

Optimality justifications: new foundations for foundation-oriented epistemology

Gerhard Schurz¹ 

Received: 31 May 2016 / Accepted: 1 March 2017
© Springer Science+Business Media Dordrecht 2017

Abstract In this paper a new conception of foundation-oriented epistemology is developed. The major challenge for foundation-oriented justifications consists in the problem of stopping the justificational regress without taking recourse to dogmatic assumptions or circular reasoning. Two alternative accounts that attempt to circumvent this problem, coherentism and externalism, are critically discussed and rejected as unsatisfactory. It is argued that optimality arguments are a new type of foundation-oriented justification that can stop the justificational regress. This is demonstrated on the basis of a novel result in the area of induction: the optimality of meta-induction. In the final section the method of optimality justification is generalized to deductive and abductive inferences.

Keywords Optimality justification · Meta-induction · Foundation-oriented epistemology · Coherentism · Internalism · Externalism

1 Introduction

According to the traditional conception, knowledge is *justified* true belief. Thus it is justification which distinguishes knowledge from accidentally true belief resulting from lucky guesses. In contemporary epistemology the conception of knowledge as justified true belief is still widely held; however, the traditional foundationalist and internalist understanding of justification has been challenged.

✉ Gerhard Schurz
schurz@phil.uni-duesseldorf.de
<http://www.philosophie.hhu.de/schurz>

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

Since the beginning of the philosophical era of enlightenment in the sixteenth century, the concept of *justification* has played a central role in epistemological debates. The leading idea of justification in the era of enlightenment was foundation-oriented and internalist: To reach knowledge our system of beliefs should be justified not by religious or otherwise authority, but by reason—by means of a system of arguments by which all our beliefs can be soundly derived from a small class of *basic* beliefs and principles that are considered as immediately evident to everybody. This idea of justification was shared by the rationalistic wing (e.g., Descartes, Leibniz, Kant) and the empiricist wing (e.g., Locke, Hume, Mill) of enlightenment epistemology.

Historically the rise of enlightenment epistemology happened in parallel to the demise of the authority of religious world views. Within centuries of religious controversies and wars the confidence of intellectuals into the truth-conduciveness of religious belief systems collapsed and the insight into their irrationality spread. The enlightenment idea of knowledge as a system of beliefs based on rational discourse was and still is a major theoretical basis of the western conception of a constitutional democracy. Though being foundation-oriented and partly ‘foundationalistic’ (see below) the enlightenment concept of justification is everything else but ‘fundamentalistic’ in the sense of being based on (religious) creed instead of reason. It is decidedly anti-fundamentalistic in its rejection of all sorts of authoritarian dogmas as legitimate ingredients of justification and replaces them by the “forceless force of the better argument” (to use a phrase of Habermas).

In the development after the enlightenment era that led to the philosophical situation of ‘(post-)modernity’, the foundation-oriented program of epistemology came increasingly under attack. The chief criticism was not that it is misguided but that it is pretentious: its noble claim of presumption-free and universally acceptable standard of justifications is illusionary; too good to be true. The major challenge for the foundation-oriented program of justification is the problem of finding means to stop the justificational *regress*, i.e., the apparent necessity to base each justification upon premises which are themselves in need of justification.

There are two *dimensions* in which this regress problem arises: *First*, at the ‘horizontal’ level of beliefs (or statements) that are traced back to more and more basic beliefs—in modern terminology, this is the problem of 1st order justification. In regard to this dimension the enlightenment epistemologists ultimately arrived at the minimalist class of introspective and analytic beliefs, which they considered as the only beliefs that can be regarded as immediately evident. *Second*, at the level of arguments whose reliability has to be demonstrated by means of certain meta-arguments—in modern terminology, this is the problem of higher order justification. Hume’s skeptical arguments against the possibility of a non-circular justification of induction made it clear that a sustainable solution to this problem is extremely difficult, if not impossible.

In the first part this paper (Sects. 2–3) I will develop a new conception of foundation-oriented epistemology that is supposed to withstand the shortcomings of traditional foundationalism. On the basis of the problem of induction, I defend this conception against two important alternative programs in contemporary epistemology: (1) coherentism and (2) externalism. In the second part (Sects. 4–6) I will sketch a new method of solving the problems of circularity and regress: the method of epistemic optimality justifications. I illustrate this method by means of a new account to the problem

of induction that has been developed elsewhere: the optimality of meta-induction (Sect. 4). In Sect. 5 it is explained how the optimality account can be generalized to deductive logic and to abductive arguments and in Sect. 6 the place of the optimality account in the landscape of epistemological positions is located.

2 Foundation-oriented epistemology: explication and major problems

The account of foundation-oriented epistemology that is developed in this paper departs from traditional foundationalist accounts in three respects:

- (1) A careful distinction is made between *foundation-oriented* and *foundationalistic* approaches. Classical foundationalist epistemologies demand that the basic beliefs be epistemologically certain or necessary (cf. Dancy 1985, Chap. 4.1). Most contemporary epistemologists reject the infallibility requirement as too strong. Foundation-oriented approaches allow that even basic beliefs may be fallible and revisable. Still, basic beliefs enjoy an epistemological priority, insofar as (a) they are *more entrenched* than non-basic beliefs, and (b) they figure as *informational inputs* in the dynamical network of beliefs.
- (2) The account is committed to the idea of *meliorative* epistemology (cf. Shogenji 2007, Sect. 1; Schurz 2008c)—the idea that epistemology should help improving the epistemic practice of people, which was an important part of the enlightenment program. Connected with this idea is the hope that disagreements between competing world views and political parties can be solved peacefully by means of rational argumentation. In contrast, prominent defenders of contemporary analytic epistemology are skeptical towards the possibility of solving fundamental disagreement by rational argument (cf. Sosa 2010).
- (3) The apparently unsolvable problem of traditional foundationalisms, the problems of circularity and regress, is handled by the novel method of epistemic optimality justifications.

In the core of the proposed account is the notion of foundation-oriented justification. Based on the previous considerations this notion is explicated as follows:

- (FOJ) Explication of *foundation-oriented justification*: A system of justifications is *foundation-oriented* in the internalist sense iff it satisfies the following requirements:
- (R1) It attempts to justify all beliefs by means of chains of arguments whose ultimate premises consist of basic beliefs that are immediately evident.
 - (R2) Thereby it avoids complete justification circles because they are epistemically worthless.
 - (R3) Its justifications intend to be complete in the sense of providing higher order justifications of the reliability of the argument patterns employed in (R1), or at least of their optimality in regard to the goal of reliability.

Concerning (R1): Contemporary foundation-oriented epistemologies come in two variants, internalist and externalist (cf. Fumerton 1995, Chap. 3). We understand the notion of a “foundation-oriented justification” in the traditional *internalist* sense, as a

system of arguments terminating in premises expressing immediately evident beliefs. For externalists (e.g. Goldman 1986) the system of justification consists in (unconditionally or conditionally) reliable cognitive processes that need not necessarily be accessible to the epistemic subject. Our preference for the internalist account is based on meliorative considerations: the justifications of our belief system must be cognitively accessible, because inaccessible justifications are epistemically useless (see Sect. 3).

