



Conceptual Spaces: A Solution to Goodman's New Riddle of Induction?

Sebastian Scholz¹

Received: 30 November 2023 / Revised: 23 April 2024 / Accepted: 8 May 2024
© The Author(s) 2024

Abstract

Nelson Goodman observed that we use only certain ‘good’ (viz. projectible) predicates during reasoning, with no obvious demarcation criterion in sight to distinguish them from the bad and gruesome ones. This apparent arbitrariness undermines the justifiability of our reasoning practices. Inspired by Quine’s 1969 paper on Natural Kinds, Peter Gärdenfors proposes a cognitive criterion based on his theory of Conceptual Spaces (CS). He argues the good predicates are those referring to natural concepts, and that we can capture naturalness in terms of similarity. In contrast to Quine, he does not primarily rely on logic, but geometry. He frames his account as a descriptive project, however, and it is not obvious how it addresses the normative dimension of Goodman’s Riddle. This paper develops a charitable reconstruction of his argument, based on the idea that the instrumental success of our projectible concepts is grounded in their cognitive-pragmatic naturalness. It then explores three lines of reasoning against the argument: Evolutionarily motivated skepticism, the miracles argument, and the relation between instrumental and pragmatic success. I conclude that in its current form, the CS account fails to provide any justification of reasoning beyond appealing to its instrumental success, and that a metaphysically robust notion of naturalness helps to achieve the desired goal.

Keywords Concepts · Induction · Naturalness · Reasoning · Scientific realism

1 Introduction

Philosophy is riddled with perplexing conundrums, and Goodman’s Riddle is a particularly intriguing specimen. Goodman had noticed that we use only certain ‘good’ (viz. *projectible*) predicates during reasoning, with no obvious demarcation criterion

✉ Sebastian Scholz
S.Scholz@hhu.de

¹ Department of Philosophy, Heinrich-Heine-University Düsseldorf, Düsseldorf, Germany

in sight to distinguish them from the bad and gruesome ones. This apparent arbitrariness undermines the justifiability of our reasoning practices – which feels important, as reasoning is generally seen as a big part of what makes us human.

Since Goodman posed his riddle in 1955, many solutions have been advanced, but the issue has proven persistent. A notably inventive approach comes from recent cognitive research. Inspired by W. O. Quine's, 1969 paper on *Natural Kinds*, Peter Gärdenfors argues the good predicates are those referring to natural concepts, and that we can capture naturalness in terms of similarity. In contrast to Quine, however, he does not primarily rely on logic, but *geometry*. His innovation is to represent similarity as a function of distance in so-called Conceptual Spaces (CS). As similar objects form clusters in these peculiar theoretical entities, properties and concepts can be defined as spatial regions – and regarded as *natural* if they display a degree of topological cohesion, as captured by the geometric property of *convexity*. The resulting demarcation criterion is praised by the author as a comprehensive solution to the riddle. But does the approach deliver what is promised? Depending on what we demand of a solution, it might not.

As a cognitive scientist, Gärdenfors is primarily concerned with explaining and predicting the behavior of cognitive agents, as well as constructing artificial agents – so he frames his account as a descriptive project. Goodman's Riddle, however, clearly has a normative dimension: It is about the *justification* of reasoning. This raises the questions of whether and how the CS account addresses the normative dimension of the riddle. In what follows, I pursue two goals. First, I will develop a charitable reconstruction of Gärdenfors' argument and thus constructively contribute to the improvement of the approach. I think the approach *does* have a clear position on the justification issue, but it is somewhat difficult to pin down and should be made more explicit. Second, I will explore three lines of reasoning that might be raised against the argument, two of which I maintain are sound. My conclusion is the approach *does not* provide justification – beyond appealing to the instrumental success of inference, which already Hume or Goodman did. In a nutshell, this is because Gärdenfors relies on a cognitive notion of naturalness, although we need ontological naturalness to make sense of the (presumed) fact that geometrically cohesive concepts are instrumentally successful in reasoning. My argument is not about convexity specifically, but it targets cognitive criteria of naturalness in general. I advocate for a conception of naturalness that takes both cognitive and ontological aspects into account.

The next sections provide background on Goodman's Riddle and what we might count as a solution. I will then introduce the CS account and work out the solution it proposes, before developing my own argument in three steps.

2 The Riddle

Goodman's New Riddle of Induction is a successor of – though logically independent from – Hume's old Problem of Induction. Hume had famously argued we base our inductive generalizations not on reason, but habit. Whenever we try coming up with reasons to justify induction, circularity looms, as we have to presuppose the unifor-

mity of past and future cases we are trying to establish.¹ Goodman extended the scope of the problem by noting that “some regularities do and some do not establish such habits” (Goodman, 1983: 82). Say, we find a bunch of green emeralds inductively confirming the hypothesis that *all emeralds are green*. This regularity and generalization habit are familiar and very intuitive. The trouble is we can formulate an alternative hypothesis, namely that *all emeralds are grue* – where *grue* is defined to be satisfied by the *very same* emeralds as follows: *green and observed before t, or blue and unobserved before t*. What we have before us, then, are actually at least two regularities and generalization habits. Crucially, both habits have equal evidential support, yet we prefer one over the other. How could we ever justify settling on a specific hypothesis, if we can always come up with alternative hypotheses that have equal degree of confirmation? This is the heart of the riddle and the very reason why it had such a big impact in the philosophy of science. Goodman had informally proved that the standard *instantial model of confirmation*, according to which “positive instance of a generalization lends some support to that generalization [...]” (Slater & Borghini, 2013: 5) is at best incomplete.

The property of hypotheses to be confirmable by their instances, or of predicates to be transposable from one set of cases to another, is what Goodman calls “projectibility”. “Green” is supposed to be projectible, while “grue” is not. As I take reasoning to be a cognitive rather than purely linguistic enterprise, I will sometimes speak more loosely of projectible concepts, but the idea is the same.²

2.1 What Counts as Solution?

Let us now sharpen our intuitions on what we might count as a solution. It is quite clear that a mere description of our inferential practices is not what we are after. Nobody doubts that, as a matter of fact, we make certain inductions, or project certain predicates. If, by contrast, somebody came up with robust, non-circular reasoning to justify our practices – viz. the ‘strong’ justification Hume was after – then this would quite certainly constitute a solution. There is, however, a lot of middle ground to cover, and the boundaries between description and prescription are sometimes muddy. Consider the following passage about Goodman’s own take on the issue:

Goodman’s solution to the new riddle of induction resembles Hume’s solution in an important way. Instead of providing a theory that would ultimately justify our choice of predicates for induction, he develops a theory that provides an account of how we in fact choose predicates for induction and projection. [...] [He] makes projectibility essentially a matter of what language we use and have used to describe and predict the behaviour of our world. (Cohnitz & Rossberg, 2020: 5.4).

