Synthese
DOI 10.1007/s11229-017-1506-9

S.I. : INFERENTIALISM

Inferentialism and knowledge: Brandom's arguments against reliabilism

José L. Zalabardo[1]

Received: 26 August 2016 / Accepted: 19 July 2017
© The Author(s) 2017. This article is an open access publication

Abstract I take issue with Robert Brandom's claim that on an analysis of knowledge based on objective probabilities it is not possible to provide a stable answer to the question whether a belief has the status of knowledge. I argue that the version of the problem of generality developed by Brandom doesn't undermine a truth-tracking account of noninferential knowledge that construes truth-tacking in terms of conditional probabilities. I then consider Sherrilyn Roush's claim that an account of knowledge based on probabilistic tracking faces a version of the problem of generality. I argue that the problems she raises are specific to her account, and do not affect the version of the view that I have advanced. I then consider Brandom's argument that the cases that motivate reliabilist epistemologies are in principle exceptional. I argue that he has failed to make a cogent case for this claim. I close with the suggestion that the representationalist approach to knowledge that I endorse and Brandom rejects is in principle compatible with the kind of pragmatist approach to belief and truth that both Brandom and I endorse.

Keywords Reliabilism · Truth tracking · Probabilistic tracking · Problem of generality · Pragmatism · Representationalism · Knowledge · Belief · Truth · Robert Brandom · Inferentialism

1 Representationalism and pragmatism

Suppose we have a range of objects, the As, and we want to explain the meaning of classifying them as Bs or not Bs. What form should this explanation take? One obvious

✉ José L. Zalabardo
j.zalabardo@ucl.ac.uk

1 Philosophy Department, University College London, Gower Street, London WC1E 6BT, UK

Published online: 29 July 2017

possibility is to specify what an A has to be like in order to count as a B—i.e. to identify the property or condition that has to be present in the As that we want to call Bs. Then we can explain the meaning of classifying an A as a B as the belief that the A in question instantiates this property or satisfies this condition. Let me refer to explanations of the meaning of classifying an A as a B that follow this pattern as *representationalist*.[1] A representationalist explanation of the meaning of classifying As as Bs will enable us to think of these classifications as representations—they will represent As as having the property or satisfying the condition that figures in our explanation. However, I want to use the label *representationalism* to denote the explanatory strategy, not the conception of the explanandum as a representational activity. And I want to leave the possibility open of conceiving of a classificatory practice as playing a representational role even if its meaning has not been explained along the lines of the representationalist model—by identifying the property or condition that has to be present in the As that we want to call Bs.

A second approach to the task of explaining the meaning of classifying an A as a B proceeds by specifying the rules that govern this classificatory practice. On this approach, to classify As as Bs or not Bs is to classify them according to these rules. Clearly any representationalist account can be turned into an explanation of this kind. If an A will count as a B just in case it instantiates property P, then the practice of classifying As as Bs or not Bs can be characterised as governed by the following rule: classify an A as a B just in case you believe that it instantiates property P. The converse, however, doesn't hold. If we explain the meaning of classifying As as Bs or not Bs by specifying the rules that govern this classificatory practice, the rules that we employ in this explanation might not yield a characterisation of what an A has to be like in order to count as a B. Let me say that an explanation of the meaning of classifying As as Bs or not Bs is *pragmatist* when it proceeds by specifying the rules that govern this classificatory practice and these rules don't yield a characterisation of what an A has to be like in order to count as a B. We get an account of this kind, for example, if the rules that govern the practice of classifying As as Bs make reference to the state of information of the classifier, but we think of the question, whether a given A is a B, as having a unique answer for all classifiers and states of information.

The contest between representationalist and pragmatist approaches to the explication of a classificatory practice articulates some important philosophical debates. Take, for example, the practice of classifying actions as morally right or not morally right. Moral realists can be seen as adopting a representationalist approach to the explication of this practice. They will start by specifying the property whose presence in or absence from an action determines whether it is morally right. Then classifying an action as morally right will be explained as the belief that this property is present in the action in question. Expressivists, by contrast, will favour a pragmatist strategy. They will characterise the classification of actions as morally right or not as the practice that is governed by certain rules—in the most basic version of the view by the following rule: classify an action as morally right just in case you feel moral approval towards it. On this approach, to classify an action as morally right or not morally right is to

[1] The term has been used in something like this sense by Huw Price. See, for example Price (2004).

classify it according to this rule. But the rule doesn't yield a specification of what an action has to be like in order to count as morally right. The morally right actions are not necessarily those towards which I feel moral approval. My moral sense is subject to change, and I don't take these changes to bring about changes in which actions are morally right. Likewise, my moral sense is often different from those of my peers, but I don't take the moral rightness of an action to be relative to the classifying individual.

