Optimisation in a Synchronised Prediction Setting[*]

Christian J. Feldbacher-Escamilla

Autumn 2017

Abstract

[419] The standard approach to solve prediction tasks is to apply inductive methods such as, e.g., the straight rule. Such methods are proven to be access-optimal in specific prediction settings, but not in all. Within the optimality-approach of meta-induction, success-based weighted prediction methods are proven to be access-optimal in all possible continuous prediction settings. However, meta-induction fails to be access-optimal in so-called *demonic* discrete prediction environments where the predicted value is inversely correlated with the true outcome.

In this paper the problem of discrete prediction environments is addressed by embedding them into a synchronised prediction setting. In such a setting the task consists in providing a discrete and a continuous prediction. It is shown that synchronisation constraints exclude the possibility of demonic environments.

Keywords: *meta-induction*, *binary prediction games*, *qualitative belief*, *degrees of belief*, *Lockean thesis*, *calibration*

1 Introduction

Hume's problem of induction poses a serious problem for the scientific task of making predictions. On the one hand, predictions about unobserved data or events are made by help of knowledge about already observed data or events. On the other hand, there is no strict logical relation between observed and unobserved data or events that would allow for a direct epistemic justification for such an inductive transmission.

[420] Among the several proposals that were put forward to address this problem is the so-called approach of *meta-induction* (cf. Schurz 2008). This

[*][This text is published under the following bibliographical data: Feldbacher-Escamilla, Christian J. (2017). "Optimisation in a Synchronised Prediction Setting". In: *Journal for General Philosophy of Science* 48.3, pp. 419–437. DOI: 10.1007/s10838-017-9379-7. All page numbers of the published text are in square brackets. The final publication is available at https://link.springer.com/. For more information about the underlying project, please have a look at http://cjf.escamilla.academia.name.]

approach justifies inductive methods via their ability to catch up with any prediction method whatsoever in the long run. To be a little bit more precise, the meta-inductive justification of induction is twofold: First, a specific meta-method for selecting predictions of any accessible method is proven to be access-optimal in the long run regarding predictive success. A meta-method is called 'access-optimal', if its predictive success is at least as high as that one of the best accessible object-methods in the setting is. In a second step, taken for granted the past success of classical inductive methods—something that is, e.g., also not scrutinised by Hume himself (cf. Howson 2003, p.4)—it is shown that also selecting these methods for predictions of unobserved data and events allows for access-optimal predictive success. This holds as long as there is no alternative method accessible which outperforms classical inductive methods.

There are several restrictions of the justification provided by meta-induction. So, e.g., in contrast to the justification of deductive methods, meta-inductive justification is not an absolute, but a relative one: It does not show that induction will be predictively successful, but only that it is predictively access-optimal; this means that there is no guarantee of a high number of correct or accurate predictions; it is only guaranteed that the number of correct or the accuracy of the predictions is at least as high as that of the best *accessible* object-methods. Furthermore, strict access-optimality results hold only for the long run, i.e. for an infinite series of predictions.

Regarding the meta-inductive results on access-optimality there is another restriction we will focus on in this paper. It is the problem of discrete prediction settings, i.e. series of predictions about events whose outcomes are not continuous or real-valued, but discrete (in the most specific case binary). The problem is that strict access-optimality does not hold for such settings. And if strict access-optimality does not hold, then the meta-inductive justification for selecting inductive methods on the basis of their past success fails.

In order to address this problem, we are going to discuss proposed solutions to it and highlight their advantages as well as their shortcomings. Then we are going on to embed the problem in a broader prediction setting where it is supposed that the methods are providing not only discrete predictions, but also continuous ones. We will argue for the thesis that rules for bridging discrete or qualitative and continuous or quantitative predictions prevents meta-inductive suboptimality in the discrete setting.

For this purpose we will start with a short sketch of the framework that is used for making predictions, namely the framework of meta-inductive prediction games (section 2). Afterwards the set of such prediction environments is restricted to discrete prediction settings. The above mentioned problem for such settings is posed and two solutions that are found in the literature are discussed (section 3). Finally, we will expand the discussion by an investigation of the problem in a synchronised continuous and discrete prediction setting (section 4).

2 Meta-Inductive Predictions

The theory of meta-induction generalises Hans Reichenbach's *best alternatives approach* (cf. Schurz 2008, sect.2). Reichenbach proposed to consider the problem of induction not with respect to the strong requirement of [421] proving that inductive methods are successful, but with respect to the much weaker, but epistemically still highly relevant, requirement of proving that inductive methods are the best methods accessible for making predictions. In his very influential "*Experience and Prediction. An Analysis of the Foundations and the Structure of Knowledge*" (1938) Reichenbach argues as follows:

1. "If we cannot realize the sufficient conditions of success, we shall at least realize the necessary conditions. If we were able to show that the inductive inference is a necessary condition of success, it would be justified; such a proof would satisfy any demands which may be raised about the justification of induction." (p.348)

2. "Let us introduce the term "predictable" for a world which is sufficiently ordered to enable us to construct a series with a limit." (p.350)

3. "The principle of induction [i.e. the straight rule which transfers the observed frequency to the limit] has the quality of leading to the limit, if there is a limit [i.e. if the world is predictable]." (p.353)

4. "But is it the only principle with such a property? There might be other methods which also would indicate to us the value of the limit. [...] Imagine a clairvoyant who is able to foretell the value p of the limit in such an early stage of the series [where the straight rule still fails];" (p.353)

5. "The indications of the clairvoyant can differ, if they are true, only in the beginning of the series, from those given by the inductive principle. In the end there must be an asymptotical convergence between the indications of the clairvoyant and those of the inductive principle." (p.354)

6. "If there is any method which leads to the limit of the frequency, the inductive principle will do the same;" (p.355)

7. [Hence, asymptotical convergence or long run equality with the inductive principle is a necessary condition for success in predictible worlds.]

8. [Hence, the inductive principle is justified regarding predictible worlds.]

Now, Reichenbach's solution to the problem of induction is a very simple, but also narrow one: If the world is predictable in the sense that for any distribution under investigation there is a limiting frequency, then a method that is defined as approaching this frequency in the limit (as, e.g., is guaranteed by the straight rule (cf. Howson 2003, p.72)), will "lead to the limit". It is clear that the whole analytical argument is based on the specific interpretation of 'a series is predictable' as 'there exists a limit of the series' (cf. premise 2) and that by this the meaning of 'induction' is some kind of "smuggled into" the meaning of 'prediction'.