¹ Hume’s argument is more intricate, of course; it takes the form of a dilemma, and circularity is the crux of one of its horns. There are also interpretations that do not make reference to the aforementioned uniformity principle; here, however, the condensed portrayal suffices. Cf. Henderson, 2020.

² I understand concepts here as mental representations that are postulated within cognitive psychology to explain higher cognitive competences such as categorization and reasoning.

First, note that the authors regard both Hume's and Goodman's approaches as solutions. Second, the quote could be read as if both philosophers deemed the normative dimensions of their riddles entirely irresolvable. But things are not as straightforward: Just because there is no "ultimate" justification does not mean there is none at all. Hume, for instance, was historically interpreted as a skeptic about the possibility of justification. He introduced his notion of habit as a descriptive explanation of our inductive practices, which is more than a mere description, but does not by itself address the normative question. And yet, not even Hume recommended against relying on induction in day-to-day life – so he had to think some 'weak' form of justification was available, even if just in terms of practical, everyday *success* (cf. Henderson, 2020: 5.1).

With Goodman, the idea that some of our projections are successful even takes center stage. His own answer, which is subject to the above quote, is *entrenchment*: As some of our past projections happened to be successful, the predicates and hypotheses that were involved stuck around and became entrenched in language. Or as Quine puts it: "In induction nothing succeeds like success." (Quine, 1969: 129). The notion of success provides weak justification and makes normativity part of the equation, so to speak. One should not lose sight of the fact, however, that the validity of the norms determining what counts as success cannot be derived from the description of our reasoning practices alone. Instead, for Goodman, practice itself plays a role in this, resulting in a complex overall picture that is reminiscent of the late Wittgenstein:

What we have in Goodman's view, as, perhaps, in Wittgenstein's, are practices, which are right or wrong depending on how they square with our standards. And our Standards are right or wrong depending on how they square with our practices. This is a circle, or better a spiral, but one that Goodman [...] regards as virtuous. (Putnam, 1983: ix).

Neither the details of Goodman's approach, nor the extensive criticisms that have been leveled at it, are of concern here. What all this goes to show is the title question of this paper does not have a categorical answer – we must conditionalize on what we demand of a solution. It is reasonable to require that it should somehow address the normative concern at the heart of the riddle, but if strong justification is believed to be unavailable, then finding necessary and sufficient conditions of projectibility that capture our practices may suffice. Much of the literature on projectibility can be read along these lines,³ and Gärdenfors' convexity criterion, too, could be interpreted accordingly. I do not think, however, that this would do justice to the approach, as it tries to achieve more by invoking *naturalness*.

³ Earman, 1985 is a good example, who formally distinguishes several problems of induction (only some of which involve gruesome predicates) and discusses candidate necessary and sufficient criteria.

2.2 The Connection with Naturalness

Gruesome predicates are weird and unwelcome. Philosophers like to label them “unnatural”, “miscellaneous”, or “gerrymandered”, and have tried casting them out to prevent them from causing trouble (this is of course a caricature). There is no straightforward way to achieve this, however, because as Gärdenfors, 1990 points out, from a purely logical standpoint, a predicate is a predicate is a predicate. Appealing to the simplicity of the more welcome ones does not help either, because what counts as simple depends on the language in which the predicates are defined. One obvious idea is to appeal to the referents of the predicates instead. Suitable candidates include properties or kinds – however, there may be many more of those around than there are projectible predicates (cf. Lewis, 1983). According to *class nominalism*, for instance, properties are sets of individuals, so every random (or gruesome) combination of things will count as a property – which certainly doesn’t help with the riddle. Platonists, by contrast, can at least in theory restrain the number of referents by appealing to *universals* (i.e., multiply instantiable entities). But anyone who wants to be metaphysically thriftier, or who looks for some practically applicable and not just theoretical demarcation criterion, is bound to have a hard time sorting this out – as does W. O. Quine, whose austere naturalistic ontology refrains from universals and *abstract objects* (i.e., spatiotemporally distributed entities), but allows for physical objects and sets because they are indispensable for science and mathematics (cf. Hylton & Kemp, 2023).

In his influential paper on natural kinds, Quine attempts demarcation via *similarity*, the idea being that “[t]wo green emeralds are more similar than two grue ones would be if only one of the grue ones were green.” (Quine, 1969: 116). Similarity, in turn, he takes to be closely related to the notion of natural kinds, even going as far as to say they are “substantially one notion” (Quine, 1969: 119). In other words, he wants to delineate among all the groupings of things an elite set of *natural* groupings of *similar* things which are depicted by the projectible predicates. Traditionally, however, naturalness is about the world having an *objective* or *mind-independent* structure for us to “carve” with our everyday or scientific “conceptual cutlery” (Slater & Borghini, 2013: 25). Thus, natural entities are usually interpreted realistically: They have been conceived of as Platonic universals, Aristotelian essences, homeostatic property clusters (Boyd, 1999), perfectly natural properties (Lewis, 1983, 1986), and many more. But due to his austere naturalism – plus his pragmatic leanings, which we will hear more about later – for Quine, these options are (or would be) out of the question. Instead, he attempts analyzing naturalness in terms of similarity, and vice versa, while also relating both concepts to more readily understandable logical or set-theoretical terms. But while he produces many valuable insights, especially into evolutionary aspects of the issue (again, more on that later), he does not find a satisfying analysis, noting that “[...] there is something logically repugnant about [similarity]” (Quine, 1969: 117) – and naturalness, for that matter.

At this point, several remarks on terminology are overdue. First, depending on the philosophical task at hand, cleanly distinguishing between natural kinds, properties, propositions, etc. is mandatory. Here this is less important, as I am primarily concerned with the contrast between *cognitive* and *ontological* naturalness of any

kinds of entities. I will thus sometimes speak interchangeably of natural *concepts* or *representations* to indicate the cognitive notion, versus natural *kinds* or *properties* to indicate the ontological notion. The nature of this distinction will be elaborated in Sect. 3.2. It should be noted that Gärdenfors treats properties as cognitive entities, which is a source of potential terminological confusion that I will do my best to avoid.

Second, I will reserve the term *strong realism* for the ontologically committed view that there are mind-independent entities which are the natural kinds. This is the position I will argue for; it gives the natural entities an *intrinsic* metaphysical status. According to Bird & Tobin, *weak realism* is by contrast the “ontologically uncommitted view that our classifications are often natural.” (Bird & Tobin, 2023: 1.1.1). Yet, in this discussion, I will include positions such as Quine’s, which affirm the reality and objectivity of the kinds, but do not assign them an intrinsic metaphysical status. Quine instead highlights their dependence on our scientific interests. Given that this stance incorporates elements traditionally associated with *conventionalism*, which is “the view natural kinds don’t exist independently of the scientists and others who talk about them” (Bird & Tobin, 2023: 1.1.2), it can be misleading to classify him as a strong realist, despite the technical possibility of doing so.

