has the same probability density is itself strongly biased. For infinite sequences, this assumption entails that one is subjectively certain (i.e., believes with probability 1) that (a) the sequence is non-computable, and (b) has a limiting frequency of 1/2.[12] However, the sequences for which a prediction algorithm can be better than a random guesser are precisely those that don't fall into the classes (a) or (b). In other words, a uniform prior distribution over all possible sequences entails that one is subjectively certain that our world is completely irregular, so that induction has no chance.

It is well-known that if one assumes a prior distribution that is not uniform over all possible sequences, but rather over all possible frequency limits of sequences, then one validates the famous Laplacean rule of induction, $p(e_{n+1}=1 \mid f_n(1) = k) = (k + 1)/(n + 2)$, where "$f_n(1)$" denotes the number of 1's among the first n events (cf. Howson and Urbach 1996, p. 55ff). According to this prior distribution, the optimal binary prediction rule predicts $e_{n+1} = 1$ iff $f_n(1) > 1/2$ (cf. Schurz 2008, §3). In computer science, the idea underlying the Laplacean induction rule has been impressively generalized by Solomonoff (1964, §4.1), who proves (among other things) that if the prior probability of a sequence is inversely proportional to its algorithmic complexity, then one validates Laplace's rule of induction.

In our eyes, Wolpert's and Solomonoff's results confirm the fact that Bayesian results are always dependent on assumed prior distributions that are never 'unbiased' or 'information-less'. This fact underscores the significance of the non-probabilistic optimality results in dynamic online learning. Still there seems to be a contradiction between Wolpert's no free lunch theorem and the explained dominance-interpretation of our Theorem 4, which asserts that the predictive long-run success of AW is in no environment worse, but in 'regular' environments better than the long-run success of any accessible non-inductive method. This contradiction disappears if one recalls that the class of regular sequences has probability zero according to Wolpert's prior distribution and, thus, doesn't count in the expectation value. However, since evolution would be impossible in worlds without regularities, we have every reason to take these *probability-zero* worlds seriously. This is what meta-induction does: We cannot prove that meta-induction will be successful, but we can prove that if anything will be successful, meta-induction will.

[12] (a) follows from the fact that there are uncountably many sequences but only countably many computable ones. (b) holds since the uniform prior distribution over $\{0,1\}^{\infty}$ implies $p(e_i|e_j) = p(e_i) = 1/2$, i.e., the distribution is IID, which entails (b) by the strong law of large numbers.

Local Improvements Over Attractivity-Based Meta-Induction at the Cost of Universal Access-Optimality

Are there any meta-strategies whose performance exceeds AW's? This is an open question. So far our findings support the following conclusions:

1. If these strategies are *inaccessible* to AW, and select from the same toolbox as AW, the answer is: *Yes*, but their surplus success over AW is local, restricted to specific types of environments, and comes at the cost of forfeiting universal access-optimality.
2. If these strategies are themselves accessible to AW (i.e., belong to AW's toolbox), the answer is "*No*" in the long run (by Theorem 4), although their short-run performance may exceed that of AW locally, at the cost of losing universal access-optimality.

The apparent fact that improvements over AW come at the cost of losing universal access-optimality is another 'revenge effect' of ecological rationality, which we will illustrate using two examples: (1) success-weighted prediction methods, and (2) improvements in short-run performance.

A class of weighting methods that has been frequently studied in the research on adaptive rationality is *success-based* weighting. A particularly simple success-based weighting method is "Franklin's rule", here abbreviated as SW, which identifies the players' weights with their normalized success-rates:

$$\mathrm{pred}_{n+1}(\mathrm{SW}) = \frac{\sum_{1 \le i \le m} \mathrm{suc}_n(\mathrm{P_i}) \cdot \mathrm{pred}_{n+1}(\mathrm{P_i})}{\sum_{1 \le i \le m} \mathrm{suc}_n(\mathrm{P_i})}$$

The corresponding intermittent variant of SW, abbreviated "iSW", identifies the players' weights with their normalized validities, and ignores (in the above sum) those players who didn't deliver a prediction in the respective round n.