However, one can try to weaken the assumptions made by Reichenbach and still prove that following an inductive method is still a "necessary condition" for predictive success in the sense that all other accessible methods that are most successful converge with that inductive method. Exactly this is done within the approach of meta-induction (cf. Schurz 2008): Here Hume's problem is framed as the problem of providing a successful prediction of the outcome of an event E_{t+1} based on information about the outcome of events $E_1, E_2, \dots, E_t$. Similar to Reichenbach's proposal, 'predictive success' is not defined as a 'correct or true prediction', but as 'best prediction accessible among all alternatives'. Contrary to Reichenbach's proposal, there are no constraints whatsoever on the series of events $E_1, E_2, \dots$; there might be a limiting frequency of the distribution of properties within this series or not—it might be predictable in the sense of premise 2 of Reichenbach's argument or completely random. Also different from Reichenbach's framing of the problem, within the approach of meta-induction it is argued for the predictive success of induction on a meta-level instead of an object-level: Whereas the (inductive) straight rule considered by Reichenbach is applied to the outcomes of the series of events $E_1, \dots, E_t$ in [422] order to predict the outcome of event E_{t+1}, the meta-inductive method is applied to the outcomes of this series taken together with the predictions made by all alternative methods—this is the reason why it is called a 'meta-method'. The underlying idea of applying the meta-inductive method is that of selecting among predictions those whose predicting methods (e.g., the straight rule) were most successful in past. It can be proven that such a selection procedure is most successful in the long run, i.e. its predictive success converges with that of the best prediction methods considered in the selection procedure. In this way one can say that the meta-inductive method infers from past success future success; it is successful induction over success rates. But here some more details: The framework of meta-induction are so-called *prediction games* where one has the following ingredients (cf. Schurz and Thorn 2016, sect.3, notation adjusted):

- $E_1, E_2, \dots$ is an infinite series of events whose outcomes $e_1, e_2, \dots$ are within $[0, 1]$
- $Pr_1(E_t), \dots, Pr_n(E_t)$ are the predictions of E_t (also elements of $[0, 1]$) of n so-called *object-level agents*
- $Pr_{mi}(E_t)$ is the prediction of E_t of a so-called *meta-level agent*

That the outcomes of the events are restricted to the interval $[0, 1]$ is just an idealisation which can be easily met by normalising outcomes: If, e.g., $E_1, E_2, \dots$ is a series of weather-events with temperature values as outcomes, one just has to normalise via the highest possible temperature value in order to embed the predictions into this setting.

As we have said above, a meta-inductive method selects or "cooks up" a prediction via past success rates; in order to keep track of success one might first measure the score of a prediction about event E_t by method Pr_i via measuring the inverse of the error of the prediction:

$$score(Pr_i, E_t) = 1 - |e_t - Pr_i(E_t)|$$

And then one just sums up all the scores of Pr_i's predictions on $E_1, E_2, \dots, E_t$ and averages over them:

$$success(Pr_i, E_t) = \frac{\sum_{s=1}^{t} score(Pr_i, E_s)}{t}$$

Now, there are several ways to try to infer via past success a prediction which guarantees relative future success. One way, e.g., is to simply copy the most successful method's prediction (in case there are several equal successful ones, one might just randomly select among them). This so-called *simple meta-inductivist* predicts most successful in case there is a single expert-setting, i.e. there is one method that turns out to be the most successful and there are not at least two prediction methods that are equally successful, but oscillating more or less close to the truth (cf. Schurz 2008, sect.4). However, a method that is guaranteed to be access-optimal in all situations is the so-called *attractivity weighted meta-inductivist* which assigns a weight to each method whose performance was in the past at least as successful as itself. Note that a method is *access*-optimal if it "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" Access-optimality has to be distinguished from absolute optimality, which means that a method predicts optimally in comparison to all candidate prediction methods, regardless of whether they are accessible to it or not (cf. Schurz and Thorn 2016, end of sect.2). For the attractivity weighted meta-inductivist it holds that the higher the past success of an attractive method, the higher [423] is also its weight. In detail the weight of a method Pr_i for method Pr_j regarding event E_{t+1} is defined as follows (cf. Schurz and Thorn 2016, sect.6):

$$weight(Pr_i, Pr_j, E_{t+1}) = \frac{max(0, success(Pr_i, E_t) - success(Pr_j, E_t))}{\sum_{k=1}^{n} max(0, success(Pr_k, E_t) - success(Pr_j, E_t))}$$

Note that the weights are success-based only and that methods that are worse performing than Pr_j get weight 0. Based on these weights, the weighted-average meta-inductive method predicts by weighting the relatively successful methods' predictions as follows (cf. Schurz 2008, sect.7):

$$Pr_{mi}(E_{t+1}) = \sum_{k=1}^{n} weight(Pr_k, Pr_{mi}, E_{t+1}) \cdot Pr_k(E_{t+1})$$

In case there are no better performing methods, but also at the very beginning (E_1) the meta-inductivist's prediction consists in the (unweighted) average of all predictions.

As we have explained above, the idea behind the meta-inductive method Pr_{mi} is quite simple; interestingly it proves to be very powerful regarding the task of justifying induction in the sense proposed by Reichenbach: It can be shown that there are quite narrow bounds of Pr_{mi}'s predictive success regarding the best predicting methods. By transforming theorems of the machine learning literature it is proven in (Schurz 2008, sect.7), that the upper and lower bounds of Pr_{mi}'s success are as follows:

$$\begin{aligned} success(Pr_{mi}, E_t) &\leq 1 \\ success(Pr_{mi}, E_t) &\geq success(Pr_i, E_t) - \sqrt{\frac{n}{t}} \quad \text{for all } 1 \leq i \leq n \end{aligned}$$

Interesting is the lower bound that states that Pr_{mi}'s success will differ by at most $\sqrt{n/t}$ (n being the number of methods under consideration; t being the number of predictions already performed) from every single method's success rate and thereby also from the best performing methods' success rates. A consequence of this result is the following theorem on the *long run acccess-optimality* of meta-induction:

$$\lim_{t\to\infty} success(Pr_{mi}, E_t) - max(success(Pr_1, E_t), \ldots, success(Pr_n, E_t)) \geq 0$$

So, the meta-inductivist's success rate and that of the best performing methods converge in the limit. As we have seen above, that is exactly what Reichenbach has described as a necessary condition for predictive success. And what is more, this result does not depend on any constraints of the event series under investigation.

Given the access-optimality result of meta-induction, Hume's problem of induction is tackled in an a posteriori, but still non-circular way (cf. Schurz 2008, p.282): It holds analytically that meta-induction is an epistemically justified inductive method for the selection of prediction methods. Now, object-level inductive methods as, e.g., the straight rule, have proven to be some of the most successful prediction methods in past and by this are selected by the meta-inductive rule for further predictions. Hence, one is epistemically justified in using the classical object-level inductive methods for making predictions (as long as no other method outperforms them regarding predictive success).

As we mentioned already in the introduction, there are several restrictions of the justification provided by meta-induction. First, it is only an optimality-justification instead of a maximality-justification. I.e., it does not demonstrate predictive success absolutely, but [424] only with respect to accessible prediction methods. If the latter perform bad, often also the former will perform bad. However, in *epistemic engineering* it is commonly accepted to consider such justifications as legitimate (cf., e.g., Reichenbach 1938, pp.349; Schurz

2008, pp.281ff). As Hume has argued: There is no justification of induction in an absolute sense; but this does not mean that there is also no justification of induction in the sense of showing that it is the best one can do. For this reason we agree here with the suggestion to accept optimality-based justifications in epistemology. What is more, in the light of the former impossibility, the possibility highlighted by meta-induction seems to be even more valuable.

Also, one might object that the provided justification is only relative to (publicly) *accessible* methods. But, again, if one accepts optimality based justifications at all within the epistemic realm, then it is almost trivial that all considered prediction methods have to be accessible: Comparing predictive success of prediction methods presupposes the accessibility of these methods.

What one might consider more pressing is the fact that strict access-optimality holds only in the long run, i.e. for infinitely long series of predictions. However, as we have seen above, there are some clear bounds on the meta-inductivist's predictive success or failure. So, although there holds no strict access-optimality for the short run, unavoidable losses in a finite sequence of predictions can be *controlled* for (cf. Reichenbach 1938, pp.354; and Schurz 2008, p.286).