(R1)'s requirement that every justification chain must terminate in immediately evident beliefs excludes infinite regresses. The class of 'immediately evident' beliefs is understood in a *minimalistic* sense that is characterized below; this minimalistic understanding distinguishes the foundation-oriented program from dogmatic accounts.

Concerning (R2): This requirement excludes justification circles and discriminates foundation-oriented accounts from coherentist accounts, which allow circular justifications. In Sect. 3 we demonstrate that complete justification circles are epistemically worthless, because with their help one may pseudo-justify mutually inconsistent epistemic recommendations.

Concerning (R3): This condition requires the justification of the reliability of the employed arguments. It is sustained in all traditional and many contemporary internalist accounts, but rejected by all externalists and even some internalist accounts. We shall argue below that condition (R3) is particularly important for internalist accounts with meliorative purposes. Recent defenses of this condition can be found, for example, in Fumerton's account of inferential justification (1995, p. 36, p. 85), according to which being justified in believing p on the basis of believing evidence e entails (1) being justified in believing e and (2) being justified in believing that e makes p probable (i.e., that the argument from e to p is reliable). Another variant of (R3) is White's reliability principle (2015, p. 219), according to which a rational person can only be justified in believing a proposition p if she is justified in believing that the methods that led her to believe p are reliable.

Following from (R1) and (R2), higher order justifications must themselves be non-circular and free from invoking an infinite regress. For this reason we prefer to speak of "higher order" instead of "second-order" justifications. The latter notion (introduced by Alston 1976) invites the question why one should not also require third or fourth order justifications (etc.); but obviously the regress of meta-levels has to be stopped at some level. In Sects. 5–6 we try to show how this is possible by means of optimality justifications. The turn to optimality justifications is explicitly reflected in our formulation of (R3); it seems 'tiny' but constitutes the crucial novelty of our proposal. To be precise, an argument is characterized as reliable (to a degree of $r > 0.5$) iff the objective probability of its conclusion, given its premises, is high (has a degree of at least r); and it is characterized optimal within a class of competing arguments if this probability is maximal in the class of competing arguments.

Figure 1 illustrates the major components of an internalist foundation-oriented epistemology and at the same time reveals its major problems. Every foundation-oriented model of justification must, *first*, specify a class of *basic beliefs* which are taken as immediately (or at least as *prima facie*) evident and are not need of further justification. *Second*, it must specify argument patterns by which derived or non-basic beliefs can be traced back to basic beliefs. The three major types of argument

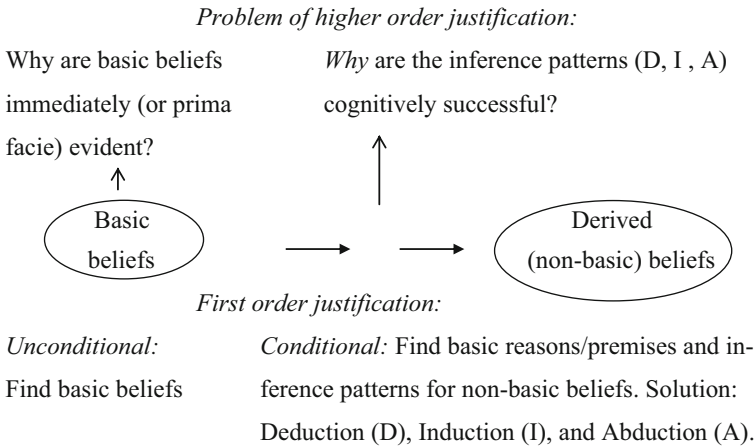


Fig. 1 Major components and problems of internalist foundation-oriented epistemology

patterns which have been established in analytic epistemology are *deduction*, *induction* and finally *abduction*, or inference to the best explanation. Specifying basic beliefs, and specifying deductive, inductive or abductive reasons for one’s derived beliefs, is what we call the task of first order justification—unconditional for basic beliefs, and conditional for derived beliefs. This part of justification is not only required according to philosophical but also according to common sense standards of rationality.

For a *complete* justification higher order justifications are required. The higher order justification of basic beliefs must explain *why* a particular class of beliefs that is regarded as basic can legitimately be considered as immediately (or at least as prima facie) evident; this is usually called the “problem of basic beliefs”. The higher order justification of inference patterns has to explain why the patterns of deduction, induction and abduction can be legitimately regarded as cognitively successful, in the sense of being reliable or at least optimal in regard to reliability. This problem is usually meant by the “problem of higher order justification” in the more narrow sense, i.e., applied to inferences. For deductive argument patterns the task of higher order justification is prima facie unproblematic, since we know by the semantic definition of logical validity that deductive arguments preserve truth in *strictly all* cases. This kind of justification is not enough if we want to defend classical logic against alternative logics (more on this in Sect. 5.1), but for the time being we take (classical) deductive logic for granted. On the other side, for induction and abduction the task of finding a non-circular higher order justification is extremely difficult, if not impossible.

The classical solution to the problem of basic beliefs has already been suggested by Augustinus and was taken over by early enlightenment philosophers such as Descartes, Locke, Berkeley and Hume. This solution was *minimalistic* insofar as it considered *only* two kinds beliefs as truly basic, i.e., free of doubt and not in need of further justification:

- (1) The *introspective* beliefs, which express facts about one’s conscious *experiences* (including sense experiences and inner experiences), without any implications

about the existence and constitution of a subject-independent external world. Our beliefs about an external reality may be in error. For example, the tree in front of me which I see right now may be a hallucination, a mere dream, or whatever. But what I know for sure is that *I have* this perceptual experience *now*—in my beliefs about my own conscious experiences I cannot be in error.

- (2) The *analytic* beliefs, by which we understand believed propositions that are either true because of the laws of (classical) logic, or because of accepted semantic definitions or meaning postulates for extra-logical concepts. Thus, analytically true sentences are true for purely logical or semantic reasons, independent from the factual constitution of the world. In contrast, *synthetic* sentences say something about the factual constitution of some part of ‘the world’ (which is the ordinary external world for synthetic realistic sentences and the world of ‘my’ experiences for synthetic introspective sentences).

Philosophers have raised objections to the classical solution of the problem of basic beliefs. It has been argued that even introspective beliefs are prone to error. For example, introspective beliefs about one’s past experiences rely on memory and memory is fallible. Moreover, if introspective beliefs are formulated in a public language, then one may be in error about the semantics of that language (Lehrer 1990, pp. 51–54, 64f). I agree with these objections. But I think that one should restrict introspective beliefs to one’s present experiences and one’s own private language. In this reconstruction, memory beliefs are introspective beliefs about one’s present memories (e.g., “I remember I have seen a table”). If introspective beliefs are restricted in this way, then I think they constitute at least optimal candidates for immediately evident basic beliefs. We need not assume that they are infallible, as foundationalistic positions do assume. It is sufficient to recognize that in *almost all* cases we can rely on our introspective reports; exceptions are only given when a brain or a mind acts in a completely schizophrenic way (cf. Fumerton 1995, p. 71).