While SW is better than AW and ITB in some environments,[13] it is not universally access-optimal: In some environments, SW's success-rate may drop far below the maximal success rates of the object-level players or cues, and the same holds for iSW (see next section). For real-valued prediction games, this can be seen from the following example: Assume two (linearly scored) forecasters P_1 and P_2 who invariably underestimate the value of the predicted event, where P_2's success rate is permanently greater than that of P_1 which is in turn greater than zero. Then it is easy to prove that SW's success rate will be permanently smaller than that of P_2, because P_1's weight will never become zero.

Attractivity-based weighting methods are protected against the preceding sort of suboptimality, because they attach a weight of *zero* to all players whose success rate is (significantly) smaller than their own. This measure ensures that AW's success

[13] See next section. Katsikopoulos and Martignon (2006) proved that under the condition of known validities and "naive Bayes environments" (conditionally independent cue validities and uniform prior), the logarithmic version of iSW that takes $\log(\mathrm{val}(P_i)/(1 - \mathrm{val}(P_i)))$ as the weight of cue P_i is probabilistically optimal among all possible methods.

rate (after some convergence time) grows above that of the suboptimal player P_1. From that time point on, AW imitates solely the optimal player P_2, approximates P_2's success, and is thus guaranteed to be universally access-optimal.

A second area in which the *revenge of ecological rationality* manifests itself is *short run* performance. The worst case short run loss of AW is $\sqrt{m/n}$, which is only small if the number of competing methods m is small compared to the number of rounds n of the prediction game—the so-called "prediction horizon". The worst-case bound for the short-run loss of AW can be improved, but only if one knows the prediction horizon in advance, which is a restrictive condition on prediction games. Under this condition it is optimal to use exponential attractivity-weights (cf. Cesa-Bianchi and Lugosi 2006, p. 16f).

AW with known prediction horizon h and exponential attractivity weights (EAW):

Weight of player P at times $n \leq h$: $w_n(P) =_{def} e^{\sqrt{8 \cdot \ln(m/h)}} \cdot at_n(P)$
Worst-case loss of EAW: $\sqrt{\ln(m)/(2 \cdot n)}$

The short run loss can be further reduced if the environment satisfies further conditions. An example is the *conditionalization* of the success rates of the players or cues to certain properties which reliably indicate a significant change of the environment. The study of these possibilities is left to another paper.

In prediction games which have a unique best player after a certain winning time (Theorem 1), AW's predictions and success rates converge against the predictions and success rates of ITB, because as soon as AW becomes better than some player P, AW ignores P in the weighted average, until AW eventually assigns a weight of 1 to the best player. (Similarly, iAW approximates TTB in intermittent games.) However, AW's convergence to ITB takes some time. In the short run, ITB enjoys an advantage over AW (and TTB over iAW) in those prediction games in which the cue validities converge quickly to a unique ordering. So again AW has to pay a price for its long-run optimality.

The 'revenge of ecological rationality' puts us in a certain dilemma: On the one hand, we have a meta-level strategy AW which is universally access-optimal in the long run. On the other hand, we have methods whose performance may exceed that of AW in the short run—and if they are inaccessible to AW even in the long run—but at the cost of losing universal access-optimality. We propose to solve this dilemma by the following *division of labor*: If the performance of a prediction strategy exceeds that of AW in some local environments, then we should not use this strategy as our meta-level selection strategy, but we should put it into the toolbox of locally adapted candidate methods. As one's meta-level selection strategy, one should employ a strategy which is known to be access-optimal. In situations of dynamic online learning this is AW (or iAW). If we find a meta-strategy S* (e.g., SW) which is more successful than AW in some environments, then we should not try to improve our success by replacing AW by S* at the meta-level, but by putting S* into the toolbox of candidate methods and applying AW to this *extended* toolbox.