Regardless what stance one takes regarding these restrictions, there is one problem which is located at the core of the meta-inductivist's aim of justifying induction: it is the problem of dealing with discrete prediction environments.

3 Discrete Prediction Environments

The results formulated in the preceding section hold for prediction games with a continuous prediction space, i.e. both, the outcomes as well as the predictions are within the interval $[0,1]$. However, things are becoming more complicated if the prediction space is discrete. For reasons of simplicity we restrict our consideration to a binary prediction space (cf. Cesa-Bianchi and Lugosi 2006, chpt.4). So, the outcomes of events $E_1, E_2, \ldots$, i.e. $e_1, e_2, \ldots$, are within $\{0,1\}$ and also the prediction methods have to provide binary predictions. However, our argument also holds for discrete predictions with more than two basic elements in the prediction space; discretising and bridging becomes a little bit more complicated then. For simple expression we add to the above framework of prediction games the following ingredients:

- $Bel_1(E_t), \ldots, Bel_n(E_t)$ are qualitative predictions on E_t (elements of $\{0,1\}$) of the n object-level agents
- $Bel_{mi}(E_t)$ is a qualitative prediction about E_t by the meta-level agent

Now, how could we define the qualitative or discrete meta-inductive method Bel_{mi}? One could try to do this analogously to the quantitative or continuous definition above, by just weighting the qualitative predictions of $Bel_1, \ldots, Bel_n$ according to their past success rate. However, it is clear that such a weighting on it's own would not produce a discrete or binary prediction (think, e.g., on

$n = 2$ with $Bel_1(E_{t+1}) = 1, Bel_2(E_{t+1}) = 0$ and success-based weights of 0.5 for both methods which produces a weighted-average prediction of 0.5). So, in order to produce a discrete prediction, one has to cut off or round the real-valued [425] prediction to the closest discrete-value. Such a simple success-based binary prediction method would be, e.g.:

$$\begin{aligned} Bel_{mi}(E_{t+1}) &= 1 \text{ if } Pr_{mi}(E_{t+1}) > 0.5 \\ &= 0 \text{ if } Pr_{mi}(E_{t+1}) \leq 0.5 \end{aligned}$$

However, such a meta-inductive prediction method can no longer be proven to be access-optimal in the long run. The reason is that at least two best performing object-level methods might oscillate in their success rates, but due to rounding, i.e. discretising, the meta-inductive method always prefers the wrong one. Theoretically seen, the reason for this problem is that a loss-function defined via such a rounding procedure is no longer convex (cf. Cesa-Bianchi and Lugosi 2006, sect.4.1; and Schurz 2019, chpt.6.7). So, e.g., in our definition of the score above we have used an inversion of the so-called *natural loss function* (where just the absolute difference between the true outcome and the predicted outcome is taken). Now, access-optimality results are proven only for so-called *convex* loss functions (very briefly: if a function is convex in one of it's arguments, than there is no reversal of order between multiplying the argument and multiplying the value of the function by constants); the natural loss function, e.g., is a convex loss function and therefore access-optimality can be proven. But a loss function that would fit with rounding as is used in the definition of Bel_{mi} above is no longer convex and a prediction method such as Bel_{mi} can even be demonstrated to be suboptimal in some infinite series of predictions of events (cf. Schurz 2019, prop.6-10).

More generally, success-based prediction fails in all so-called *demonic scenarios*. René Descartes described such scenarios as follows:

> "Accordingly, I will suppose not a supremely good God, the source of truth, but rather an evil genius, supremely powerful and clever, who has directed his entire effort at deceiving me. I will regard the heavens, the air, the earth, colors, shapes, sounds, and all external things as nothing but the bedeviling hoaxes of my dreams, with which he lays snares for my credulity." (Descartes 1637/1998, par.22f)

In the best alternative approach the idea remains the same, although the structure is little bit more complex: Here demonic scenarios are defined as settings in which at least one method is predictively successful, i.e., the success rate is in the long run above 0, but in which the meta-level method is either not successful at all, i.e. its success rate is 0, or it is at least sub-optimal with respect to some accessible methods. Metaphorically speaking, in a demonic setting the "evil genius" is feinting the meta-level method to a higher degree than the object-level methods. We assume for a demonic setting that the success rates of the object-methods are limited. Although this is a strong assumption, it is still weaker than Reichenbach's assumption about the limit of the series of the true event outcomes. Furthermore, we could not think of a discrete setting where the meta-inductivist performs sub-optimally, although the success rates of the object-methods are not limited. And finally, this assumption allows us to generalise our approach to demonic scenarios with any finite number of object-methods. Whether there are any discrete settings where the meta-inductivist performs sub-optimally, although the object-methods' success rates are not limited, is a topic of further research. According to our extrapolation of Descartes' demon to the best alternative approach, a demonic scenario has the following properties:

1. The object-method's success rates are limited: For all $1 \leq i \leq n$ there exists $lim_{t\to\infty}(success(Pr_i, E_t))$.

2. [426] The meta-inductive method performs long run sub-optimally: There is an $1 \leq i \leq n$ such that: $lim_{t\to\infty}(success(Pr_{mi}, E_t) - success(Pr_i, E_t)) < 0$
   Note that by this characterisation it is supposed that also the success rate of the meta-inductivist or at least its upper bound with respect to the best performing method(s) is limited.

   (From these two assumptions it follows that at least one object-method is predictively successful.)

In the continuous prediction setting the access-optimality results show that there is no demonic scenario possible. Even if the world were demonic in the sense that whenever Pr_{mi} predicts an event's outcome ($Pr_{mi}(E_t) = r$), then the true outcome differs maximally from the prediction ($e_t = 1$ if $r < 0.5$ and $e_t = 0$ if $r \geq 0.5$), even for this case the meta-inductive access-optimality results show that also the object-level methods cannot perform better. So they also would not gain success above that of the meta-inductivist in the long run.

Things are different in the discrete prediction setting: Here it is possible that object-level methods are predictively successful, although due to necessary rounding the meta-inductivist method remains unsuccessful. (Cesa-Bianchi and Lugosi 2006, p.67) provide a nice example for such a demonic scenario:

- Let $Bel_1(E_t) = 1$ and $Bel_2(E_t) = 0$ for all t

- Now, let the true binary series of outcomes $e_1, e_2, \ldots$ be demonic with respect to Bel_{mi}, i.e. $Bel_{mi}(E_t) = 1 - e_t$ for all t

Then Bel_{mi} must oscillate between preferring Bel_1 and preferring Bel_2 and $\lim\limits_{t\to\infty} success(Bel_1, E_t)) = \lim\limits_{t\to\infty} success(Bel_2, E_t)) = 0.5$, whereas $\lim\limits_{t\to\infty} success(Bel_{mi}(E_t)) = 0$.

In the demonic scenario an equal success rate in the long run among the object-level methods is the main cause for the meta-inductivist's switching. The success rates of at least the best-performing object-methods have to be equal, because otherwise the meta-inductivist would in the long run simply switch to that method with the highest success rate (if there is one best method in the setting, then Pr_{mi} converges to the meta-method *imitate the best*—cf. Schurz and Thorn 2016, sect.7). Only due to the equal success rates of the best performing object-methods and the meta-inductivist's oscillating between them the method is prone to demonic failure.