Now, according to Judith Crane, “[p]hilosophical treatments of natural kinds are embedded in two distinct projects. [...] The kinds studied in the philosophy of science approach are projectible categories that can ground inductive inferences and scientific explanation. The kinds studied in the philosophy of language approach are the referential objects of a special linguistic category—natural kind terms—thought to refer directly.” (Crane, 2021: 12177). Our present concern is the former project, and although Crane does not mention him explicitly, we may think of Quine as one of its key historical figures. Couched in modern terminology, thus, his 1969 paper is about *grounding* projectibility in naturalness – although he was ontologically more parsimonious than many recent grounding enthusiasts, so this is to use ‘grounding’ in a broad sense akin to ‘metaphysical explanation’. In order not to open a big can of worms, I will not get into the details of the grounding debate. However, there is an important line of thought for us to distill here. Grounding reasoning in the structure of the natural world comes down to at least two things: Taking the latter to be more fundamental than the former (cf. Fine, 2001 on the relation between grounding and fundamentality), and, crucially, taking the former to be justifiable insofar as it corresponds to the latter. This second point is perhaps so banal that it is rarely stated explicitly, but it is implicit in commonplace talk about how natural kinds “enable us” (Khalidi, 2013: 72) to make inductive inferences, for an instance. The justification strategy to derive here is that, simply put, *if we can provide a sound metaphysical explanation for why some predicates are projectible, then we have all the more reason to trust them*. Whether this approach can generate anything resembling a ‘strong’ justification is unclear and will highly depend on one’s meta-metaphysical views, but it is certainly more ambitious than the familiar appeal to instrumental success. And I think this is precisely the move Gärdenfors is trying to make along the normative dimension of the riddle.

3 Conceptual Spaces

Gärdenfors' diagnosis of Quine's 'failed' analysis of naturalness is that he relied too heavily on the methodology of the Vienna Circle: "Using logical analysis, the prime tool of positivism, is of no avail [...]. [We] have to go below language." (Gärdenfors 2011: 3). Instead of focusing on the symbolic, or propositional form of inference, he suggests looking into how we represent the conceptual knowledge that is involved. And the format of our conceptual representations, he contends, is *geometric* in nature.⁴ It is this move that allows him to treat similarity as spatial distance, namely as distance in CS.

Conceptual Spaces are cognitive and mathematical theoretical entities. Cognitive, as they are postulated to explain and predict the behavior of cognitive agents, and mathematical, as they are based on axiomatic geometry and usually come with metrics (Euclidean or other, the details do not matter here). Moreover, CS typically come with *quality dimensions*, whose "primary function [...] is to represent various 'qualities' of objects." (Gärdenfors, 2009: 5). Dimensions can be *separable* or *integral*: The visual color domain, which is Gärdenfors' paradigm case, consists of the dimensions *hue*, *saturation*, and *brightness*. Possible or real visual objects, as represented by points in the space, cannot have a value in one of these dimensions without having one in the others, and in this sense the latter are inseparable, viz. integral. With this conceptual apparatus in place, *properties* are defined as regions in CS-domains, *concepts* as regions across multiple domains. The property of being green, for instance, is but a region in the visual domain.

3.1 Proposed Solution

To see how this might help with the riddle, recall Quine's take on projectibility. The upshot was projectible predicates are those referring to natural classes of similar entities. I just outlined how Gärdenfors translates similarity into spatial proximity. As a result, similar objects belonging to a natural class will form clusters in CS, the degree of cohesion of which can be taken to indicate the degree of naturalness of the respective classes. Gärdenfors discusses several topological cohesion criteria, but the one he settles on is *convexity*: A region is convex, iff for any points A and B within the region, there is no point C between A and B that is not in the region as well. The resulting definitions read as follows:⁵

- (P) A *natural property* is a convex region in some domain.
- (C) A *natural concept* is represented as a set of convex regions in a number of domains together with a prominence assignment to the domains and information about how the regions in different domains are correlated.

⁴ It must be noted that the CS approach leaves room for other, symbolic or sub-symbolic representational formats, so it is actually pluralistic about the structure of mental content.

⁵ Gärdenfors employs slightly differing definitions, these are taken from (2008).

To sum up the proposal in a nutshell: Convexity figures as a demarcation criterion for projectible predicates via demarcation of the natural representations which are their referents.

Two things are important to note: First, the connection between convexity and naturalness is meant to be empirical, not analytical. If some other cognitive criterion would fit the data better, Gärdenfors would presumably go with that, without his superordinate account being in jeopardy. Second, convexity is meant to be a necessary, but not sufficient criterion. Being non-black, for instance, is a convex region of the visual space, but does not feel very natural as a property (the example is taken from Hempel's Paradox of Confirmation, which is the other riddle Quine discusses). But for Goodman's case, at least, the criterion seems to get rid of the undesirable predicate: Being green is convex and satisfies the necessary condition of naturalness. Being grue, by contrast, is non-convex in a visual space plus time dimension, as the corresponding region disconnects at time= t . There is some discussion on whether this is convincing (cf. Hernández-Conde, 2017, or Strößner 2022), however, my argument will not depend on the intelligibility of convexity specifically, but apply to cognitive naturalness criteria more generally, so there is no need to engage in it here.

Let me now reconstruct Gärdenfors' argument to the conclusion that convexity in CS is the solution to Goodman's Riddle. As was previously mentioned, CS are largely about explaining and predicting behavior. Since projecting some predicates but not others is part of our 'inferential behavior' as cognitive agents, convexity is supposed to account for that. A genuinely descriptive premise can be derived, analogous to Hume's notion of habit:

P1. Our most projectible concepts are convex (cognitively-descriptively)

⁶Since the issue expressed in P1 is an empirical rather than a conceptual matter, I will not try resolving it here. The same is true of the second premise, which – analogously to Goodman's notion of entrenchment – is about instrumental success:

P2. Convex concepts have been instrumentally most successful in reasoning

⁷This premise is still descriptive, but it makes normativity part of the equation, as success is a normative notion. Why presume Gärdenfors would subscribe to P2? On a general level, he repeatedly emphasizes his affiliation with *scientific instrumentalism*, the view that scientific theories are tools for predicting and controlling experiences rather than true descriptions of reality. But you can also see that by the fact that he wants to apply convexity to build reasoning AI, so he has to think the approach leads to somewhat successful reasoning.