Our proposed division of labor between general selection strategies and locally adapted methods solves the explained dilemma, at least in the long run: Since AW imitates S* in those and only those environments in which S*'s success is greater than that of the other players, AW enjoys S*'s local long-run advantage without suffering from its drawbacks in other environments. In the short run, however, AW may still suffer a certain delay and thus a loss when imitating S* in those environments to which S* is locally adapted. A complete avoidance of the revenge of ecological rationality in the short run is, thus, impossible.

Meta-Induction Within the Monash University Footy Tipping Competition: Results of an Empirical Study

In the final part of this paper, we describe a retrospective empirical application of the discussed meta-inductive strategies to the results of one of the world's longest-running prediction competitions, the Monash University footy tipping competition (MUFTC). The event sequence that we consider consists of the 3-valued results of 1514 matches of the Australian Football League (1, 0, or tie) over 8 seasons from 2005 to 2012, recorded at the MUFTC website.[14] The prediction of each match constitutes one round of the prediction game. Our tournament included the predictions of all 1071 human participants, as well as the predictions of different meta-inductive strategies that were applied to these predictors (while the meta-inductive strategies were mutually inaccessible). Predictors had to specify the winning probability of the first of two teams. We scored their predictions by a linear loss function.

The MUFTC tournament presented a *significant challenge* for meta-inductive strategies for the following reasons:

1. Individual participation was highly intermittent: Only 69 (31) out of the 1071 players made predictions for at least 1/2 (3/4) of the matches. Many players entered the game in the midst of the season for several rounds and distracted the meta-inductive algorithms by lucky initial successes. For this reason, we considered "subgames" of the competition where the player set was reduced, for example to the subset of 69 players who made predictions in at least 1/2 of the rounds.
2. The success rates of the players were quite close together (0.49–0.62). Their validities differed more strongly, but because of many players predicting only a few times, the validity dispersion was not conclusive.
3. Even for players who predicted very frequently the success rates were not much higher than the success of random guessing. So there were no strong effects that the meta-inductivist strategies could exploit.
4. The number of competing players (1071) was not much smaller than the number of rounds (1514), which means that from the perspective of AW, the tournament was not a long run but only a *short run* experiment.

[14] http://www.csse.monash.edu.au/~footy/.

Table 2 Theoretical worst-case loss of AW and EAW and empirical loss of AW compared to the maximal success rate, in the subgame with 31 players who predicted at least 75 % of the time

Round	AW: $\sqrt{m/n}$	EAW: $\sqrt{\ln(m)/(2 \cdot n)}$	AW: obtained
20 "extremely short"	1	0.29	0.025
100 "very short"	0.56	0.13	0.026
500 "short"	0.25	0.06	0.006
1500 "medium"	0.14	0.034	0.005
10,000 "long"	0.05	0.01	(expected: 0)

Nevertheless we obtained many interesting and surprising results.

Result 1 The actually obtained short run losses were much smaller than the theoretically calculated worst-case short run losses. This result is shown in Table 2.

While in 6 out of the 8 seasons there was a *different* best player, AW and ITB was among the top players, in each season. Moreover, in all considered subgames, the success rate of AW and ITB was close to the maximal success rate of the human participants. The same was not true for SW and CSW. This fact reflects the access-optimality of AW and ITB in the considered subgames.

Result 2 In spite of the low dispersion of success rates, ITB was almost always better than the success-based weighting method SW (recall fn. 2). ITB was also better than the method CSW, which stands for "chance-adjusted success-weighting". This is a variant of SW proposed in Jekel et al. (2012), which identifies the weight of a predictor with its success rate minus the success rate of random guessing.