An epistemic consequence one might draw is that depending on the nature of events—whether they are continuous or discrete in nature—one is able to provide a justification of induction or not. In this respect one might see a parallel between the epistemic problem at hand and one that was discussed more than three centuries ago regarding the role of continuous or discrete nature in the applicability of formal methods: So, e.g., Gottfried Wilhelm Leibniz famously argued for the assumption of a continuous nature: "Everything goes by degrees in nature, and nothing by leaps, and this rule regarding changes is a part of my law of continuity" (Leibniz 1896, Book IV, Sect.16, p.552); one main motivation behind his argument for continuity in nature was to make the—by him developed—formal mathematical tools applicable to a study of nature:

> "[...] Leibniz considered "transitions" of any kind as continuous. Certainly he held this to be the case in geometry and for natural processes, where it appears as the principle *Natura non facit saltus*. According to Leibniz, it is the Law of Continuity that allows geometry and the evolving methods of the infinitesimal calculus to be applicable in physics." (Bell 2013, sect.4)

[427] Bearing such an analogy in mind, is the meta-inductivist also forced to argue for a continuous, more induction-friendly, nature (cf. Schurz 2008, p.299)? It seems not. At least it seems not to be the most preferable way to go.

What seems to be more preferable is to find ways of approximating the access-optimality results for a continuous setting within a discrete one. In the literature two approaches are proposed. We will discuss them in the following two subsections. Afterwards we are going to expand the investigation to synchronised continuous and discrete settings with the aim of ruling out demonic scenarios via such a synchronisation.

3.1 The Randomisation Approach

The randomisation approach is common in machine learning and tries to overcome the gap between access-optimality in a continuous and discrete setting via randomly picking out a prediction in such a way that the outcome is still biased towards an access-optimal prediction method (cf. Cesa-Bianchi and Lugosi 2006, chpt.4). The idea is as follows: In predicting an event outcome one does not consider only past event outcomes, but all possibilities of past and present event outcomes; then one defines a prediction method that—regarding the binary setting—randomly predicts 0 or 1, but is—regarding a continuous setting—biased towards the ideal calculated value of the continuous setting. So, averaging over all possibilities, the method approaches the ideal calculated value in the finite case and reaches it in the long run. The details are as follows—this presentation is in accordance with (Schurz 2019, chpt.6.7.1): In order to explain the randomisation approach in detail, we expand the prediction setting further by the following elements:

- $Bel_{rmi}(E_t)$ is a qualitative prediction on E_t by a randomising meta-level agent
- $\mathbb{E}_1, \mathbb{E}_2, \ldots$ is an infinite series of an infinite series of events; we identify $\mathbb{E}_1$ with the infinite series of events above: $\mathbb{E}_1 = E_1, E_2, \ldots$; and we use sub-sub-indices to pick out specific events: $\mathbb{E}_{1_1} = E_1$; analogously we refer to the outcome of the single events by $\mathbb{e}$, as, e.g., in $\mathbb{e}_{1_1} = e_1$; finally, in the definitions of the score, success, and weight of an agents' prediction the series of events is always restricted to that provided in the argument place;

The binary randomising meta-inductive agent Bel_{rmi} predicts within the limits of:

$$\mathbb{P}(Bel_{rmi}(E_t) = 1) \approx Pr_{mi}(E_t)$$

Here $\mathbb{P}(Bel_{rmi}(E_t) = 1)$ is the frequency of Bel_{rmi} predicting 1 in round t of all $|\{0,1\}^t|$ possible series of binary event outcomes. Take, e.g., the outcomes of series of events as given in table 1 with the object-level predictions Bel_1 and Bel_2 (where $\mathbb{e}_1 = e$ is still considered to be the true series of outcomes, the other series of outcomes $\mathbb{e}_2$–$\mathbb{e}_8$ are the past outcomes, that are up to $t = 3$ possible; up to $t = 4$ there are 16 series possible, including the true outcome, etc.). Then a

	$t=1$	$t=2$	$t=3$	...
$\mathbb{e}_6$	0	0	0	
$\mathbb{e}_2$	0	0	1	
$\mathbb{e}_3$	0	1	0	
$\mathbb{e}_4$	0	1	1	
$\mathbb{e}_5$	1	0	0	
$\mathbb{e}_1$	1	0	1	
$\mathbb{e}_7$	1	1	0	
$\mathbb{e}_8$	1	1	1	
Bel_1	1	1	1	1
Bel_2	0	0	0	0
Pr_{mi}	0.5	1.0	0.5	0.6

Table 1: Example of predictions of two object-methods (Bel_1, Bel_2) and one success-based weighting meta-method (Pr_{mi})

randomising meta-inductive method within the above stated limits would predict, e.g., according to table 2. Of course there are numerous other randomising meta-inductive methods possible; important is only that their (weighted) average over all possible event series $\overline{Bel_{rmi}(\mathbb{E})}$ coincides with the prediction made according to weighted-average meta-induction (Pr_{mi}).

	$t=1$	$t=2$	$t=3$	$\dots$
$Bel_{rmi}(\mathbb{E}_1)$	1	1	0	1
$Bel_{rmi}(\mathbb{E}_2)$	1	1	0	1
$Bel_{rmi}(\mathbb{E}_3)$	1	1	0	1
$Bel_{rmi}(\mathbb{E}_4)$	1	1	0	1
$Bel_{rmi}(\mathbb{E}_5)$	0	1	1	1
$Bel_{rmi}(\mathbb{E}_6)$	0	1	1	0
$Bel_{rmi}(\mathbb{E}_7)$	0	1	1	0
$Bel_{rmi}(\mathbb{E}_8)$	0	1	1	0
$\overline{Bel_{rmi}(\mathbb{E})}$	0.5	1.0	0.5	0.6

Table 2: Example of a ranomised success-based meta-method: Such a method predicts in the binary case on average as often 1 as its real-valued prediction would be. So, e.g., given the real-valued predictions of table 1, it predicts in 50% of t_1-cases (where $Bel_1(E_1) = 1$ and $Bel_2(E_1) = 0$) 1 and in 50% of such cases 0. Analogously for all other cases. Which prediction the meta-method makes in the end is chosen randomly/arbitrarily, but biased towards the real-valued prediction. For the optimality result important is the fact that the exact choice of the meta-method $Bel_{rmi}(\mathbb{E}_{i_t})$ is probabilistically independent from the true outcome $\mathbb{E}_{1_t}$.