⁶ P1 speaks of 'concepts' instead of 'predicates' to indicate Gärdenfors' focus on cognition over language. Another way to express this would be to say that our most projectible predicates *refer* to convex concepts. Within *cognitive semantics*, which feeds into the theoretical background of the CS account, linguistic expressions are generally thought to refer to cognitive entities.

⁷ If a solution to Hume's Problem of Induction is available, then this premise can be generalized to future cases. This is supposed to be a circumspect reconstruction of Gärdenfors' argument, however, so a less assumptive formulation was chosen.

Note that P2 can be considered an *explanans* for P1 as *explanandum*: It is *because* of their instrumental success that convex concepts are projectible. Now, the way I think CS take things a step further is by invoking naturalness as an *explanans* for P2 as *explanandum*:

P3. P2 holds because of the convex concepts' naturalness.

This is the crucial step described above which amounts to grounding instrumental success in naturalness to achieve better justification: Convex concepts are reliable, because they are natural, and that's why we should trust them. This is Gärdenfors' answer to the normative part of the riddle, which has so far only been implicit in his writings.

The success of this argument will of course heavily depend on how one conceives of naturalness, and it is here where my critique will hopefully have some bite. I also owe some substantiation of the sentiment that Gärdenfors would buy into a version of P3. At any rate, it is required to delve a bit deeper into the metaphysical and epistemological intricacies of the approach.

3.2 Metaphysical and Epistemological Ramifications

As previously stated, most traditional theories of naturalness treat natural entities as mind-independent and objective, i.e., they advocate for some version of strong realism. Gärdenfors, however, attacks mind-independence:

[...] when it comes to inductive *inferences* it is not sufficient that the properties exist out there somewhere, but we need to be able to grasp the natural kinds by our minds. In other words, what is needed to understand induction, as performed by humans, is a *conceptualistic* or *cognitive* analysis of natural properties. (Gärdenfors 2011: 3).

But cognitive analyses of this sort are sometimes explicitly dismissed by contemporary metaphysicians such as David Lewis, whose writings continue to have considerable influence on the naturalness debate:

Nor should it be said [...] that as a contingent psychological fact we turn out to have states whose content involves some properties rather than others, and that is what makes it so that the former properties are more natural. (This would be a psychologicistic theory of naturalness). (Lewis, 1983: 377).

We are thus dealing with two distinct notions of naturalness – a *cognitive* and an *ontological* one. I will take ontological naturalness to mean strong realism, although nominalist accounts could be accommodated as well. The important point here is naturalness depends on the structure of mind-independent reality.⁸ Cognitive natural-

⁸ The criterion of mind-independence is a *proviso* and might turn out to be problematic (cf. Khalidi, 2016). I am confident that there are other (non-cognitive) ways of spelling out ontological naturalness (e.g., in

ness, by contrast, first appears to be a type of full-blown conventionalism. However, as will be explored throughout this section, one can coherently devise a cognitive analysis of naturalness without giving up the notion that the world plays *some* role in shaping our natural concepts, i.e., cognitive naturalness does not necessitate *strong* conventionalism. Such analysis is precisely what Gärdenfors aims for when he says it is not sufficient that the properties exist out there – but does not rule out that they do.

But how exactly is this supposed to work? Does that mean his approach is an intermediate position between conventionalism and realism, and if so, doesn't it run the risk of collapsing into one side or the other? This point is important and deserves some scrutiny. It runs parallel to the debate around *perspectival realism*, where some philosophers try to account for the perspectival nature of, say, how we theorize about kinds, but without giving up on non-perspectival aspects, viz. realism – which has proven difficult to achieve. This is a remarkable cross-connection I can merely allude to here.⁹ To see how the CS account navigates the issue, let us first get a clear picture of its conventionalist leanings.

Criteria (P) and (C) tie naturalness to convexity, i.e., to the structure of our cognitive representations. It is fair to say, then, that this characterization *localizes* naturalness in the heads of cognitive agents. And indeed, Gärdenfors explicitly endorses the affirmative version of Hilary Putnam's famous slogan "[...] meanings just ain't in the head!" (Putnam, 1975a: 227), that is he endorses *semantic internalism*. He does make an adjustment to account for the social dimension of meaning, as indicated by the plural form "heads" (cf. Gärdenfors, 1999a), but this appears to fit neatly with the view that meanings – including those of natural kind terms – are a matter of (social) convention and have nothing to do with mind-independent facts. One of the hallmarks of semantic internalism is the notion that *intensions fix extensions* (cf. Schrenk, 2016: 238–39). Within philosophy of science, this view can amount to the meanings of theoretical terms (including natural kind terms) being wholly determined by their respective background theories, which implies what is known as *Kuhnian relativism*.¹⁰ And indeed, Gärdenfors is compelled to address this very issue:

But is not the choice of a conceptual space arbitrary? Since a conceptual space may seem like a Kuhnian paradigm, aren't we thereby stuck with an unavoidable relativism? After all, anyone can pick her own conceptual space and in this way make her favorite properties come out natural in that space. (Gärdenfors, 2000, Sect. 3.7).

Let us consider for a moment where the quality dimensions of CS come from: They are not given by God, but chosen by scientists who try to model cognition. They are

terms of fundamentality), however, this needs to be done in a different paper.

⁹ There have been attempts to spell out the notion of scientific perspective within the CS framework, e.g. Kaipainen & Hautamäki 2015. The authors speak of perspectives on so-called "ontospaces", but that does not seem to involve a strong metaphysical commitment.

¹⁰ Kuhnian relativism makes truth, knowledge, or meaning relative to Kuhnian *paradigms* (cf. Kusch, 2021). Thomas Kuhn introduced the latter concept in his *Structure of Scientific Revolutions* as "universally recognized scientific achievements that for a time provide model problems and solutions to a community of practitioners." (Kuhn, 2012: *Preface*).

often the result of mathematical regression analyses (e.g. multi-dimensional scaling) being applied to psychological data such as similarity judgements, aiming for a high fit between judgements and spatial distances in the model. If criteria P and C tie naturalness – and thus projectibility – to these models, doesn't that introduce precisely the kind of capriciousness you would expect of epistemic relativism? Gärdenfors speaks of a “[...] *metaproblem* of inductive inferences: What criteria can be used to choose between competing conceptual spaces?” (Gärdenfors, 1990: 94).

This is the nature of the challenge, and it might be answered as follows. First, it should be noted that internalism about linguistic meaning must be distinguished from internalism about mental content. One can consistently tie the meanings of natural kind terms to exclusively their intensions without claiming the same for their mental counterparts – which is precisely the path Gärdenfors seems to take by asking “*Does Semantics Need Reality?*”, and replying:

[...] “not directly.” Once we accept the conceptual structure of an individual as given, the semantic mapping [...] can be described without any recourse to the external world. But a second part of the cognitivist answer is “indirectly,” since the conceptual structure is built up in an individual in interaction with reality. (Gärdenfors, 1999b: 14).