Figure 5 illustrates this effect, charting the performance of the four meta-inductive methods SW, CSW, AW and ITB, which employ a persistent success evaluation, amidst the 69 players who predicted at least 1/2 of the time. We see that SW and CSW are clearly below AW and ITB, while CSW is slightly above SW. Moreover ITB is slightly above AW, with diminishing success differences for increasing numbers of rounds. At the end of the game, ITB and AW come close to the maximal success rate of the human players, represented by the solid black line: Instances of the triangular icon indicate a switch of the top human player. The dotted black line represents the success rate of that human player who had the highest success at the end of the game (and is occluded by the black line, once that player achieves the maximal success rate).

Result 3 (*Intermittent success evaluation*) Human predictors with the highest success rates also tended to make the most predictions. Conversely, predictors whose validities differed significantly from their success rates predicted rarely. So in most rounds the players with the highest success rate were identical with the players with the highest validity. Therefore there were not many effects to be exploited by the intermittent versions of the meta-inductive algorithms (TTB, iAW, iSW, and iCSW) in prediction games with the 69 *most frequent* human predictors.

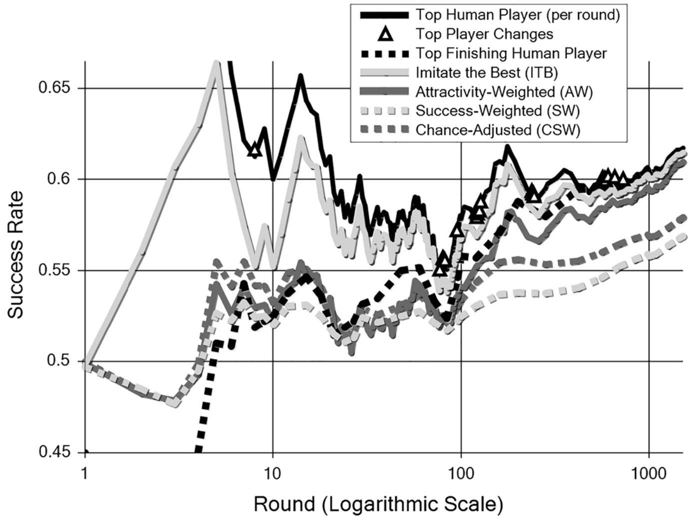

Fig. 5 Results of the footy tipping competition, subgame with 69 players who predicted at least 50 % of the time

To study the impact of intermittent success evaluation, we investigated a subgame consisting of the 50 human players with the highest validities. The validities of these players differed significantly from their success rates, because they predicted with low frequency. In this subgame, we expected the intermittent versions of the meta-inductive strategies to demonstrate their advantage over persistent strategies. This expectation was confirmed. Figure 6a shows the result of the described subgame, with four interesting sub-results:

1. The intermittent meta-inductive strategies performed *much* better than their persistent versions, which for the sake of comparison are shown in Fig. 6b.
2. The success rates of the intermittent meta-inductive strategies climbed above the maximal success rate of the human players (as explained in the "Take-the-Best (TTB)" section).
3. In the beginning of the game, the intermittent weighting methods were better than TTB. Here TTB suffered due to the ongoing appearance of new players: When a new player P's first prediction had a high score (by luck), but his second prediction scored low, TTB imitated P's second prediction and earned a loss. We experience here a kind of the TTB-adversarial scenario which was theoretically described in the "Imitate-the-Best (ITB)" and "Take-the-Best (TTB)" sections (Theorem 2).
4. iSW had an even higher success rate than iAW in this game. This demonstrates, empirically, the effect which was theoretically described in the previous section, that in some environments SW's performance may be greater than that of AW. Of equal interest is the fact that, in this game, the chance-adjusted method iCSW performed slightly worse than iSW, which refutes the conjecture