Now, assume that the pattern of the binary sequences in $1 \leq t \leq 3$ goes on this way; as can be seen in the tables above, a randomising meta-inductivist method would not approach the best predictor's success rate in every possible event series. However, one might wonder whether something weaker can be shown, like that on average randomisation in a discrete setting based on a continuous success-based prediction is long run [428] access-optimal? Indeed, if one defines a measure for expected success via the success of Bel_{rmi} with respect to an event series $\mathbb{E}_k$ weighted by the probability of $\mathbb{E}_k$ itself (which is a function of Bel_{rmi}'s predictions), then one can transfer the optimality result for expected success to the discrete setting. The quite complicated formula for expected success is as follows (the big product produces the value for the probability of each event series $\mathbb{E}_k$):

$$expsuccess(Bel_i, E_t) = \sum_{k=1}^{|\{0,1\}^t|} \prod_{l=1}^{t} (1 - Bel_i(\mathbb{E}_{k_l}) - e_{k_l} + 2 \cdot Bel_i(\mathbb{E}_{k_l}) \cdot e_{k_l}) \cdot success(Bel_i, \mathbb{E}_{k_t})$$

Under a specific assumption, in the discrete setting the bounds for expected success of the randomising meta-inductivist are analogous to that of the weighted-average [429] meta-inductivist's prediction in the continuous setting (cf. Schurz 2019, sect.6.7.1, prop.6-11):

$$
\begin{aligned}
expsuccess(Bel_{rmi}, E_t) &\leq 1 \\
expsuccess(Bel_{rmi}, E_t) &\geq expsuccess(Bel_i, E_t) - \sqrt{\frac{n}{t}} \quad \text{for all } 1 \leq i \leq n
\end{aligned}
$$

Again, strict access-optimality holds for the long run only. Crucial for these bounds is an independence assumption stating that the true outcome and the prediction of the randomising meta-inductivist are probabilistically independent in the following way:

$$
\mathbb{P}(Bel_{rmi}(E_t) = 1 | E_t = 1 \ \& \ Pr_{mi}(E_t) = r) = \\
\mathbb{P}(Bel_{rmi}(E_t) = 1 | Pr_{mi}(E_t) = r) \ \text{ for all } r \in [0, 1]
$$

A feature of randomisation in discrete settings is its structural closeness to the continuous case. However, considering the independence assumption above it is clear that a demonic scenario is ruled out only by stipulation. Furthermore, the relativisation of the optimality result to expected predictive success instead of predictive success *per se* opens another dimension into the infinite whose trend is even opposed: Whereas in the continuous case strict access-optimality is restricted to the long run, i.e. to infinite series of predictions, in the randomising approach access-optimality is restricted to the long run as well as to weighted averaging among the set of possible outcomes; since the number of possible event outcomes increases with the number of predictions, in the long run, information about expected success decreases.

Now we go on to consider another proposal that is about access-optimality of predictive success *per se* in a discrete setting.

3.2 The Meta-Meta Approach

In (Schurz 2008, sect.8) a set of qualitative meta-inductive prediction methods is defined which, on average, transforms access-optimality results for a continuous setting to the discrete realm. The idea is as follows: If one wants to approach a value of a continuum by help of discrete values, one may arrange discrete values around the value of the continuum in such a way that the average of the discrete values is close to the value of the continuum. E.g., one can approach/reach $0.5 \in [0, 1]$ by averaging over the elements of $\{0, 1\}$. Similarly for $0.75 \in [0, 1]$ by averaging over elements of $\{0, 1\}$: $0.75 = (0 + 1 + 1 + 1)/4$. Now, in a discrete setting, like the binary setting, only discrete predictions, e.g., binary predictions, are admissible. So, every method can predict only a value out of $\{0, 1\}$. However, the number of prediction methods is in principle not fixed. This can be exploited by a meta-strategy by settling around the value of a continuum 0/1-predicting methods in such a way that on average the value of the continuum is approached. So, e.g., if the calculated ideal prediction is 0.25, then the meta-method can approach it by averaging over one 1-predictor and three 0-predictors: $0.25 = (1 + 0 + 0 + 0)/4$. In the binary prediction setting no meta-method can exploit this fact directly, because averaging over the predictions leaves the binary value space. However, on a meta-meta level where one can compare successes of object- and meta-methods as, e.g., we are doing, a meta-meta-method can average over the single prediction method's success and can exploit this on the meta-meta-level.

[430] In order to indicate such a meta-meta-method, we add to the discrete prediction setting a group of binary meta-inductive methods:

- $Bel_{mi_1}(E_t), \ldots, Bel_{mi_k}(E_t)$ are the qualitative predictions on E_t of k meta-level agents

Now, Gerhard Schurz has found an interesting way of emulating real-valued success-based predictions in the discrete setting by defining the meta-inductive predictions as follows ($[\cdot]$ rounds to the next integer, as, e.g. $[0.75] = 1, [0.25] = 0, [0.5] = 1$):

$$\begin{aligned} Bel_{mi_i}(E_t) &= 1 \quad \text{if } i \leq [Pr_{mi}(E_t) \cdot k] \\ &= 0 \quad \text{otherwise} \end{aligned}$$

So, if, e.g., $k = 10$ and the ideal (continuous) predicted value $Pr_{mi}(E_t) = 0.75$, then the first seven meta-inductivists predict 1 ($1, \ldots, 7 \leq 0.75 \cdot 10$), and the remaining three meta-inductivists predict 0 ($8, \ldots, 10 > 0.75 \cdot 10$). By this a meta-meta-inductivist can exploit the meta-inductivists' predictions by averaging and approximating 0.75 by 0.7. In this case, using only a subset of four meta-inductivists would perform better. It turns out that, although each meta-inductivist's success rate is not bounded by the object-level methods' success rates, the average of them is (cf. Schurz 2008, p.299):

$$\begin{aligned} \overline{success(Bel_{mi}, E_t)} &\leq 1 \\ \overline{success(Bel_{mi}, E_t)} &\geq success(Bel_i, E_t) - \sqrt{\frac{n}{t}} - \frac{1}{2 \cdot k} \quad \text{for all } 1 \leq i \leq n \end{aligned}$$

So, for the long run, i.e. the limiting case, the distance of the average shrinks as a function of the number of meta-level methods k to $1/(2 \cdot k)$. The trick of this "collective average-weighting method" is to mimic access-optimality in the real-valued case, which is achieved by, lets say, $Pr_{mi}(E_t) = r$, via bringing $\overline{Bel_{mi}(E_t)}$ as close as possible to r. It is, so to say, lifting the binary predication game onto a meta-meta level of a prediction game with $1/(2 \cdot k)$ (with arbitrary high k) as the smallest approximatable unit.

Averaging success rates means that the meta-inductive agents $Bel_{mi_1}, \ldots, Bel_{mi_k}$ have to share their success. According to the general impossibility result regarding a demonic setting one cannot define a meta-method which takes as input the predictions of $Bel_{mi_1}, \ldots, Bel_{mi_k}$ and produces a prediction on its own. So, one might say that in order to deal with the problem of discrete predictions from a meta-inductive perspective one is forced to act as a collective.

The main advantage of this approach is to be found in its applicability to any prediction setting whatsoever. One also does not have to exclude a demonic scenario by stipulation, as is done in the randomization approach. Although in such a setting all meta-inductivist methods might perform suboptimally, on average these methods approximate access-optimal performance. However, it guarantees *approximation* of access-optimality only in the long run. As we have seen, the lower bound of the average success rate is in the long run $max(success(Bel_1, E_t), \ldots, success(Bel_n, E_t)) - 1/(2 \cdot k)$. Now, although k might be chosen arbitrarily high, one cannot achieve equal success rates in the long run. In order to achieve strict access-optimality one might think of introducing infinitely many meta-inductive methods and by this get: $lim_{k \to \infty} 1/(2 \cdot k) = 0$. However, as is shown in (Arnold 2010), a meta-method cannot be based on infinitely many methods in order to achieve access-optimality. So, for some demonic scenarios even the collective of meta-inductive methods will predict suboptimally.

[431] To sum up, discrete prediction settings allow for demonic scenarios; randomisation allows for weak access-optimality in the sense of convergence of expected success rates with that one of the best object-level method accessible, but at cost of stipulating independence between meta-inductive prediction and true outcome, thus stipulating that demonic scenarios are impossible. On the other hand, collective meta-induction allows for an approximation of average success as accurate as one wishes; however, strict convergence is not always possible and by this also a collective of meta-inductive methods performs suboptimal in at least some demonic settings, even in the long run. This facts seem to suggest that in order to approach the problem of induction within a discrete setting one has to enrich the structure of the problem and try to prove meta-inductive access-optimality or the impossibility of a demonic setting for such an enriched structure. This is the line of argumentation we are following in the next section by considering the problem of demonic settings within a synchronised prediction environment.