To be sure, the debate between realists and expressivists about moral discourse is not only about which strategy we should adopt for explicating moral discourse. It is also about the function of the discourse—about what we are doing when we ascribe "morally right" to an action. Whereas for the moral realist with these ascriptions we represent actions as being a certain way, the expressivist sees these ascriptions instead as expressing the attitude of moral approval towards the actions. However, as I mentioned above, the two issues—explanatory strategy and the function of the discourse or practice—are in principle independent of each other. It is the former—explanatory strategy—that I want to concentrate on here, remaining open to the possibility that a classificatory practice that has been explicated along pragmatist lines might still be characterised as in the business of representing things as being a certain way.

The contest between representationalism and pragmatism can also be applied to the practice of classifying beliefs, judgments or assertions as true or false. Substantive theories of truth can be characterised as versions of representationalism. They will start with a specification of the condition that one of these items has to satisfy in order to count as true. Then classifying one of them as true will be explained as the belief that the truth-property is present in the item in question. But the practice of truth ascription can also be explicated in terms of the rules that govern it—using rules that don't yield a specification of the condition that an item has to satisfy in order to count as true. Elsewhere I have defended an account of the practice of assessing judgments as true or false that follows this template (Zalabardo 2016).

The contest can also be formulated with respect to the practice of classifying true beliefs according to whether or not they have the status of knowledge. Representationalists, on the one hand, will seek to explicate these classifications by identifying the condition whose satisfaction or otherwise by a true belief will determine whether it should be classified as knowledge. Then the ascription of the status of knowledge to a true belief will be explained as the belief that this condition is satisfied by the true belief in question. Pragmatists, on the other hand, will seek to explicate these classifications by reference to the rules that govern them, using rules that don't yield an informative characterisation of the condition that a true belief has to satisfy in order to count as knowledge.

Robert Brandom's views on knowledge place him squarely on the pragmatist side of this divide. In *Articulating Reasons*, he presents a contrast between two forms that a theory of knowledge can take that is broadly parallel to the contrast I've drawn between representationalism and pragmatism:

> Epistemology is usually thought of as the theory of knowledge. But epistemological theories in fact typically offer accounts of when it is proper to *attribute* knowledge: for instance, where there is justified belief, or where true beliefs have resulted from reliable belief-forming processes. Now, a theory of knowl-

> edge can take this form. The two might be related as formal to material mode, in Carnap's terminology; instead of asking what X's are, we ask when the term 'X' is properly applied. But the two need not be versions of the same question. In the case of knowledge, I think they stand in a more complex relationship (Brandom 2000: pp. 117–18).[2]

Brandom's own account of knowledge proceeds by focusing on the conditions under which the term 'knowledge' is properly applied:

> In calling what someone has 'knowledge', one is doing three things: *attributing* a *commitment* that is capable of serving both as premise and as conclusion of inferences relating it to other commitments, *attributing entitlement* to that commitment, and *undertaking* that same commitment oneself (Brandom 2000: p. 119).

In his account of when we can attribute entitlement, Brandom incorporates an insight that he extracts from Alvin Goldman's reliabilism:

> [...] entitlement may be attributed on the basis of an assessment of the reliability of the process that resulted in the commitment's being undertaken (Brandom 2000: p. 118).

But while he is receptive to reliabilist ideas in this context, Brandom is firmly opposed to a reliabilist version of representationalism—to a reliabilist account of the condition that a true belief has to satisfy in order to count as knowledge. According to Brandom, reliabilism doesn't have the resources to accomplish this task. His discussion is focused on Goldman's process reliabilism, but he clearly wants to draw a more general conclusion:

> When we understand properly the sense in which facts about the reliability of a mechanism can be objective, we see that appeals to objective probability fall short of enabling fully naturalistic accounts of knowledge (Brandom 2000: p. 112).

Process reliabilism is not the only way in which we can appeal to objective probability in our account of knowledge. In Zalabardo (2012), I defended an analysis of knowledge in which tracking the truth is a sufficient condition for a true belief to have the status of knowledge, with truth-tracking construed, not in terms of subjunctive conditionals, as in Nozick's original proposal (1981), but, following Roush (2005), using conditional probabilities.

My primary goal in the present paper is to provide a critical assessment of Brandom's arguments against reliabilism—from the point of view of the version of the truth-tracking account that I favour. I am going to argue that this analysis of knowledge is not undermined by Brandom's arguments. I will then contend that the representationalist account of knowledge sustained by this analysis is perfectly compatible with a pragmatist explication of belief and truth ascriptions.

[2] See also Brandom (1994: p. 212).