Likewise, analytically true beliefs are obvious candidates for basic beliefs. The truth of logically true beliefs follows from the semantic laws characterizing logical concepts. Similarly, the truth of extra-logical meaning postulates (such as “bachelors are unmarried men”) follows from accepted semantic conventions for non-logical concepts. While the classical empiricists (from Locke to Hume) confined the class of a priori knowable propositions to the analytic ones, the classical rationalists (e.g., Descartes and Leibniz) thought that even some synthetic truths (e.g., certain laws of physics) can be justified on a priori grounds. The same claim was made by Kant, though for radically different philosophical reasons. However, most of the principles that the rationalists and Kant considered as a priori truths were refuted on empirical grounds by modern physics. Thus it seems us more reasonable to prefer the minimalistic solution, in accordance with the empiricist tradition, according to which the only beliefs whose truth is doubtlessly a priori are analytic statements.

The main problem with the *minimalist* solution of the problem of basic beliefs does not lie in the problem of justifying analytic beliefs, nor does it lie in possibility of erroneous introspective beliefs. Rather, it lies in the difficulty of *inferring* from this *small* class of basic beliefs anything *non-trivial* about that part of the world which lies *outside* of our consciousness, including our internal future as well as the external

world. Deductive inferences are clearly insufficient for this task, because by means of deductive logic one cannot infer conclusions that relevantly contain predicates not contained in the premises.¹ Thus what one can (relevantly) infer from introspective beliefs by means of deductive logic are merely other introspective beliefs.

In order to make inferences from beliefs about actual experiences to beliefs about the future or about general laws one needs *induction*. Moreover, to make an inference from introspective experiences to beliefs about an external reality which causes or explains these experiences one needs *abduction*, or inference to the best explanation. Note that we understand here the notion of induction in the narrow or ‘Humean’ sense, as an inference in which an observed regularity is transferred to either a new future instance (inductive prediction) or to the entire domain of individuals in space-time (inductive generalization). In contrast, by an abductive inference we understand an inference from an observed effect to a hypothetical cause or explanation which is as plausible as possible in the given background knowledge.²

In conclusion, the real problem of the ‘minimalistic’ solution to the problem of basic beliefs is that it shifts the burden of justification to the inference patterns of induction and abduction. So the task of providing higher order justification for these two inferences becomes enormously important. However, up to now no generally satisfying higher order justification of induction and abduction has been found in the philosophical literature. In the next two sections we focus on the problem of justifying induction, or Hume’s problem, and consider the problem of justifying abduction in the outlook.

3 Comparing foundation-oriented epistemology with contemporary alternatives: the problem of induction

Hume’s major skeptical challenges against the possibility of giving a rational justification of induction (1748, Chaps. 4, 6) can be summarized as follows:

- (1) Obviously, inductive inferences cannot be directly justified by *observation*, because the conclusions of inductive inferences are propositions about unobserved events.
- (2) It is likewise obvious that inductive inferences cannot be justified by deductive *logic*—for it is logically possible that from tomorrow on our world behaves completely differently than it has behaved so far. Therefore no inference from propositions about observed events to propositions about unobserved ones can be logically or analytically valid.
- (3) Induction cannot be justified by the standard method of empirical science—by induction from observation. This is the most important point in Hume’s skeptical

¹ If the conclusion of a deductive inference contains a predicate that doesn’t occur in the premises, then this predicate is completely irrelevant, in the sense of being replaceable by any other predicate *salva* validitate of the inference. This follows from the theorem of uniform substitution for predicates (cf. Schurz 1991). *Examples* (the underlined predicates are replaceable *salva* validitate): $Fa \models Fa \vee \underline{Ga}$, $\forall x(Fx \rightarrow Gx) \models \forall x(Fx \wedge \underline{Hx} \rightarrow Gx)$ (etc.).

² Inductive and abductive inferences are often subsumed under the umbrella notion of inductive inferences in the wide sense (cf. Pollock 1986, p. 42).

reasoning. To argue that the inductive method will be successful in the future because it has been successful in past applications would mean *justifying induction by induction*—which is a *circularity*. Since circular arguments presuppose what they purport to justify, they are without any justificatory value (in line with requirement (R2)).

- (4) Of course, inductive inferences are not strict entailments. But their justification should be able to show that they are *reliable* in the sense of being truth-preserving in a high *majority* of cases. As Hume argued in (1748, Sect. 6), this probabilistic reformulation—contrary to what some philosophers have proposed—is of no help as well. In order to justify that an inductive inference of the form “Most observed Fs have been Gs, therefore the next F will be a G” is truth-preserving in *most future cases*, we must presuppose that the relative frequencies of the events in the past can be transferred to the events in the future. This is nothing but a probabilistic version of the inductive generalization rule.

These have been the reasons which led Hume to the skeptical conclusion that induction is not capable of having a rational epistemic justification at all, but is merely the result of *psychological habit* (1748, part 5). Let us become clear how harsh Hume’s challenge against rationality really is. On the one hand, all results of empirical science, from physics to psychology, are based on induction. On the other hand, all sorts of human prejudice and superstition, from rain dancing to burning witches, are based on ‘psychological habit’. If there is no substantial difference between induction and psychological habit, then the enterprise of enlightenment rationality breaks down. Along these lines, Russell once remarked that if Hume’s problem cannot be solved, then “there is no intellectual difference between sanity and insanity” (Russell 1946, p. 699).

There have been several attempts in the analytic philosophy of the twentieth century to find possible or improved ways of justifying induction, but it seems that so far, none of these attempts has been successful. A similar diagnosis applies to the problem of justifying abduction, apart from the fact that this problem is even more difficult than that of justifying induction (see Sect. 5.2). If the problem of establishing non-circular higher order justifications of inductive (or abductive) inferences cannot be solved, then no belief which transcends our own consciousness can have a justification meeting the foundation-oriented demands and foundation-oriented epistemology must end in skepticism.

Faced with such a desperate situation, non-foundationalistic epistemologists have searched for possible ways to circumvent the task of higher order justification by weakening the justificational standard. In the rest of this section we illustrate on the basis of the problem of induction what in our view constitutes the major shortcoming of these two alternatives: their failure of being meliorative.

3.1 Coherentism

Here circular justifications are accepted. The kind of circularity that is of particular importance for coherentist account to the problem of higher order justification is *rule-circularity* (as opposed to premise-circularity; cf. Ladyman and Ross 2007, p. 75). An

argument is rule-circular if the truth of its conclusion is presupposed by the underlying inference rule. A typical rule-circular argument is the ‘inductive justification of induction’.