4 An Approach via Synchronisation Constraints

As we mentioned in the preceding section, we suggest structural enrichment for solving the problem of suboptimal success-based predictions in a discrete setting. The structure we are interested in is a synchronised setting. So, we combine a discrete prediction setting with a continuous one and put forward some synchronisation constraints. We aim at showing that, given these constraints, demonic scenarios are impossible. Demonic scenarios underlay the meta-inductivists suboptimality. So, arguing for the impossibility of such demonic scenarios in a synchronised setting is the same as to argue for the optimality of meta-induction in all reasonable synchronised settings.

The formalism we introduced above was chosen in such a way that we just have to re-interpret it. By this we can put forward the constraints for the new structure of the problem (or one might also call it 'the structure of the new problem'):

- The discrete prediction setting consists of:
  - $E_1, E_2, \ldots$ an infinite series of binary events whose outcomes $e_1, e_2, \ldots$ are within $\{0, 1\}$
  - $Bel_1(E_t), \ldots, Bel_n(E_t)$ which are, as before, the binary predictions on E_t (elements of $\{0, 1\}$) of the n object-level agents
  - $Bel_{mi}(E_t)$ which is, also as before, the binary prediction on E_t by the meta-level agent
- The continuous prediction setting consists of:
  - The same series of binary events
  - $Pr_1(E_t), \ldots, Pr_n(E_t)$ which are real-valued predictions on E_t (elements of $[0, 1]$) of the *n object-level agents*
  - $Pr_{mi}(E_t)$ which is a real-valued prediction on E_t of the *meta-level agent*

Now, *Bel* is interpreted as qualitative belief or acceptance in the sense that $\ulcorner Bel_i(E_t) = 1 \urcorner$ is supposed to mean that according to method i $e_t = 1$ or just simply: agent i believes that E_t will take place; analogously $\ulcorner Bel_i(E_t) = 0 \urcorner$ means that agent i believes that E_t will not take place. Note that we assume here that beliefs are complete in the sense that for every event E_t the agent either believes that it will take place or believes that it will not take [432] place. In principle one might try to relax the completeness condition by allowing agents to abstain from judgement. But then, of course, the question arises of how to adequately take into account abstention in scoring. We will stick to the idealisation of completeness, this the more since under specific circumstances expert knowledge also spreads from object-level methods to meta-level methods in a setting with restricted access only, where restricted access might be equalised with incompleteness (for details cf. Thorn and Schurz 2012, sect.7f).

Similarly we re-interpret *Pr*: $\ulcorner Pr_i(E_t) = r \urcorner$ is now supposed to mean that i's degree of belief that E_t will take place is r. (NB: In the discussion above $\ulcorner Pr_i(E_t) \urcorner$ is to be interpreted as the predicted value of E_t by i, which satisfies completely different truth-conditions, although due to the simple logical structure it is formally equivalent). Regarding scoring, this re-interpretation seems to be fine inasmuch as scoring can be directly related to betting behaviour. If outcomes are binary, it holds that the more an agent tends to extremes $(0, 1)$, the higher are also her chances in scoring well. But at the same time also her risk is of not scoring at all—so scoring tends also to the extremes then. And the more an agent tends to the indefinite (0.5), the safer she scores, but also the smaller the scores. In the case of a constant degree of belief of 0.5, expected predictive success will be equal to a randomisation among all possible event series as, e.g., is the case of flipping a fair coin—which is, to say the least, not a remarkably good benchmark.

Given such an expanded structure, what rationality constraints can be put forward? In the light of the re-interpretation provided above it seems to be appropriate to put forward synchronisation principles between the qualitative and the quantitative series of beliefs. The first principle we think of is a synchronisation principle that acts event-wise between these systems. It is a very specific case of the so-called *Lockean thesis* and states that a degree of belief above a specific threshold is necessary and sufficient for qualitative belief or acceptance. Since we are dealing with complete qualitative belief, the *natural* threshold is 0.5. Otherwise the situation could arise that one qualitatively believes a proposition and disbelieves its negation, although one's degree of belief in the proposition is strictly lower than the degree of belief in the negation, which sounds at least paradoxical (cf., e.g., Leitgeb's critique of Lin & Kelly's approach in Leitgeb 2017). Since this synchronisation principle is event-wise, we call it a 'synchronous synchronisation principle'. It is as follows:

$$\begin{aligned} Bel_i(E_t) &= 1 \text{ if } Pr_i(E_t) > 0.5 \\ &= 0 \text{ otherwise} \end{aligned} \tag{SynSync}$$

The case of $Pr_i(E_t) = 0.5$ is, epistemically speaking, not clearly regulated—regarding complete belief one might belief or disbelief that E_t will take place. For our argument below one can uphold a principle (SynSync*) similar to (SynSync) above, where $Pr_i(E_t) = 0.5$ enforces one to set $Bel_i(E_t) = 1$. What matters only is that all cases of $Pr_i(E_t) = 0.5$ are treated the same way, i.e. enforce either $Bel_i(E_t) = 0$ or $Bel_i(E_t) = 1$. (SynSync) can be also expanded to discrete settings with more than two admissible qualitative predictions. If, e.g., there are three qualitative predictions admissible ($\{0, 0.5, 1\}$), then one could state that $Bel_i(E_t) = 1$, if $Pr_i(E_t) > 2/3$, $Bel_i(E_t) = 0.5$, if $1/3 < Pr_i(E_t) \leq 2/3$, and $Bel_i(E_t) = 0$, if $Pr_i(E_t) \leq 1/3$. However, such a general bridging between the quantitative and the qualitative realm needs further argumentation; in the literature often so-called non-epistemic values are cited for such a bridging (cf., e.g., Longino 2008). Note that the binary meta-inductive prediction method Bel_{mi} of section 3 satisfies this constraint by definition. For all object-level methods in a demonic scenario (SynSync) poses no problem, since they can [433] easily pick out a (partial) probability function that satisfies for each event this constraint. In the demonic example mentioned above the agents with $Bel_1(E_t) = 1$ and $Bel_2(E_t) = 0$ might simply equate their constant qualitative belief with their quantitative one.

Beside this constraint we suggest another one for diachronic considerations. The idea is as follows: An agent i might believe or disbelieve and have an event-wise synchronised degree of belief above or below the threshold 0.5 regarding the event's taking place or not. However, the event-wise synchronisation does not oblige i to synchronise her degrees of belief according to her qualitative predictions in the long run. Take, e.g., an agent i with the alternating acceptance behaviour and equal degrees of belief according to table 3. Although both epistemic attitudes towards E_t are event-wise synchronous, one might ask whether it is rational for i to stick to her degrees of belief also in the

	t=1	t=2	t=3	t=4	t=5	t=6	$\cdots$
$Bel_i(E_t)$	1	0	1	0	1	0	$\cdots$
$Pr_i(E_t)$	1.0	0.0	1.0	0.0	1.0	0.0	$\cdots$

Table 3: Example of qualitative and quantitative predictions

long run or whether she should at some point in time s adopt her degrees of belief also according to her past prediction behaviour? We think that in order to be diachronic synchronous too, agents should also calibrate—although it should be highlighted that this is already a much stronger assumption than that of synchronous synchronisation.