Moreover, recall that convexity is intended only as a necessary naturalness criterion. That is why he can claim his analysis “[...] is compatible with, though does not require, strong metaphysical realism.” (Gärdenfors, 1990: 90).

One may wonder, however, whether this is true and whether he does need some version of strong realism in order to avoid relativism, after all. Partly, this concern is addressed by appealing to coherence (instead of convention) with an “extended net of empirical knowledge” (Mormann 1993: 236) – which certainly makes things less arbitrary, but might still be compatible with relativism. The more important argument is once more inspired by Quine – who has phrased the key idea as poignantly as it gets: “Creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind.” (1969, p. 126). The reason we can trust our inferences is natural selection would have wiped us out a long time ago if we could not. By applying this line of thought to our internal cognitive make-up, Gärdenfors installs evolution as a ‘bridgehead’ between cognition and world. He fills this general point with life by emphasizing the cognitive *efficiency* of convex concepts (cf. Douven & Gärdenfors 2019 for elaboration of this point). With respect to the metaphysics of naturalness, this leaves us with weak realism that is enriched with *pragmatic* elements:

Via successful and less successful interactions with the world, the conceptual structure of an individual will adapt to the structure of reality. It must be emphasized, however, that this does not entail that the conceptual structure *represents* the world. (Gärdenfors, 2000: 156).

The bottom line is this: While the world *does* shape our cognition, we cannot say much at all about *its* structure. Our relation to it is not conceived of in terms of truth

or correspondence, as would be typical for strong realism, but as a matter of practical problem solving. It is basically viewed as a Kantian thing-in-itself, and indeed Gärdenfors labels his epistemological position as neo-Kantian (cf. Gärdenfors, 2000: fn. 161, 171). A more precise version of P3 can now be derived:

P3_{CP}. P2 holds because of the convex concepts' *cognitive-pragmatic* naturalness.

By now it should be clear that P1 and P2 alone are not sufficient for a charitable reconstruction of the argument. Naturalness is meant to contribute decisively to the avoidance of relativism, which distinguishes the proposal from Goodman's (in the normative dimension, that is, differences in their descriptive accounts should be obvious). P3_{CP} adequately captures the spirit of this approach.

To sum up, the metaphysics and epistemology of naturalness in CS build upon Quine's take on the riddle. Natural entities are thought to be real and objective, but still dependent on what practical use they have for us. While Quine tends to emphasize the roles of science and language, Gärdenfors highlights the role of knowledge representation. Cognitive naturalness certainly has a high affinity to conventionalism, but within the framework of Quinean pragmatism a weakly realistic version of the notion can be obtained.

4 The Case for Strong Realism

The remainder of the paper will investigate three lines of reasoning that can be advanced against the argument I just presented. One is as obvious as it is interesting, but it undermines my own position as well – fortunately, it can be refuted. The second one argues positively in favor of strong realism. The third one expresses a concern, rather than providing a knock-down argument. All three points extend into several major debates, so what follows is an explorative rather than an exhaustive discussion. Let me go through them one by one.

4.1 Adaptivity of Poor Reasoning

The theory of evolution is powerful, but one must always be careful not to strain it. It is a common mistake to assume that evolutionary processes automatically lead to progress, and perhaps we are dealing here with an epistemological instance of this fallacy. Gärdenfors and Quine seem to put a lot of hope into evolution when they effectively claim the adaption of cognitive agents to their environment leads to reliable inductive reasoning. But as has been argued many times (cf. Stich, 1990: 55–74 for a systematic philosophical discussion): *unreliable reasoning can be highly adaptive*. Just think of the rich psychological literature on the countless ways in which human cognition has evolved to be biased. If naturalness is tied to what facilitates survival and reproduction, maybe we should not expect it to remedy Goodman's Riddle.

Unfortunately, this point does not only target P3_{CP}, but also the brand of scientific realism I want to advocate here. If the evolutionary perspective gives us reason

distrust of our cognitive faculties, then we surely must also doubt we have reliable access to mind-independent reality: “[...] a fancier style of representing is advantageous *so long as it is geared to the organism’s way of life and enhances the organism’s chances of survival*. Truth, whatever that is, definitely takes the hindmost.” (Churchland, 1987: 549).¹¹ Or as Gärdenfors might frame the issue: Organisms adapt to their cognitively reshaped, *proximal* environments, which they encounter in the form of practical problems (food retrieval, etc.). They have no access to a common *distal* environment, which thus takes on the status of a Kantian Thing-in-itself.¹² There may be an insurmountable gap between what is evolutionarily useful and what the world is really like, call this the evolutionary argument against realism, or EAAR.¹³

Several countering strategies are on the market, but I will focus on a response that would be natural for the CS account to give and then make my way to the second argument against P3_{CP} – which happens to be a cardinal argument against EAAR. Both points I am going to make have to do with one of my favorite insights, namely that *science works*. The first is more direct, countering the claim of unreliability. The second is more ambitious, trying to account for why our faculties might be reliable by suggesting that they capture the structure of the world.

Now, the approach to take is to first point out that even if much of our ‘primordial’ reasoning is flawed, we have developed means to correct for these flaws. Couched in Gärdenfors’ terminology, one could say that some of our *phenomenal* (i.e., subjective, derived from psychological data) CS may be prone to error, but many of our *scientific* (i.e., idealized, objective) spaces are more reliable.

As Quine (1969) notes, it is a sign of mature science that notions of similarity become less and less important, being replaced by theoretically more sophisticated concepts. [...] In this way science builds upon our more or less evolutionarily determined conceptual spaces, but in its most mature form becomes independent of them. (Gärdenfors, 2000: 83).

We are no longer deceived by superficial similarities of dolphins and fish, for instance, but figure the former are actually mammals. This rejoinder appeals to the instrumental success of science and restores hope in the pragmatic utility of our natural concepts: They are at least reliable enough to provide a basis for science to operate on. Note, however, that this move turns Gärdenfors’ argument – as reconstructed above anyway – on its head. All the emphasis is on instrumental success now, which was

¹¹ This is not to say Churchland is a scientific anti-realist, she just exemplifies this line of reasoning nicely.

¹² Gärdenfors, 2000 cites Thompson, 1995 on this, who develops an empirical argument against the (tele-functional) supposition of distal representational objects: “The distal properties detected in color vision form a heterogeneous collection whose type-divisions at the physical level do not match the type-divisions at the perceptual level.” (Thompson, 1995: 6).