That a rule-circular argument may have epistemic value has been claimed by many contemporary epistemologists.³ However, it can be shown that rule-circular arguments can be used to ‘justify’ irrational rules and even rules with mutually contradictory conclusions. Especially powerful is the following objection of Salmon (1957, p. 46) against the inductive justification of induction. Salmon shows that the same type of rule-circular argument that ‘justifies’ the reliability of induction can also be used to justify the rule of anti-induction. The latter rule predicts (roughly speaking) the opposite of what has been observed in the past; so it predicts the opposite of what is predicted by the rule of induction:

(2) *Rule-circular justification of induction:*

Premise: Past inductions have been successful.

Therefore, by rule of induction:

Inductions will be successful in the future.

Rule-circular justification of anti-induction:

Premise: Past anti-inductions have been not successful.

Therefore, by rule of anti-induction:

Anti-inductions will be successful in the future.

Both circular ‘justification’ patterns have precisely the same structure, the premises of both arguments are true, and yet they have opposite conclusions. This shows that rule-circular argument patterns are *pseudo*-justifications: they cannot have epistemic value, because with their help one can pseudo-justify argument patterns with mutually contradictory conclusions.

3.2 Externalism

A proposition is called *internal* for an assumed subject iff it designates a state of the subject’s mind/brain that *is* cognitively *accessible* to this subject. This doesn’t imply that all of the subject’s internal states are *actually* conscious, but it implies that all of them *can* be brought to consciousness, for example by memory retrieval.⁴ On the other hand, facts or states of affairs are called *external* if they belong to the external reality outside of the cognitively accessible part of the subject.

More or less all philosophical accounts of justification until the 1960ies considered justification as an internal concept. The obvious reason for this understanding is that possessing a justification for one’s belief refers to a cognitive state or disposition of the epistemic *subject*. In contrast, externalist accounts characterize justification as a

³ Braithwaite (1974), Black (1974), Van Cleve (1984), Papineau (1993, Sect. 5), Goldman (1999, p. 85), Psillos (1999, p. 82)

⁴ This definition of “internalism” is also called accessibility-internalism, as opposed to state-internalism (Fumerton 1995, pp. 60–66).

property of the external world that need not be accessible to the subject. Most general is the externalist concept of justification proposed by Goldman (1986): He defines a belief to be externally justified iff this belief was formed by a cognitive process that is reliable in our world, which means that under normal circumstances this process leads to true beliefs with high objective probability. Goldman's account is also called *reliability-externalism*.

The trouble of purely externalist justifications—i.e., those that are not 'backed up' by an internalist justification—is their lack of cognitive accessibility, which deprives them from their meliorative function. Let me illustrate this point by means of the externalist treatment of rule-circular 'justifications'.

The stance towards the rule-circular 'justification' of induction and anti-induction highlights the characteristic differences between coherentism, foundation-'orientism' and externalism:

- For coherentist positions both arguments are acceptable justifications, which was criticized above as unacceptable.
- For foundation-oriented internalism, both arguments are pseudo-justifications.
- Justification-externalism leads to a third view: at most one of these two arguments can be acceptable as an externalist justification, but whether this is the case depends on external facts.

An example in point is van Cleve's externalist version of the inductive justification of induction. Van Cleve (1984, p. 562) makes the externalist 'correctness' of this argument dependent on the truth of a general fact—the reliability of induction—that is not stated as a premise and may be epistemically inaccessible. He writes: "... the antecedent on which this [justification] depends—that induction is reliable ... is an *external* antecedent. It makes knowledge possible not by being known, but by being true". As a consequence of this position, not only induction but also anti-induction may possess an externalist 'justification' of this sort, namely exactly if it is true that anti-induction is reliable (which is possible in worlds with permanently oscillating event-frequencies). Our argument can be summarized as follows:

(3) *Rule-circular justification of induction ... as in (2)*

Rule-circular justification of anti-induction: ... as in (2)

The internalist concludes from the perfect symmetry that both 'justifications' are epistemically worthless. In contrast, for the externalist both justifications can be 'correct' in the following sense:

The circular justification of induction is correct in worlds where inductive inferences are reliable.

The circular justification of anti-induction is correct in worlds where anti-inductive inferences are reliable.

However, without possessing a cognitively accessible higher order justification of induction one can impossibly know that the left argument is externalistically correct and the right one incorrect. This concludes our critique of justification-externalism from the meliorative perspective.

4 The optimality of meta-induction

According to the proposed foundation-oriented program of epistemology the major share of the justification load rests on the two types of content-expanding non-deductive arguments: inductions and abductions. Essential for this account is therefore the possession of higher order justifications of these inferences that can escape the problem of the justificational regress without resorting to dogma or circularity. In other papers (cf. Schurz 2008b, 2009) I have developed a new type of higher order justification for inductive inferences that I call *optimality justification*. Optimality justifications do not attempt to ‘prove’ that a cognitive method (here induction) is reliable—something that, by Hume’s arguments, cannot be done—but, rather, that it is *optimal*, i.e., that it is the best that we can do in order to achieve our epistemic goal, which in the case of induction is predictive success. In this section I explain the epistemic optimality account on the basis of the problem of induction. In the following section I will explain how this account may be generalized to abductive inferences and to systems of logic.

Reichenbach (1949, Sect. 91) was the first philosopher who suggested something like an optimality account: he attempted to demonstrate that induction is the best what we can do for the purpose of predictive success. Reichenbach’s attempt failed, because—as pointed out by Skyrms (1975, Chap. III. 4)—nothing in Reichenbach’s “best alternative” account can exclude that a clairvoyant may be better in predicting random sequences than an empirical inductivist. More generally, results in formal learning theory show that no prediction method can be universally optimal at the level of *object-induction*, that is, of induction applied to the task of predicting events in arbitrary possible worlds (cf. Kelly 1996, p. 263; Sterkenburg 2017). In contrast, my account is focused on the concept of *meta-induction*, i.e., induction applied at the meta-level of competing prediction methods.

Meta-induction tracks the success rate of all prediction methods whose predictions are *accessible* and predicts a weighted average of the predictions of those methods that were most successful so far. What my account attempts to show is that there is a meta-inductive strategy that is predictively optimal among all prediction methods that are (simultaneously) accessible to the epistemic agent. Since the restriction to accessible methods is crucial for the optimality theorem, Schurz and Thorn (2016) call this kind of optimality *access-optimality*. Remarkably, the access-optimality of meta-induction holds in *all* possible worlds, even in radically ‘non-uniform’ worlds or in ‘paranormal’ worlds that host perfect clairvoyants.