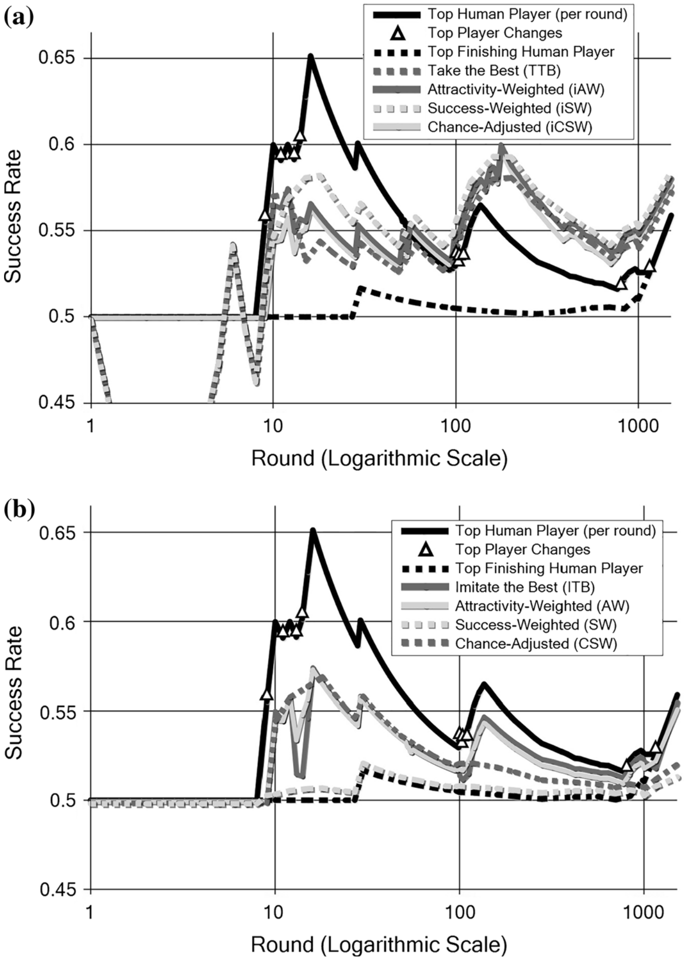

Fig. 6 Footy tipping subgame with the 50 human predictors with the highest validities. **a** Intermittent success evaluation, **b** persistent success evaluation

in Jekel et al. (2012, p. 12), according to which iCSW is expected to be generally better than iSW.

For the sake of comparison with the usual setting of prediction studies in adaptive rationality research (cf. fn 3), we ran two variants of our data set.

In *Footy Variant 1*, we re-ran the footy tipping prediction games under the condition that the success rates and validities were estimated by random sampling from future events. Here we found that both ITB and AW performed better in the case of continuous success updating via online learning, as compared to success evaluation based on random sampling. The success ordering between the meta-inductive methods (ITB, AW, SW and CSW) remained the same as in Fig. 5.

In *Footy Variant 2*, we transformed the footy tipping prediction games into a binary comparative format by the following scale transformation: The cue-difference of cue (player) P_i is $+1/0/-1$ iff its real-valued (probabilistic) prediction lies in the interval [0.55,1]/(0.45,0.55)/[0,0.45], respectively. Our major result was that a scale transformation from the real-valued into the binary-comparative format worsened the success of TTB compared to weighting methods. This result agrees with the finding reported in Katsikopoulos et al. (2010, p. 1265).

Detailed reports of Footy variants 1 and 2 are left to future studies.

Summary and Conclusions

The study of prediction games under conditions of dynamic online learning raises new questions about adaptive rationality. We discussed competing research programs concerning the question of the local adaptivity versus the generality of prediction strategies, and argued for a *division of cognitive labor* between locally adapted methods and *meta-inductive* selection strategies. We presented mathematical theorems concerning the optimality of different meta-inductive strategies in regard to their *accessible* object-level methods, a property that we called *access-optimality*. We showed that one-favorite methods such as ITB and TTB are access-optimal in environments with stabilizing success-orderings, but in environments with adversarial success-oscillations (keyword “bubble economy”) their success breaks down. We then showed that there is a meta-inductive strategy, attractivity-based meta-induction AW, which is universally access-optimal in the long run. The performance of AW can be improved upon, but only locally, and at the cost of forfeiting universal access-optimality. We called this situation the “revenge of ecological rationality”.