How should calibration in such a setting work? Of course we do not suggest to oblige the agent to calibrate directly according to the past outcomes—this would lead to a constraint of applying the straight rule. The straight rule is to be considered as a possibility for an object-level method; but it (or convergence in the limit with it) is not to be considered as a necessary condition for rationality. We have seen in the discussion of section 2 that Reichenbach's argument for the straight rule holds only in "predictable cases", but not in general as, e.g., in demonic scenarios. On the other hand, calibration according to past predictions alone seems to be a constraint too weak to be upheld. It would not ground the agent's degrees of belief to (experimental) outcomes at all. In the example given above the agent was obliged to calibrate her degrees of belief in the long run, i.e. starting at a point in time s, to 0.5; and this regardless of the past outcomes. What seems to be more reasonable is to demand calibration with respect to one's own predictions *and* the past outcomes. Recall that according to the definition above, *success* combines both via keeping track of an agent's *true* predictions in comparison to *all predictions* made so far by her. For this reason we suggest as a middle ground between purely outcome-oriented calibration and calibration that is based on past predictions only: success-oriented calibration for diachronic synchronicity. So, we think that in the long run an agent's degrees of belief should be calibrated by her success rate in the following way:

There is a s, such that for all $r \geq s$:

$$\begin{aligned} Bel_i(E_r) = 1 \;\Rightarrow\; Pr_i(E_r) &= \lim_{t\to\infty} success(Bel_i, E_t) \\ Bel_i(E_r) = 0 \;\Rightarrow\; Pr_i(E_r) &= 1 - \lim_{t\to\infty} success(Bel_i, E_t) \end{aligned} \tag{DiaSync}$$

The principle (DiaSync) states about quantitative belief, in order to be diachronically synchronised, the following: There is a point in the series s such that for all events following E_s, i.e. for all E_r with $r \geq s$, the quantitative belief regarding E_r's taking place ($Pr_i(E_r)$) [434] equals the limiting success rate regarding Bel_i. It is clear that (DiaSync) holds only, if the success rate is limited. This means that there is a point in the series where the success rate is fixed, where the object method reached an "equilibrium" regarding closeness of the predictions to the truth. The idea is that s is after such a limiting point.

(DiaSync) can be expanded also to a discrete setting where the number of admissible predictions is greater than two, not only in $\{0, 1\}$. So if, e.g., the admissible quantitative predictions are in $\{0, 0.5, 1\}$, then for $Bel_i(E_r) = 1$ and $Bel_i(E_r) = 0$ things may remain as in (DiaSync); and with respect to $Bel_i(E_r) = 0.5$ the degree of belief in E_r's taking place may be equalised with the value in-between them; the third value of such a discrete setting may then be plausibly interpreted as *suspension of judgement* (the outcome may be interpreted as indetermined). That there is always a plausible interpretation for a qualitative value in such an extended diachronic synchronisation principle is, of course, not guaranteed. But if there is a "bridge" between the qualitative and quantitative system under investigation, then it seems that one can also make sense of an extended diachronic synchronisation principle.

In our description of demonic scenarios in section 3 we have stipulated that in such scenarios the success rates of the object-methods are limited. So, (DiaSync) is supposed to hold for quantitative beliefs in such scenarios. If an agent believes that an event E_r will take place, then her degree of belief in E_r's taking place should cohere with her past performance in predicting E-events. And if an agent believes that an event E_r will not take place, then her degree of belief in E_r's not taking place should—completeness of belief presupposed—equal the inverse of her degree of belief in E_r's taking place. We have argued above that just considering the event outcomes in calibration would be inadequate since it would enforce the straight rule. Such a calibration principle might be considered as a *purely* empirical constraint. On the other hand, just calibrating according to one's past predictions seems to be without any empiristic spirit at all. An agent would be considered diachronically rational if she just sticks to her method. In case the used method is *a priori*, also the calibration principle would lead from *a priori* predictions to *a priori* ones. Hence, one might consider such a principle as rationalistic in spirit. By stipulating diachronic coherence of predictions through equalising degrees of belief with limiting success rates (if they exist), we think one gets the right spin from both camps: One remains in an empirically informed way with one's method.

It should be mentioned here that the diachronic synchronisation principle is used not as a reflection principle in our argument. It is not intended that *de facto* an agent should be supposed to update her degrees of belief according to her success rate—since this information is available only for the limiting case such an application would be too much to ask for. However, we, talking about demonic scenarios and having knowledge about the limiting case, may reasonably put forward constraints also for this case. And we think that from this perspective (DiaSync) is reasonable to ask for. Note that also the meta-inductive solution to the problem of induction holds strictly speaking only for the limiting case—only for this case it can be shown that the meta-inductive weighting method is among the best accessible methods within the setting (although, of course, the short run results demonstrate some kind of "epistemic controllability" by help of meta-induction). In order to uphold access-optimality (DiaSync) just adds another consideration to the limiting case: Meta-induction remains access-optimal also in a setting where discrete predictions are coupled with continuous ones, if all the agents within the setting are diachronically coherent, i.e. calibrated.

According to this proposal, the alternating predictions above would force an agent to calibrate her degrees of belief depending on the outcomes of the events as given in table 4. [435] In the first case, predictions are in complete

	t=1	t=2	t=3	t=4	t=5	t=6	$\cdots$	$\cdots$
$Bel_i(E_t)$	1	0	1	0	1	0	$\cdots$	$\cdots$
E_t	1	0	1	0	1	0	$\cdots$	$\cdots$
$success(Pr_i, E_t)$	$\frac{1}{1}$	$\frac{2}{2}$	$\frac{3}{3}$	$\frac{4}{4}$	$\frac{5}{5}$	$\frac{6}{6}$	$\cdots$	1
E_t	0	1	0	1	0	1	$\cdots$	$\cdots$
$success(Pr_i, E_t)$	$\frac{0}{1}$	$\frac{0}{2}$	$\frac{0}{3}$	$\frac{0}{4}$	$\frac{0}{5}$	$\frac{0}{6}$	$\cdots$	0
E_t	1	1	1	1	1	1	$\cdots$	$\cdots$
$success(Pr_i, E_t)$	$\frac{1}{1}$	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{2}{4}$	$\frac{3}{5}$	$\frac{3}{6}$	$\cdots$	$\frac{1}{2}$
E_t	0	0	0	0	0	0	$\cdots$	$\cdots$
$success(Pr_i, E_t)$	$\frac{0}{1}$	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{2}{4}$	$\frac{2}{5}$	$\frac{3}{6}$	$\cdots$	$\frac{1}{2}$

Table 4: An example of predictions and their corresponding success-rates showing that object-predictors in a demonic scenario are not calibrated according to (DiaSync)

agreement with outcomes, so it seems to be plausible that an agent trusts completely in her prediction method (regarding qualitative belief) in the long run; analogously in the second case, where predictions are in complete disagreement with outcomes; here it seems to be plausible that an agent distrusts her prediction method (regarding qualitative belief) completely in the long run. Finally, in the third and fourth case, where just 50% of the predictions are correct, an agent should trust in her method (regarding qualitative belief) no more, but also no less, than trusting in flipping a fair coin. Note that the last two cases represent the object methods in our example of a demonic scenario (we just switched the values predicted by the methods with that of the event outcomes here).