Technically the account of meta-induction is based on the notion of a prediction game:

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

- (1) An infinite sequence $(e) := (e_1, e_2, \dots)$ of events $e_n \in [0, 1]$ coded by real numbers ranging between 0 and 1, possibly rounded according to a finite accuracy. For example, (e) may be a sequence of daily weather conditions, football game results or stock values. Each time n corresponds to one round of the game.
- (2) A finite set of prediction methods or ‘players’ $\Pi = \{P_1, \dots, P_m, MI\}$. In what follows we identify ‘methods’ with ‘players’. In each round it is the task of each

player to predict the next event of the event sequence. “MI” signifies the meta-inductivist and the other players are the ‘non-MI players’ or ‘candidate methods’. They may be real-life experts, virtual players implemented by computational algorithms, or even ‘clairvoyants’ who can see the future in ‘para-normal’ possible worlds. It is assumed that the predictions of the non-MI players are accessible to the meta-inductivist.

Each prediction game constitutes a *possible world*, or in cognitive science terminology a possible environment. Apart from the above definition we make no further assumptions about these possible worlds. The sequence of events (e) can be arbitrary: a deterministic sequence, a random sequence or Markov chain, or a ‘chaotic’ sequence whose finite frequencies don’t converge to limits at all. We also do not assume a fixed list of players—the list of players may vary from world to world, except that it always contains MI, and some *fallback* strategy of MI in situations in which there are no other accessible players. Naturally this fallback strategy will be some object-inductive method, but it may also be blind guessing or whatever. The only restriction concerning the set of non-MI players is that it is *finite*; this restriction will be justified at the end of this section.

The *predictive success rate* of a method P is defined by means of the following chain of definitions:

- $\text{pred}_n(P)$ is the prediction of *player* P for time n which is delivered *at* time $n - 1$ (like the events, predictions are coded by real numbers between 0 and 1).
- The deviation of the prediction pred_n from the event e_n is measured by a normalized loss function $\text{loss}(\text{pred}_n, e_n)$ ranging between 0 and 1.
- The *natural* loss-function is defined as the absolute (linear) distance between prediction and event, $|\text{pred}_n - e_n|$; however, our results apply to a much larger class of loss functions (see below).
- $\text{score}(\text{pred}_n, e_n) =_{\text{def}} 1 - \text{loss}(\text{pred}_n, e_n)$ is the *score* obtained by prediction pred_n of event e_n (ranging between 0 and 1).
- $\text{abs}_n(P) =_{\text{def}} \sum_{1 \leq i \leq n} \text{score}(\text{pred}_i(P), e_i)$ is the *absolute* success achieved by player P until time n (ranging between 0 and n).
- $\text{suc}_n(P) =_{\text{def}} \text{abs}_n(P)/n$ is the *success rate* of player P at time n (ranging between 0 and 1).

The optimality theorem below holds for all *convex* loss functions, which means that the loss of a weighted average of two predictions is not greater than the weighted average of the losses of two predictions. In what follows we assume convex loss functions; they comprise a large variety of loss functions including all linear, polynomial and exponential functions of the natural loss function.

The simplest meta-inductive strategy is called *Imitate-the-best* and predicts what the presently best non-MI player predicts. It is easy to see that this meta-inductive strategy cannot be universally access-optimal: its success rate breaks down when it plays against non-MI methods that are *deceivers*, which means that they lower their success rate as soon as their predictions are imitated by the meta-inductivist (cf. [Schurz 2008b](#), Sect. 4). A realistic example is the prediction of stock values in a ‘bubble economy’: Here the prediction that a given stock will yield a high rate of return leads many investors to put their money on this stock and by doing so they cause it to crash.

Nevertheless there exists a meta-inductive strategy that is provably universally optimal. This strategy is called *attractivity-weighted meta-induction*, abbreviated as wMI. This method predicts a weighted average of the predictions of the non-MI players, using their so-called “attractivities” as weights. The attractivity of a player P (at a given time) is the surplus of P’s success rate in relation to wMI’s success rate. From the viewpoint of wMI, this attractivity is called *regret* and attractivity-based meta-induction is a variant of regret-based learning (Cesa-Bianchi and Lugosi 2006).

(5) Predictions of wMI (attractivity-weighted meta-induction):

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

- $\text{at}_n(P_i)$ is the attractivity of a player P_i for wMI at time n , defined as $\text{at}_n(P_i) =_{\text{def}} \text{suc}_n(P_i) - \text{suc}_n(\text{wMI})$, if this expression is positive; else $\text{at}_n(P_i) = 0$, and
- if $n=1$ or the denominator is zero, wMI’s predicts by her fallback method.

Observe that a player’s attractivity is set to zero if his success rate is lower than that of wMI. This is a necessary condition for wMI’s access-optimality: it guarantees that wMI’s success approximates the success rate of the best non-MI player, because as soon as wMI’s success surpasses the success rate of some non-best player, wMI ignores its predictions. However, wMI is even access-optimal when there is no best non-MI player, but the success rates of the non-MI players are endlessly oscillating around each other. More generally, let “ maxsuc_n ” denote the non-MI-players’ maximal success rate at time n . Then the following universal optimality theorem for wMI has been proved:⁵

(6) *Theorem:* (universal access-optimality of wMI):

For every prediction game ((e), $\{P_1, \dots, P_m, \text{wMI}\}$) the following holds:

- (i) (Short run:) ($\forall n \geq 1$): $\text{suc}_n(\text{wMI}) \geq \text{maxsuc}_n - \sqrt{m/n}$.
- (ii) (Long-run:) $\text{suc}_n(\text{wMI})$ converges to the non-MI-players’ maximal success for $n \rightarrow \infty$.

According to theorem (6)(ii) attractivity-weighted meta-induction is long-run optimal for *all* possible event sequences and finite sets of (simultaneously accessible) prediction methods. In the short run, weighted meta-induction may suffer from a possible loss, compared to the leading player. This loss is caused by the fact that wMI must base her prediction of the next event on the *past* success rates of the candidate methods, and the hitherto most attractive methods may perform badly in the prediction of the next event. Fortunately theorem (6)(i) states a worst-case upper bound for this loss, which is small if the number of competing methods (m) is small compared to the number of rounds (n) and converges to zero when n grows large.

Theorem (6) applies to prediction games with real-valued as well as binary (or discrete) events. Even if the events are binary, wMI’s predictions are real-valued, because proper weighted averages of 0s and 1s are real-valued. How can the optimality

⁵ The proof is given in Schurz (2008b, Sect. 7, Theorem 4), based on results in regret-based learning theory (Cesa-Bianchi and Lugosi 2006).

result of theorem (6) be transferred to *binary* games whose predictions must be binary? There are two methods by which this can be done:

- (1) *Randomization*: Here one assumes that the meta-inductivist predicts $e_n = 1$ with a probability that equals the optimal real-valued prediction of wMI (Cesa-Bianchi and Lugosi 2006, Sect. 4.1).
- (2) *Collective meta-induction*: Here a *collective* of meta-inductivists approximates real-valued predictions by the mean value of their binary predictions (Schurz 2008b, Sect. 8).