¹³ Plantinga 1993 develops an evolutionary argument against naturalism (EAAN) based on the idea that naturalists should believe in evolution, and evolution undermines the reliability of our beliefs – including our beliefs in naturalism and evolution (see Beilby, 2002 for a collection of essays on the issue). I do not want to invoke the details of his argument, however, but highlight the general line of thought is well-known in the philosophical literature.

supposed to be grounded by naturalness. Against this background, the chances of providing more than an instrumentalist justification of reasoning seem dim.

Let me illustrate this line of thought with an example. *Animism* is roughly the tendency to attribute agency to inanimate objects. From the perspective of evolutionary psychology, it can be seen as a reasoning bias that evolved “[...] because those interpretive bets that aim highest (by attributing the most organization and hence significance to things and events) have the greatest potential payoffs and lowest risks. For example, it is better for a hiker to mistake a boulder for a bear than to mistake a bear for a boulder.” Guthrie, 1993: 6. It is therefore plausible to assume that evolution has programmed us to systematically misjudge the degree of organization of our environment, which can lead to generalized doubts about our epistemic and reasoning capacities. To counter this, instrumentalists can point out that science enables us to study biases such as animism, which may put us in a position to avoid them. Moreover, on a general level, science often works remarkably well, which also speaks against generalized skepticism. This countering strategy is also available to Gärdenfors, but adopting it would suggest tying the justification of scientific reasoning to purely instrumental considerations, after all. Appealing to cognitive naturalness is not an option here, as this is the very notion that invites skepticism in the first place. But it would feel contrived to involve naturalness in the discussion on Goodman’s Riddle, but not in a reply to evolutionarily motivated skepticism, as both undermine trust in our reasoning capacities. This suggests that Gärdenfors can either bite the bullet of evolutionarily motivated skepticism, or his theory of naturalness collapses into instrumentalism.

4.2 No (Naturalness-) Miracles

I suggest to improve on the account by making naturalness more metaphysically robust. Yes, many of our natural concepts and scientific generalizations *are* reliable – and this is because they capture nature’s structure. Reasoning that attributes agency to inanimate objects is flawed precisely because it does not capture the relevant structure. Ghosts and their ilk are not natural kinds, even if they are represented by convex CS regions – which they presumably are, by the way. While it is at least feasible we’ve got everything wrong, this seems unlikely, given some of us can send rockets to Mars, vaccinate against Polio, and so on. We simply unveiled enough of the world to achieve these successes – at times *despite* evolution having shaped our cognition. It is only with real natural kinds, properties, etc. in place that the reliability of reasoning makes sense in the first place. This is not engaging in speculative metaphysics, mind you, but a scientifically informed *inference to the best explanation*, or IBE argument.

While this approach is metaphysically more demanding than Quine’s, it is meant to be aligned with his project of *naturalized epistemology*. According to Kornblith (1993), this project is “[...] addressed to two questions: (1) What is the world that we may know it?; and (2) What are we that we may know the world?” Crucially, both answers “must dovetail in important ways” (Kornblith, 1993: 2), and they must be given without recourse to any *prior philosophy* (cf. Hylton & Kemp, 2023: 5). The latter comes down to taking science seriously, and concerning question (2), cognitive

psychology tells us that our access to the external world is *mediated* by the concepts we categorize and reason with. Even if our scientific representations become ever more idealized, objective, etc. – we cannot attain the infamous god’s eye view. Therefore, I do not advocate reducing the issue of naturalness to its ontological dimension, viz. to question (1), but opt for some explication that contains both a necessary cognitive component – whether it turns out to be convexity or anything else that is empirically robust – *and* a necessary ontological component.¹⁴ For the CS account of Goodman’s Riddle, this would mean to replace P_{3CP} with:

P_{3CO} P2 holds because of the convex concepts’ *cognitive-ontological* naturalness.

Prima facie, the ontological part has no *practical* consequences for *descriptively* distinguishing natural from non-natural concepts, or projectible from non-projectible predicates. For that matter, Gärdenfors would likely insist “that adding a metaphysical component (the natural kinds *an sich*) to the conceptual spaces does not improve our understanding of natural kinds.” (Gärdenfors, 1990: 90). But it does have important *theoretical* implications for the justification of our projections, as it makes sense of the very fact that some of our concepts have been reliable to project. It is not enough to point to the evolutionary utility, cognitive efficiency, etc. of the concepts, as their utility and efficiency are revealed at the inescapable touchstone of reality. There is, however, a remarkable paper where Gärdenfors appears to share this view.

But let me take a quick step back. The above argument is, of course, a reiteration of the famous no miracles argument, according to which realism “is the only philosophy that doesn’t make the success of science a miracle.” (Putnam, 1975b: 73). Its strengths and weaknesses are well known – for instance, that it expresses a powerful intuition shared by many philosophers.¹⁵ Or, on the other hand, that there are valid concerns regarding the so-called *base rate fallacy* (cf. Chakravartty, 2017: 2.1). I will not get into the debate here, but what I want to point out is Kornblith (1993) makes an even stronger claim than Putnam does. He develops a *transcendental* argument based on the *modal* claim that reliable induction is impossible without the world playing its part – and curiously enough, Gärdenfors & Stephens, 2017 quote him on this:

“We might not have unbiased contact with the world, but it is still the real world that provides the sensory input we get, and “[i]t is precisely because the world has the causal structure required for the existence of natural kinds that inductive knowledge

¹⁴ I personally align with the views of Kornblith (1993) (except for his hardcore essentialism) and support Boyd’s Homeostatic Property Cluster (HPC) account, which emphasizes accommodation to causal structure. Boyd’s stance on the metaphysics of natural kinds is nuanced: He attributes an intrinsic metaphysical status to the property clusters, yet considers kinds to be real only to the extent that they meet the accommodation demands of *disciplinary matrices*, i.e. collective cognitive and methodological environments in which scientific inquiry and classification take place. Perhaps we can think of these cognitive environments in terms of CS – in any case, Boyd emerges as a strong realist who integrates cognitive elements into his theory.

¹⁵ Note that the *abductive* inference that the argument makes reflects the methodology of the natural sciences. This is in keeping with Quine’s naturalism, but also with the contemporary research program of *inductive metaphysics*, which emphasizes the role of both aprioristic *and* empirical methods for metaphysics (see Engelhard et al. 2021 for reference). Both feed into the methodological background of this paper.

is even possible.” (Kornblith, 1993, p. 35). There are only certain clusters of properties that are organised in a stable enough way as to stick together in natural categories enabling us to make inductive inferences.”