In the final section, we presented an empirical study based on data from the Monash University footy tipping tournament. The results of our study confirmed our theoretical results and expectations. We observed that: (1) AW and ITB (but not SW or CSW) were approximately access-optimal in regard to the human predictors in all subgames of the tournament, (2) in most subgames, ITB and AW were highly successful, (3) the performance of SW and CSW exceeded that of ITB and AW in only a few subgames, while in most subgames SW and CSW were inferior to ITB and AW, and (4) in subgames with rarely predicting players, the intermittent versions of meta-inductive strategies exhibited significantly improved performance, compared to their persistent cousins.

Acknowledgments Work on this paper was supported by the DFG Grant SCHU1566/9-1 as part of the priority program “New Frameworks of Rationality” (SPP 1516). For valuable help we are indebted to K.V. Katsikopolous, Ö. Simsek, A.P. Pedersen, R. Hertwig, M. Jekel, P. Grunwald, J.-W. Romeijn, and L. Martignon.

Appendix

Proof of Theorem 2 *For 2.1(a):* Since xMI's imitates for each time $n > 1$ some deceiver P_i, xMI's score for all times > 1 is 0, and so limsuc(xMI) = 0. *For 3.1(b):* On average, xMI imitates each player equally often, with a limiting frequency of 1/m. So the frequency of times for which each player earns a maximal score of 1 is $(m - 1)/m$. *For (2.2):* Let "p(fav-P_i)" be the limiting frequency of times for which player P_i was xMI's favorite, and limsuc(P_i|fav) be player P_i's limit success conditional on these times. Then by probability theory, limsuc(xMI) = $\Sigma_{1 \leq i \leq m}$p(fav-P_i)·limsuc(P_i|fav), which implies that limsuc(xMI) $\leq$ max({limsuc(P_i|fav):$1 \leq i \leq m$}). By the negative correlation assumption, limsuc(P_i|fav)) < limsuc(P_i) holds for all $i \in \{1, \ldots, m\}$; so limsuc(xMI) < max({limsuc(P_i):$1 \leq i \leq m$}) $=_{def}$ maxlimsuc. Hence xMI is not long-run optimal. □

Proof of Theorem 3 *For 3.1(a):* Before the convergence time s, TTB may be, in the worst case, permanently deceived by the non-MI-players (or cues), by negatively correlated success-oscillations. So TTB's worst-case success until time s is zero, whence his worst-case loss at times $n \geq s$ is s/n. After time point s, TTB's earns for each kth-best player (or cue) P_{s_k} the sum-of-scores earned by P_{s_k}, for all time points at which P_{s_k} but no player better than P_{s_k} delivered a prediction. The sum-expression in 3.1(a) is identical with the sum of these scores divided by time n. *For 3.1(b):* Additionally we assume that after time s each player's validity is better than the success of a random guess, which is 1/2 in a binary prediction game. This implies that suc_n(TTB) > $maxsuc_n - (s/n)$, where $maxsuc_n \geq suc_n$(ITB) since ITB approximates $maxsuc_n$ from below (by Theorem 1). *(3.2)* follows from (3.1) in the explained way, since $\lim_{n\to\infty}(s/n) = 0$. □

References

Arnold, E. (2010). Can the best-alternative-justification solve Hume's problem? *Philosophy of Science, 77*, 584–593.

Brighton, H. & Gigerenzer, G. (2012). How heuristics handle uncertainty. In Todd & Gigerenzer (eds.), (pp. 33–60). (Note: This article is an abbreviated version of Gigerenzer and Brighton 2009.)

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

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

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In: Gigerenzer et al. (Eds.), (pp. 97–118).

De Finetti, B. (1974). *Theory of probability*. New York: John Wiley.

Dieckmann, A., & Todd, P. M. (2012). Simple rules for ordering cues in one-reason decision making. In: Todd & Gigerenzer (Eds.), (pp. 274–306).