Theorem (6) establishes the following

(7) *A-priori justification of attractivity-weighted meta-induction*: In all possible worlds it is reasonable for the given epistemic subject X to apply the strategy wMI to all prediction methods accessible to X, since this can only improve but not worsen X's success in the long run.

Claim (7) should not be misunderstood as entailing that the application of wMI *in isolation* is the best epistemic strategy. Rather, it implies that wMI is optimal *ceteris paribus*, conditional on a given candidate set of object-level methods. Besides this, it is always reasonable *in addition* to try to improve one's candidate set. But this does not constitute an objection against the universal recommendability of applying wMI *on top* of one's candidate set.

The given justification of meta-induction is a priori and analytic, insofar it does not make any assumptions about contingent facts. Moreover, the justification is non-circular, because it does not rest on any inductive inference or assumption of inductive uniformity. The only assumptions on which the optimality justification of meta-induction rests are

- (i) that the decision-maker has the required ('normal') cognitive capabilities for computing wMI's predictions from her past observations, and
- (ii) that past observations/experiences can be reliably recorded.

Assumption (i) is epistemologically harmless. Assumption (ii) implies that the subject's beliefs concerning her *past* observations are justified. This is not epistemologically harmless, but uncritical, because this assumption does not presuppose induction and, thus, preserves the non-circularity of the optimality justification.

Theorem (6) asserts the optimality but not the dominance of attractivity-based meta-induction. Thus there may exist other meta-level methods, different from wMI, that are likewise access-optimal in the long run. For example, one can show that certain variants of wMI are long-run optimal as well and have short-run advantages in certain and disadvantages in other environments. In other words, meta-induction is universally optimal, but not universally dominant. However, as Reichenbach (1949, p. 475f) has pointed out, the optimality argument may nevertheless be considered as a sufficiently strong justification of meta-induction, insofar as meta-inductive strategies are the *only* prediction meta-strategies for which optimality can be rationally demonstrated.

By itself this justification does not entail anything about the rationality of object-level induction: it may well be that we live in a world in which a method different

from object-induction (e.g., clairvoyance) is predictively superior. However, the a priori justification of meta-induction gives us the following

(8) *A-posteriori justification of object-induction*: As a matter of fact, object-inductive prediction methods were so far much more successful than all accessible non-inductive (object-level) prediction methods. Therefore it is justified, by meta-induction, to continue favoring object-inductive prediction methods in the future.

Argument (8) is no longer circular, because a non-circular justification of meta-induction has been established independently. The argument presupposes a contingent premise about the past success rates of inductive compared to non-inductive prediction methods. Given that all scientific prediction methods are based on object-induction I think this premise is eminently plausible (although, following from its nature as a contingent premise, there is always room for debate). Observe that argument (8) only infers that object-inductive methods are better justified than *non-inductive* methods; this respects the fact that there exist several different object-inductive prediction methods whose success rate may differ in different environments. This fact constitutes the advantage of wMI in inductively uniform worlds, since wMI favors a particular object-inductive method in proportion to its attractiveness in the given environment.

We finally turn to the restriction of theorem (6) to *finitely many* methods. This restriction is a necessary condition for the proof of the *universal* access-optimality of meta-induction (without it only weaker results are provable). The finiteness restriction can be justified by the following

(9) *Argument from cognitive finiteness*: Epistemic subjects are assumed to be finite beings. Finite beings can simultaneously access (and compare) only finitely many methods of finite complexity. Therefore the optimality justification of meta-induction is not affected by the finiteness restriction.

Not all philosophers will be satisfied by this justification. They may object that human beings can represent infinite sets. However, if human beings do this, they represent infinities always by *finite* representations. Another objection may point out that the restriction to finitely many methods is inadequate because the problem of induction has to do with infinities. However, there is a world of difference between the problem of infinities at the level of competing hypotheses over infinite domains and at the level of methods. Methods are much more general than hypotheses in two respects: (i) Infinitely many different hypotheses can be generated by one and the same method if it is applied to infinitely many different event sequences. (ii) While in the situation of choosing among a finite number of competing hypothesis one may hope to get to the truth via the elimination of falsified or statistically refuted hypotheses, it is impossible to discover the ‘true method’ by elimination at the level of methods, because methods cannot be ‘falsified’ or ‘statistically refuted’.

In any case, the problem of choosing among finitely many competing methods captures the most important part of the induction problem. In all real-life decision problems one is confronted with competitions between *finitely many* methods. Without a non-circular solution to Hume’s problem we wouldn’t even be able to defend induction against a *single* competitor, e.g., against reading the future from a deck of cards, or against anti-inductive stubbornness (recall Sect. 3.2). It seems fair to conclude

that the optimality justification of meta-induction gives us an at least partial solution to Hume's problem of induction. The *core* of this solution consists in the fact that meta-induction has an unlimited learning ability: whenever this strategy is confronted with a so far better method, it will learn from it and reproduce its success. This is what makes it optimal—not among all possible but among all accessible prediction methods.

5 Generalizing optimality justifications

Our results about meta-induction bestows us a new insight for foundation-oriented epistemology: It is possible to stop the justificational regress at the level of higher order justifications by demonstrating that an epistemic inference strategy is universally access-optimal and, thus, is justified independently from further contingent assumptions. In the preceding section it was shown that attractivity-weighted meta-induction is such a strategy. In this section I give a brief explanation of how the method of optimality justifications can be generalized to two further problems of foundation-oriented epistemology: the 2nd order justification of classical logic (in comparison to non-classical logics) and the 2nd order justification of abduction (in comparison to epistemologies that reject this inference).

5.1 The justification of (classical) logic

In Sect. 2 we argued that *prima facie* the justification of (classical) deductive inferences is unproblematic, because one can prove semantically that these inferences are strictly truth-preserving. This is true, but for the proof of this semantic fact one needs again the principles of classical logic, now stated within the meta-language in which the semantic rules are expressed. For example, the semantic proof of the truth-preserving nature of the simplification rule " $p \wedge q/p$ " goes as follows: (i) $\text{True}(p \wedge q)$ implies (ii) $\text{True}(p) \wedge \text{True}(q)$ (by the definition of \wedge 's truth-table) which implies (iii) $\text{True}(p)$ by the simplification rule. Thus, to prove the simplification rule in the object language we need the simplification rule in the meta-language. Does this mean that we are again in the threatening situation of an epistemic circle or infinite regress? *No*, it merely means that semantic explications, although philosophically insightful, cannot stop the justificational regress. At some meta-language level we must stop the regress by assuming the principles of classical logic as *given*, i.e., as *basic* in the explained sense. Technically this is done by assuming an *axiomatic* system, i.e., a system of axioms and rules from which (hopefully) all other logically valid theorems can be derived.