Let us apply the framework presented above to the problem of demonic scenarios. Now, as was pointed out in section 3, in a demonic setting the success rates of the relevant, i.e. the best, object-level agents converge. So, it holds:

$$\lim_{t\to\infty} success(Bel_1, E_t) \quad = \quad \lim_{t\to\infty} success(Bel_2, E_t)$$

But then, in order to be diachronically synchronised, the degrees of belief of the agents are also calibrated equally in the long run: By (DiaSync) we get for some point s in E (in case s differs agent-wise one has to choose the "larger" one):

There is a s, such that for all $r \geq s$:

$$\begin{array}{rcl} Bel_1(E_r) = 1 \;\Rightarrow\; Pr_1(E_r) & = & \lim_{t\to\infty} success(Bel_1, E_t) \\ & & \qquad \parallel \\ Bel_2(E_r) = 0 \;\Rightarrow\; Pr_2(E_r) & = & 1 - \lim_{t\to\infty} success(Bel_2, E_t) \end{array}$$

Since in the binary case with two admissible predictions of the demonic scenario the object-level agents' success rates are 0.5, by (SynSync) we get indiscernibility of qualitative beliefs, i.e. we get for some point s in E that for all $r \geq s$:

$$Bel_1(E_r) \;=\; Bel_2(E_r)$$

But then the meta-level agent Pr_{mi} and her qualitative counterpart Bel_{mi} would at some point in E predict exactly the same way as both object-level agents Pr_1, Bel_1 and Pr_2, Bel_2 [436] do. So the object-level and the meta-level methods' success rates would converge which means that the setting cannot be a demonic one.

This result also holds for a binary demonic scenario with more than two object methods, since their success rates still have to converge in order to be attractive for the meta-inductive method; by this, again, their degrees of belief converge (DiaSync); and by this, again, their qualitative beliefs converge (SynSync). In case the admissible predictions are not binary, but discrete to a degree greater than 2, also demonic scenarios are impossible. Consider, e.g., the case where $Bel_1(E_t)$ is constantly 1, $Bel_2(E_t)$ is constantly 0, and $Bel_3(E_t)$ is constantly 0.5. Their success rates also have to converge and can be maximally $1/3$. But then, by an expanded version of (DiaSync), $Pr_1(E_r)$ (for all $r \geq$ some s) equals also a value $\leq 1/3$. However, this would be incoherent with an expanded version of (SynSync), enforcing, e.g.: $Pr_1(E_r) > 2/3$.

5 Conclusion

As we have seen in the foregoing sections, meta-induction is a quite powerful approach for the epistemic justification of inductive methods. However, the approach falls short of its aim when inductive methods are applied to discrete prediction settings. In such settings, so-called *demonic scenarios* are possible which allow for suboptimality in meta-inductive predictions in the long run.

Demonic scenarios are characterised as prediction settings where the methods that are to be justified fail in all their predictions, whereas rival methods gain at least some predictive success. In order to overcome this problem, two solutions were suggested in the literature. One is the randomisation approach according to which the discrete meta-method should predict randomly, but with a bias towards the continuous meta-method's prediction. A shortcoming of this solution is that it does yield to access-optimality only, if one presupposes probabilistic independence between the prediction of the meta-method and the true outcome which is the same as stipulating the impossibility of demonic scenarios. The other approach is the theory of collective weighted-average meta-induction which introduces a collective of meta-level methods whose average is proximally access-optimal. A shortcoming of this solution is its incapability of proving strict access-optimality and that the needed collectives are not generally available.

For this reason we suggested to enrich the formal structure of the problem by combining a discrete and a continuous prediction setting. The former is about qualitative belief; the latter about degrees of belief. To keep both systems synchronised we suggested two synchronisation principles: an event-wise synchronisation principle (SynSync) according to which qualitative and quantitative belief should be bridged via a threshold (simple version of *Lockean thesis*). And a diachronic synchronisation principle (DiaSync) according to which both beliefs should be bridged via calibration by help of success rates. Our main argument against the possibility of a demonic scenario in such synchronised settings runs as follows:

1. A demonic setting with successful agents enforces different qualitative beliefs, but equal success rates in the long run among the relevant object-level agents. (cf. sect. 3)

2. Equal success rates in the long run enforce equal calibration of degrees of belief. (cf. DiaSync)

3. [437] Equal calibration of degrees of belief enforces equal qualitative beliefs. (cf. SynSync)

4. Hence, no demonic setting satisfies synchronic and diachronic synchronisation constraints at the same time. (1–3)

So, in case of a richer structure of prediction tasks meta-inductive suboptimality can be overcome by putting forward synchronisation constraints: If all agents in the combined qualitative and quantitative setting are rational in the sense that they are synchronically and diachronically synchronised (calibrated), then no demonic scenarios are possible and by this meta-induction remains access-optimal.

Acknowledgements

This research was funded by the *Austrian Academy of Sciences* (*OeAW*, DOC grant) and the Deutsche Forschungsgemeinschaft (DFG priority program *SPP 1516 "New Frameworks of Rationality"*). I would like to thank two anonymous referees for very helpful comments. For valuable discussions on this topic I would like to thank Gerhard Schurz.

References

Arnold, Eckhart (2010). "Can the Best-Alternative Justification Solve Hume's Problem? On the Limits of a Promising Approach". In: *Philosophy of Science* 77.4, pp. 584–593. DOI: 10.1086/656010.

Bell, John L. (2013). "Continuity and Infinitesimals". In: *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*. Ed. by Zalta, Edward N. URL: http://plato.stanford.edu/archives/win2015/entries/continuity/.

Cesa-Bianchi, Nicolo and Lugosi, Gabor (2006). *Prediction, Learning, and Games*. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9780511546921.

Descartes, René (1637/1998). *Discourse on Method*. Ed. by Cress, Donald A. Indianapolis: Hackett Publishing Company.

Howson, Colin (2003). *Hume's Problem*. Oxford: Clarendon Press.

Leibniz, Gottfried Wilhelm (1896). *New Essays Concerning Human Understanding*. Translated from the Original Latin, French and German with Notes by Alfred Gideo Langley. New York: The Macmillan Company.

Leitgeb, Hannes (2017). *The Stability of Belief. How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.

Longino, Helen E. (2008). "Values, Heuristics, and the Politics of Knowledge". In: *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*. Ed. by Carrier, Martin, Howard, Don, and Kourany, Janet A. Pittsburgh: University of Pittsburgh Press, pp. 68–87.

Reichenbach, Hans (1938). *Experience and Prediction. An Analysis of the Foundations and the Structure of Knowledge*. Chicago: University of Chicago Press.

Schurz, Gerhard (2008). "The Meta-Inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem". In: *Philosophy of Science* 75.3, pp. 278–305. DOI: 10.1086/592550.

— (2019). *Hume's Problem Solved. The Optimality of Meta-Induction*. Cambridge, Massachusetts: The MIT Press.

Schurz, Gerhard and Thorn, Paul D. (2016). "The Revenge of Ecological Rationality: Strategy-Selection by Meta-Induction Within Changing Environments". In: *Minds and Machines* 26.1, pp. 31–59. DOI: 10.1007/s11023-015-9369-7.

Thorn, Paul D. and Schurz, Gerhard (2012). "Meta-Induction and the Wisdom of Crowds". In: *Analyse und Kritik* 34.2, pp. 339–366.