Feldbacher, C. (2012). Meta-induction and the wisdom of crowds. *Analyse & Kritik, 34*(2), 367–382.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science, 1*, 107–143.

Gigerenzer, G., Todd, P. M., & The ABC Research Group (Eds.). (1999). *Simple heuristics that make us smart*. Oxford: Oxford Univ. Press.

Hertwig, R., Hoffrage, U., & The ABC Research Group. (2013). *Simple heuristics in a social world*. New York: Oxford University Press.

Hoffrage, U., Hertwig, R., & Gigerenzer, G. (2000). Hindsight bias: A by-product of knowledge updating? *Journal of Experimental Psychology. Learning, Memory, and Cognition, 26*, 566–581.
Hogarth, R. M., & Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is Less "More"? *Journal of Mathematical Psychology, 49*, 115–124.
Howson, C., & Urbach, P. (1996). *Scientific reasoning: The Bayesian approach* (2nd ed.). Chicago: Open Court.
Jekel, M., Glöckner, A., Fielder, S., & Bröder, A. (2012). The rationality of different kinds of intuitive decision processes. *Synthese, 89*(1), 147–160.
Katsikopoulos, K. V., & Martignon, L. (2006). Naive heuristics for paired comparisons. *Journal of Mathematical Psychology, 50*, 488–494.
Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review, 117*(4), 1259–1266.
Kelly, K. (1996). *The logic of reliable inquiry*. New York: Oxford Univ. Press.
Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? In: Gigerenzer et al. (Eds.), (pp. 119–140).
Popper, K. (1935/2002). *Logic of discovery*. London: Routledge.
Rieskamp, J., & Dieckmann, A. (2012). Redundancy: Environment structure that simple heuristics can exploit. In Todd & Gigerenzer (eds.), (pp. 187–215).
Rieskamp, J., & Otto, P. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135*(2), 207–236.
Schmitt, M., & Martignon, L. (2006). On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research, 7*, 55–83.
Schurz, G. (2008). The meta-inductivist's winning strategy in the prediction game: A new approach to Hume's problem. *Philosophy of Science, 75*, 278–305.
Schurz, G. (2009). Meta-induction and social epistemology. *Episteme, 6*, 200–220.
Schurz, G. (2012). Meta-induction in epistemic networks and social spread of knowledge. *Episteme, 9*(2), 151–170.
Schurz, G. (2013). *Philosophy of science: A unified approach*. New York: Routledge.
Schurz, G., & Thorn, P. (2014). TTB vs. Franklin's rule in environments of different redundancy. Appendix to G. Schurz, Instrumental justifications of normative systems of reasoning. *Frontiers in Psychology*, *5*, article 625, July 2014. doi:10.3389/fpsyg.2014.00625.
Shanks, D. R., Douglas, L. M., & Holyoak, K. J. (Eds.). (1996). *The psychology of learning and motivation*. San Diego, CA: Academic Press.
Simon, H. A. (1982). *Models of bounded rationality* (Vol. 1). Cambridge, MA: MIT Press.
Solomonoff, R.J. (1964). A formal theory of inductive inference. *Information and Control, 7*, 1–22 (part I), 224–254 (part II).
Thorn, P., & Schurz, G. (2012). Meta-induction and the wisdom of crowds. *Analyse & Kritik, 34*(2), 339–366.
Todd, P. M., & Gigerenzer, G. (Eds.). (2012). *Ecological rationality: intelligence in the world*. New York: Oxford University Press.
Vickers, J. (2010). The problem of induction. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*. http://plato.stanford.edu/archives/spr2010/entries/induction-problem.
Weibull, J. (1995). *Evolutionary game theory*. Cambridge, MA: MIT Press.
Wells, A. J. (2005). *Rethinking cognitive computation: Turing and the science of the mind*. Basingstoke: Palgrave.
Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation, 8*(7), 1341–1390.