Not only is this an endorsement of strong realism, but of a quite demanding premise, as such modal claims can be difficult to substantiate. However, this stance is an utter exception in Gärdenfors’ work, and can be attributed with some certainty to his collaborator.¹⁶ Even whether the less demanding abductive inference is justified is subject to vigorous debate – however, if the premise can be made to work for scientific realism in general then it will transfer to natural kinds and their ilk, as these are the kinds of entities that are regularly postulated by the special sciences.

With that, there is one more thing I would like to add in this section. In a sense, the IBE miracles argument runs bottom to top: We start with empirical facts about creatures like us – in this case the way we make inferences with convex concepts – and end with real natural kinds as ontological overlay. We can also turn that around: Assuming that realism about kinds is true (and we are thus subject to evolution), we should expect inferentially successful and cognitively economic concepts. Both directions taken together constitute a *package deal* argument for strong realism about natural kinds.¹⁷

4.3 Cognitive-Pragmatic Naturalness – A Poor Explanation?

I have just presented the outlines of a positive argument for a cognitive *cum* ontological approach to naturalness. To round things off, I will briefly return to undermining $P3_{CP}$. Recall the notion of cognitive-pragmatic naturalness is supposed figure as an *explanans* for the instrumental success of convex concepts (viz. P2) as an *explanandum*. My concern is that $P3_{CP}$ might be more of a restatement than an explanation for P2. This is because one appeals to instrumental success, while the other to evolutionary success, and these two might be closely related. If they turn out to imply one another, then we have before us an instance of circular reasoning.

The point seems to hang on what is considered a success in each case. As previously stated, success is a normative notion and implies the presence of some function, or goal. A typical goal in the context of scientific instrumentalism is reliable prediction. There appear to be many cases in which this coincides with evolutionary goals such as survival and reproduction. Reliably predicting the toxicity of a mushroom should aid in survival, which could be taken to indicate $P2 \supset P3_{CP}$. However, if framed a bit differently, this same example might be used to disambiguate instrumental and evolutionary success. For instance, if a strategy of overly cautious mushroom categorization – yielding many false positives – enhances survival, it could be considered evolutionarily reliable *qua* being instrumentally unreliable, i.e., $\neg(P3_{CP} \supset P2)$.

Unfortunately, the issue is too complex to be resolved in this paper, as it would quickly devolve into a systematic comparison of scientific instrumentalism and prag-

¹⁶ Gärdenfors has confirmed to me in a personal conversation that he has not become a strong realist.

¹⁷ I would like to thank Markus Schrenk, who brought up the idea of a package deal in a discussion, thus complementing the IBE-argument.

matism – which is a mammoth project. But even assuming the desired disambiguation is feasible, we are still left with the diagnosis that on the CS account, grounding instrumental success in naturalness amounts to grounding it in evolutionary success. The last sections have cast doubt on whether this is a good idea. To be clear, I do not want to claim strong realism is the only game in town as a theory of naturalness. ‘Cognitive-ontological naturalness’ means I am open to nominalism as well as weak conventionalism, as long as a mind-independent component is involved. I also want to highlight that Gärdenfors’ approach demonstrates appealing to naturalness is possible for instrumentalists, pragmatists, and the New Realists who want to get rid of the traditional commitment to mind-independence (e.g., Chang’s, 2022 operational coherence, Massimi’s, 2022 perspectival realism, or Kendig’s, 2015 practice grounded kinds). It is an open question, however, whether these and comparable authors would subscribe to grounding reasoning in evolution. After all, even if circularity can be avoided and EAAN as well as independent arguments for strong realism can be refuted, the normative question remains addressed in a very familiar way. The quote from Sect. 2.1 by Cohnitz & Rossberg, 2020 can be slightly adjusted to illustrate this: “[Gärdenfors] makes projectibility essentially a matter of what [concepts] we use and have used to describe and predict the behaviour of our world” – not the world in and of itself though, but an evolutionarily reshaped version of it. This move may enable one to sidestep *Goodmanian* relativism, but at the cost of an evolutionary one. Whether this is much progress, I leave for the reader to decide.

5 Conclusion

Returning to the title question of this paper: Are Gärdenfors’ Conceptual Spaces a viable solution to Goodman’s Riddle? They most certainly have something to offer. Convexity, as a necessary condition of projectibility, leads to rich empirical predictions and – given those predictions will obtain in the long run, which I have not debated here – provides a descriptive explanation of our inferential practices, as well as a practically implementable demarcation criterion. This has to count as progress, especially from a cognitive science perspective – but it is progress in the descriptive dimension of the riddle. If progress in the normative dimension is the benchmark, then, as I have argued, CS in their current form miss it.

Section 2.1 has shown that already Hume and Goodman saw the instrumental success of some of our inferences provides a basis for their ‘weak’ justification. Gärdenfors can be interpreted as appealing to instrumental success as well (P2), but I have presented a reconstruction of his account that aims for a wider justification strategy by grounding instrumental success in cognitive-pragmatic naturalness (P3_{CP}). It is doubtful, however, whether this notion of naturalness is up to the task. Evolutionarily motivated skepticism may be avoidable, but the most obvious countering strategy appeals to the success of science (i.e., no progress in the normative dimension), as appealing to naturalness is not an option, if it is tied to evolutionary success – which is the very notion that invites skepticism in the first place. As a result, the wider justification strategy collapses into the familiar instrumentalist one. By contrast, appealing to naturalness is very straightforward on a strongly realist reading: Reasoning

is justified insofar as it captures the structure of nature. Section 4.2 has outlined a ‘package deal’ argument in favor of giving natural entities an intrinsic metaphysical standing – but without disregarding the progress that comes with cognitive criteria of naturalness. My positive case is partially based on the miracles argument from the great debate about scientific realism, where I’m in good company with philosophers like Hilary Putnam. My negative case against cognitive-pragmatic naturalness is strengthened by considerations about the relation of instrumental and evolutionary success, which the CS account would need to disambiguate more clearly for the grounding justification strategy to work.

I agree with Gärdenfors that going below the level of language to understand reasoning is the right move. But it is ill-advised to tie the justification of reasoning exclusively to how we represent the world around us. It is the structure of the natural world itself that determines which of our inferences work and which do not.