What justifies us in considering the principles of classical logic as basic? The traditional answer to this question points to the fact that there is a crucial difference between the problem of justifying induction and that of justifying deduction: While we can easily imagine possible worlds in which induction fails (whence induction needs justification), we can hardly imagine possible worlds in which logic fails, because we presuppose our logic already in the representation of these worlds. For this reason, deductive logic is basic and needs no justification.

Unfortunately this justification is not fully convincing, because it presupposes that possible worlds are represented by means of classical logic. However, there are alternatives to classical logic: *non-classical* logics do not make the classical assumptions but assume, for example, more than two truth values, e.g. “true”, “false”, and “undetermined”. How can one justify classical logic, or a system of logic at all, in view of this situation of ‘logical pluralism’?

The situation may seem hopeless, but in fact it is not, since logical systems are *translatable* into each other. For example, a three-valued non-classical logic may be translated into a two-valued classical logic by introducing three additional concepts into the language of classical logic: the propositional operators of “being true” (T), “being false” (F) and “being undetermined” (U). If S is a sentence of the three-valued logic, then the sentences T(S), F(S) and U(S) are nevertheless two-valued. Based on this fact, every semantic axiom or rule of a three-valued non-classical logic can be translated into a corresponding axiom or rule formulated in the expanded language of the classical two-valued logic. For example, Lukasiewicz’ three-valued truth table for negation is represented by the three semantic axioms $T(\neg S) \leftrightarrow F(S)$, $U(\neg S) \leftrightarrow U(S)$ and $F(\neg S) \leftrightarrow T(S)$. By representing all truth tables of Lukasiewicz’ three-valued logic via semantic axioms of this kind and adding the axiom $T(S) \dot{\vee} U(S) \dot{\vee} F(S)$ (with “ $\dot{\vee}$ ” for exclusive disjunction), we obtain the axiom system Ax_{Luk} of Lukasiewicz’ logic in the language of classical logic. Now each sentence of the three-valued logic S can be translated into the corresponding sentence T(S) of classical logic so that the property of validity is preserved, i.e., a sentence is logically true in the three-valued logic exactly if its translation is logically true in the corresponding axiomatic system in classical logics: $\models_{Luk} S$ iff $Ax_{Luk} \models_{class} T(S)$.⁶

The same translation strategy applies to all many-valued logics representable by means of finitely many truth values. For example, many *para-consistent* logics can be characterized by means of finite truth value matrices, including truth-values such as “both true and false” (Priest 1979, 2013, Sect. 3.6). A detailed elaboration of this idea is work for the future.

I conjecture that a similar translation strategy applies to all kinds of non-classical logics (even those not characterizable by finite matrices). My reason for this conjecture is that all non-classical logics known to me use classical logic in their meta-language in which they describe the semantics of their non-classical principles. Therefore there must exist ways to translate the principles of these logics into classical logic, by introducing additional operators into the language of classical logic corresponding to the semantical concepts of the non-classical logic (e.g., non-standard truth values in the case of many-valued logic).

What this argument would show, if it is correct, is that every non-classical logic can be represented within classical logic, by using an appropriate extension of the language. This would give us an optimality justification: By using classical logic our conceptual representation system can only gain but can never lose, because if another logic turns out to have advantages for certain purposes, we can translate and thus embed it into classical logic. This optimality justification does not make any presuppositions,

⁶ A different translation is proposed by Rutz (1972): He translates sentences of the three-valued logic into n-tuples of sentences of the two-valued logic.

except the existence of the two logical systems (the translation function is defined in the classical meta-language, which is of the same conceptual type as the object-language). Of course, the argument only shows that classical logic is optimal, but not that it is ‘dominant’ in the sense of being better than, say, three-valued logics. For example, it can be shown that classical logic is also translatable into three-valued logics (cf. [Rutz 1972](#)). So the defender of a three-valued logic can argue that her system is optimal, too, since she can translate every bivalent system into her three-valued logic. But this fact does not undermine the force of the optimality justification; it merely draws the picture of a situation of logical pluralism in harmony, since the alternative logical systems are inter-translatable. Still, one may prefer classical logics as psychologically more natural since they fit better with the way our mind or brain is working, but this is a different matter that will not be pursued here.

5.2 The justification of abductive inference

Abductive inference (or inference to the best explanation) becomes epistemologically indispensable as soon as the conclusion contains *theoretical* concepts that are not contained in the premises and represent unobservable (or ‘hidden’) parameters. The observation of the success records of the empirical predictions of a theory doesn’t give us any direct feedback about the fit of a theory’s theoretical structure with the unobservable structure in the theory’s domain. Such a direct feedback does not exist, because theoretical parameters are unobservable. Therefore the method of meta-induction cannot be directly applied to the abductive inference from the empirical adequacy to the truthlikeness of (the unobservable part of) a theory. The same diagnosis applies to the abductive inference from the regularities in our introspective experiences (which are inductively generalized) to the existence of an external reality with certain properties that figure as best explanation of these introspectively experienced regularities.

Schurz (2008a, Sect. 7.4) analyses the inference to external reality as a *common cause abduction*. The hypothesis of external objects provides a common cause explanation of a huge set of inter-correlations between our introspective experiences. *First*, there are the *intra*-sensual correlations, in particular those within our system of visual perceptions: There are potentially infinitely many two-dimensional visual images of the same perceptual object on the retina, but all these 2D images are strictly correlated with the position and angle at which we look at 3D objects; so these correlations have a common cause explanation in terms of external objects in a 3D space. *Second*, there are the *inter*-sensual correlations between different sensual experiences, in particular between visual perceptions and tactile perceptions, which are similarly explained by the assumption of external objects in a 3D space.

The question arises how the cognitive optimality of abductive inferences can be justified in a non-circular way. I see two ways of doing this, an *instrumentalist* and a *realist* way. Instrumentalistically we can argue that by performing abductive inferences we always take the advantage of explaining and representing our system of experiences by the best available theoretical model, i.e., by the most *simple* and most *unified* theory. Although this justification is instrumental, it goes beyond mere consideration of predictive success and considers matters of unification and economy which belong to the

dimension of cognitive costs. More precisely, the instrumental optimality justification in terms of cognitive success works as follows: Should some part of our theoretical model be false, one of two cases may obtain. *Either* we observe this in the form of an incorrect prediction; as soon as this happens we will take steps to correct our theory. In other words, abductive inferences are self-corrective and have an inbuilt learning ability. *Or* we never observe it (because our experiences are limited); then nothing happens and we continue to operate with an instrumentally optimal theory, although it is false, but in a way that cannot be empirically detected by us and, thus, will not practically harm us. Thus by performing abductive inferences to unifying theoretical models we can only gain but not lose something. I conjecture that even empiricists such as van Fraassen who reject standard accounts of IBE (cf. 1989, p. 142ff) would accept this instrumentalistic justification.