Acknowledgements The research was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft DFG) under project number 462340335. I would like to thank my colleagues and supervisors in Düsseldorf, especially Daian Bican, Giacomo Gianni, Jan Michel, Markus Schrenk, Gerhard Schurz, and Gottfried Vosgerau, for their support and most helpful discussions.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Beilby, J. K., Ed. (2002). *Naturalism defeated? Essays on Plantinga’s evolutionary argument against naturalism*. Cornell University Press
- Bird, A., & Tobin, E. (2023). ‘Natural Kinds’. In E. N. Zalta, & U. Nodelman (Eds.), *The stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/natural-kinds/>.
- Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New interdisciplinary essays*, (pp. 141–185). MIT Press.
- Chakravartty, A. (2017). Scientific realism. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2017/entries/scientific-realism/>.
- Chang, H. (2022). *Realism for realistic people: A new pragmatist philosophy of science*. Cambridge University Press. <https://doi.org/10.1017/9781108635738>.
- Churchland, P. S. (1987). Epistemology in the age of neuroscience. *Journal of Philosophy*, 84(10), 544–553. <https://doi.org/10.5840/jphil1987841026>.
- Cohnitz, D., & Rossberg, M. (2020). Nelson goodman. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy, Summer* (pp. 1–25). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2020/entries/goodman/>.

- Crane, J. K. (2021). Two approaches to natural kinds. *Synthese*, 199(5–6), 12177–12198. <https://doi.org/10.1007/s11229-021-03328-9>.
- Douven, I., & Gärdenfors, P. (2019). What are natural concepts? A design perspective. *Mind and Language*, 35(3), 313–334. <https://doi.org/10.1111/mila.12240>.
- Earman, J. (1985). Concepts of projectibility and the problems of induction. *Noûs*, 19(4), 521–535. <https://doi.org/10.2307/2215351>.
- Engelhard, K., Feldbacher-Escamilla, C. J., Gebharter, A., & Seide, A. (2021). Inductive metaphysics: Editors' introduction. *Grazer Philosophische Studien*, 98(115), January, 1–26. <https://doi.org/10.1163/18756735-00000129>.
- Fine, K. (2001). The question of realism. *Philosophers' Imprint*, 1, 1–30.
- Gärdenfors, P. (1990). Induction, conceptual spaces and AI. *Philosophy of Science*, 57(1), 78–95.
- Gärdenfors, P. (1999a). Some tenets of cognitive semantics. In J. Allwood, & P. Gärdenfors (Eds.), *Cognitive semantics: Meaning and cognition* (pp. 19). John Benjamins Publishing Company. <https://doi.org/10.1075/pbns.55.03gar>
- Gärdenfors, P. (1999b). Does semantics need reality? In A. Riegler, M. Peschl, & A. Stein (Eds.), *Understanding representation in the cognitive sciences* (pp. 209–217). Springer. https://doi.org/10.1007/978-0-585-29605-0_23
- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT Press Cambridge, Massachusetts
- Gärdenfors, P. (2008). Reasoning in conceptual spaces. *Reasoning: Studies of human inference and its foundations* (pp. 302–320). Cambridge University Press. <https://doi.org/10.1017/CBO9780511814273.018>.
- Gärdenfors, P. (April 2009). Meanings as conceptual structures. *Lund University Cognitive Studies*, 16, 1–13.
- Gärdenfors, P. (1 January 2011). Inductive reasoning: From Carnap to cognitive science. 1–16.
- Gärdenfors, P., & Stephens, A. (2017). Induction and knowledge-what. *European Journal for Philosophy of Science*, 8(3), 471–491. <https://doi.org/10.1007/s13194-017-0196-y>.
- Goodman, N. (1983). *Fact, fiction, and forecast*. Vol. 4th Edition. Harvard University Press.
- Guthrie, S. E. (1993). *Faces in the clouds: A new theory of religion*. Oxford University Press USA.
- Henderson, L. (2020). The problem of induction. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2020/entries/induction-problem/>.
- Hernández-Conde, J. (2017). A case against convexity in conceptual spaces. *Synthese*, 194(10), 4011–4037. <https://doi.org/10.1007/s11229-016-1123-z>.
- Hylton, P., & Kemp, G. (2023). Willard van Orman Quine. In E. N. Zalta, & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/quine/>.
- Kaipainen, M., & Hautamäki, A. (2015). A perspectivist approach to conceptual spaces. In F. Zenker and P. Gärdenfors (Eds.), *Applications of conceptual spaces: The case for geometric knowledge representation*, 245–258. Synthese Library. Springer International Publishing. https://doi.org/10.1007/978-3-319-15021-5_13.
- Kendig, C. (2015). Editor's introduction: Activities of kinding in scientific practice. *Natural kinds and classification in scientific practice* (pp. 1–14). Routledge.
- Khalidi, M. A. (2013). *Natural categories and human kinds: Classification in the natural and social sciences*. Cambridge University Press.
- Khalidi, M. A. (2016). Mind-dependent kinds. *Journal of Social Ontology* 2(21 August), 223–246. <https://doi.org/10.1515/jso-2015-0045>.
- Kornblith, H. (1993). *Inductive inference and its natural ground: An essay in naturalistic epistemology*. The MIT Press
- Kuhn, T. S. (2012). *The structure of scientific revolutions: 50th Anniversary Edition*. The University of Chicago
- Kusch, M. (2021). *Relativism in the philosophy of science*. Cambridge University Press.
- Lewis, D. (1983). New Work for a theory of universals. *Australasian Journal of Philosophy*, 61(4), 343–377. <https://doi.org/10.1080/00048408312341131>.
- Lewis, D. (1986). *On the plurality of worlds*. Wiley-Blackwell.
- Massimi, M. (2022). *Perspectival realism*. *Oxford studies in philosophy of science*. Oxford University Press
- Mormann, T. (1993). Natural predicates and topological structures of conceptual spaces. *Synthese*, 95(2), 219–240. <https://doi.org/10.1007/BF01064589>.

-
- Putnam, H. (1975a). *Philosophical papers: Volume 1: Mathematics, matter, and method*. Cambridge University Press
- Putnam, H. (1975b). *Philosophical papers, volume 2: Mind, language, and reality*. Cambridge University Press
- Putnam, H., & Goodman, N. (1983). Foreword to the Fourth Edition. *Fact, fiction, and forecast* (4th ed.). Harvard University Press.
- Quine, W. V. (1969). Natural kinds. *Ontological relativity and other essays* (pp. 114–138). Columbia University Press. <https://doi.org/10.7312/quin92204-006>
- Schrenk, M. (2016). *Metaphysics of science: A systematic and historical introduction*. Routledge. <https://doi.org/10.4324/9781315639116>.
- Slater, M. H., & Borghini, A. (2013). Introduction: Lessons from the scientific butchery. In M. O'Rourke, M. H. Slater, & J. K. Campbell (Eds.), *Carving nature at its joints: Natural kinds in metaphysics and science* (pp. 1–32). MIT Press
- Stich, S. P. (1990). *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation*. The MIT.
- Strößner, C. (2022). Criteria of naturalness in conceptual spaces. *Synthese*, 200(78), 36.
- Thompson, E. (1995). Colour vision, evolution, and perceptual content. *Synthese*, 104(11 July), 1–32. <https://doi.org/10.1007/BF01063672>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.