Can more than such an instrumentalist justification be given—a justification that directly infers the realistic truthlikeness of the theoretical part of an empirically successful theory? A naive argument of this sort is Putnam’s *no miracle* argument (Putnam 1975, p. 73). It argues, roughly speaking, that without the assumption of realistic truthlikeness the empirical success of science would be a sheer miracle. There are broad controversies about this argument, including Laudan’s “pessimistic meta-induction” (1981), and this is not the place to enter this debate.

In Schurz (2016, Sect. 3) it is argued that the strategy of common cause abduction can be justified in a realistic manner by assuming the principles of Markov causality. These principles imply that an observed correlation between events or dispositions that are not related as cause and effects must be produced by (unobserved) common causes. However, this realist justification of abduction rests on a certain amount of causality principles, which have in turn to be justified by means of an instrumental abduction terms of unification (cf. Schurz and Gebharder 2016). I regard it as an open question whether a non-circular optimality justification of the abductive inference to reality can be given that is stronger than an instrumentalistic justification in terms of predictive success and cognitive economy.

6 Conclusion

This concludes my brief sketch of the application of optimality justifications to other domains of foundation-oriented epistemology. Generally speaking, optimality justifications constitute new foundations for foundation-oriented epistemology. Let me finally try to locate, in a preliminary way, the place of the account of optimality justifications in the landscape of epistemological positions in the history of enlightenment philosophy. For this purpose, the philosopher Immanuel Kant shall figure as my lighthouse. Certainly the epistemic optimality strategy does not belong to pre-Kantian metaphysical accounts that were based on uncritically accepted premises, which later on turned out to be unjustified by the skeptical challenges of empirical scientists and philosophical empiricists, in particular by those of David Hume. What our account shares with the Kantian philosophy is the ‘Copernican’ turn towards the inner ‘transcendental’ dimension of knowledge, the question of its ultimate cognitive foundations, presuppositions and justifications. In contrast to Kant, however, we nei-

ther assume nor argue that certain cognitive methods or principles are a priori, in the sense that we must apply them as necessary presuppositions of cognition. Kantian a-priorism is not tenable and modern philosophy has shown time and again that no transcendental argument can prove the a priori validity or necessity of a cognitive method or principle. Even at the most fundamental level, there are *choices*: there is more than one method and more than one way to go. However, what one can still have in this situation of a ‘foundational pluralism’ are optimality justifications by means of strategies that are universally access-optimal because of their inbuilt learning capacities. This is the central innovation of the proposed account of optimality justification. In conclusion, if Kant’s philosophy is called *transcendental a-priorism*, then the account proposed in this paper can be called *transcendental optimalism*.

Acknowledgements This work was supported by the DFG (Deutsche Forschungsgemeinschaft), SPP 1516.

References

- Alston, W. P. (1976). Two types of foundationalism. *Journal of Philosophy*, 73(7), 165–185.
- Black, M. (1974). Self-supporting inductive arguments. in: Swinburne (Ed.), (pp. 127–134).
- Braithwaite, R. B. (1974). The predictionist justification of induction. In Swinburne (Ed.), (pp. 102–126).
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge: Cambridge Univ. Press.
- Dancy, J. (1985). *An introduction to contemporary epistemology*. Oxford: B. Blackwell.
- Fumerton, R. (1995). *Metaepistemology and skepticism*. London: Rowman & Littlefield Publ.
- Goldman, A. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard Univ. Press.
- Goldman, A. (1999). *Knowledge in a social world*. Oxford: Oxford Univ. Press.
- Hume, D. (1748). *An inquiry concerning human understanding* (ed. by S. Butler). Fairford: Echo Library.
- Kelly, K. T. (1996). *The logic of reliable inquiry*. New York: Oxford Univ. Press.
- Ladyman, J., & Ross, D. (2007). *Every thing must go. Metaphysics naturalized*. Oxford: Oxford University Press (With D. Spurrett and J. Collier).
- Lehrer, K. (1990). *Theory of knowledge*. Boulder: Westview Press (1st edn. 1974 as *Knowledge*).
- Papineau, D. (1993). *Philosophical naturalism*. Oxford: B. Blackwell.
- Pollock, J. (1986). *Contemporary theories of knowledge*. Maryland: Rowman & Littlefield.
- Priest, G. (1979). Logic of paradox. *Journal of Philosophical Logic*, 8, 219–241.
- Priest, G. (2013). Paraconsistent logic. In E. Zalta, E. (Ed.), *Stanford encyclopedia of philosophy*. <http://plato.stanford.edu>.
- Putnam, H. (1975). What is mathematical truth? In H. Putnam (Ed.), *Mathematics, matter and method* (pp. 60–78). Cambridge: Cambridge Univ. Press.
- Reichenbach, H. (1949). *The theory of probability*. Berkeley: University of California Press.
- Russell, B. (1946). *History of Western philosophy*. London: George Allen and Unwin.
- Rutz, P. (1972). *Zweiwertige und mehrwertige Logik*. Ehrenwirth Verlag München.
- Salmon, W. C. (1957). Should we attempt to justify induction? *Philosophical Studies*, 8(3), 45–47.
- Schurz, G. (1991). Relevant deduction. *Erkenntnis*, 35, 391–437.
- Schurz, G. (2008a). Patterns of abduction. *Synthese*, 164, 201–234.
- Schurz, G. (2008b). 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. (2008). Third-person internalism: A critical examination of externalism and a foundation-oriented alternative. *Acta Analytica*, 23(2008), 9–28.
- Schurz, G. (2009). Meta-induction and social epistemology. *Episteme*, 6, 200–220.
- Schurz, G. (2016). Common cause abduction: The formation of theoretical concepts and models in science. *Logic Journal of the IGPL*, 24(4), 494–509.
- Schurz, G., & Thorn, P. (2016). The revenge of ecological rationality: Strategy-selection by meta-induction within changing environments. *Minds and Machines*, 26(1), 31–59.

- Schurz, G., & Gebharter, A. (2016). Causality as a theoretical concept: Explanatory warrant and empirical content of the theory of causal nets. *Synthese*, 193(4), 1071–1103.
- Shogenji, T. (2007). Internalism and externalism in meliorative epistemology. *Erkenntnis*, 76(1), 59–72.
- Skyrms, B. (1975). *Choice and chance*, Dickenson, Encinco (4th ed. Wadsworth 2000).
- Sosa, E. (2010). The epistemology of disagreement. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford: Oxford University Press.
- Sterkenburg, T. (2017). Putnam's diagonal argument and the impossibility of a universal learning machine, pre-print, philsci-archive.pitt.edu, Pittsburgh.
- Swinburne, R. (1974). *The justification of induction*. Oxford: Oxford University Press.
- Van Cleve, J. (1984). Reliability, justification, and induction. In P. A. French, et al. (Eds.), *Causation and causal theories* (Vol. 4, pp. 555–567). Midwest Studies in Philosophy.
- Van Fraassen, B. (1989). *Laws and asymmetry*. Oxford: Clarendon Press.
- White, R. (2015). The problem of the problem of induction. *Episteme*, 12(2), 275–290.