2 Barn facades

The problem that Brandom invokes to undermine reliabilism is what has come to be known as the *problem of generality*, first noted by Goldman (1979) and developed in detail by Feldman (1985) and Conee and Feldman (1998). In Goldman's presentation, the problem arises from the fact that the process-token that produces a belief will be an instance of several process-types. These types might have different degrees of reliability. Hence to determine the epistemic status of the belief, the reliabilist would have to specify which of these types is epistemically relevant. Brandom presents the problem as arising from a fact about objective probabilities:

> An objective probability can be specified only relative to a reference class. And [...] the world as it objectively is, apart from our subjective interests and concerns (paradigmatically, explanatory ones), does not in general privilege one of the competing universe of possible reference classes as the correct or appropriate one (Brandom 2000: pp. 112–13).

Brandom illustrates the point with a version of the fake-barn thought experiment that Goldman had deployed against causal theories of knowledge (Goldman 1976). Brandom asks us to consider a normal perceiver in standard conditions who is looking at a barn and forms as a result the belief that there's a barn in front of him. As in Goldman's original example, Brandom asks us to imagine that the barn the subject is looking at is located in Barn Facade County. The only real barn in the county is the one the subject is looking at. The county contains, in addition, 999 barn facades, visually indistinguishable from actual barns. If the subject had been looking at one of the barn facades, he would have formed the belief that he was looking at a barn.

Goldman expects the case, thus described, to elicit the intuition that the subject's true belief that there's a barn in front of him is not knowledge, because it was formed by a belief-forming process that's not reliable—in the circumstances prevailing in Barn Facade County, where the belief was formed. But Brandom seeks to undermine the legitimacy of this verdict by adding some details to Goldman's example:

> For suppose that Barn Facade County is one of a hundred counties in the state, all the rest of which eschew facades in favor of actual barns. Then, considered as an exercise of a differential responsive disposition within the *state*, rather than within the county, our subject's process of perceptual belief-formation may be quite reliable, and hence when it in fact yields correct beliefs, it may underwrite attributions of perceptual *knowledge.* But then, if the whole country, consisting of fifty larger states, shares the habits of Barn Facade County—so that over the whole country (excepting this one state) facades predominate by a large margin—then considered as a capacity exercised in the *country*, the very same capacity will count as quite *un*reliable, and hence as insufficient to underwrite attributions of knowledge. And then again, in the whole *world,* barns may outnumber facades by a large margin. So considered with respect to that reference class, the capacity would once again count as reliable. And so on. Do we need to know about the relative frequencies of barns and facades in the solar system or the galaxy in

> order to answer questions about the cognitive status of our subject's beliefs? (Brandom 2000: pp. 115–16)

Goldman readily accepts the challenge to process reliabilism illustrated by Brandom's example, and in *Epistemology and Cognition* he tentatively puts forward a strategy for meeting it (Goldman 1986: pp. 50–51). Brandom doesn't discuss Goldman's or any other attempted solution. In his view, no solution could ever work, as nothing in the world singles out a unique level of generality as epistemically significant:

> Which is the correct reference class? Is the perceiver an objectively reliable identifier of barns or not? I submit that the facts as described do not determine an answer. Relative to each reference class there is a clear answer, but nothing in the way the world is privileges one of those reference classes, and hence picks out one of those answers (Brandom 2000: p. 116).

My goal here is not to assess Goldman's solution, or indeed the ability of process reliabilism to meet Brandom's challenge. I am going to argue instead that a different strategy for explicating non-inferential knowledge in terms of objective probabilities has the resources for overcoming Brandom's challenge.

The strategy that I want to explore belongs to the truth-tracking tradition. Truth-tracking conditions for knowledge were formulated by Nozick as subjunctive conditionals (1981). In order to track the truth, according to Nozick, a belief has to be *sensitive* and *adherent*.[3] S's belief that p is sensitive just in case if p were false S would not believe p, and adherent just in case if p were true S would believe p. But as Roush has shown, these conditions can also be naturally expressed in terms of conditional probabilities (2005). Sensitivity is now expressed as a low probability of S believing p if p is false, i.e. a low value for pr(Bel(p)| ∼p) (or, equivalently, a high value for pr(∼Bel(p)| ∼p)), and adherence as a high probability of S believing p if p is true, i.e. a high value for pr(Bel(p)|p).[4] Roush treats high levels of probabilistic sensitivity and adherence of a true belief as individually necessary and jointly sufficient conditions for knowledge by truth tracking. Her account also contemplates a second way in which a true belief can acquire the status of knowledge, besides truth tracking, but here we shall not be concerned with this aspect of her position. My goal in the remainder of this section is to consider how Brandom's challenge to reliabilism can be adapted to target probabilistic tracking, and to argue that, in this case, the challenge doesn't succeed.

Let BARN denote the proposition that there is a barn in front of our subject. The sensitivity of his belief in this proposition will be expressed by the value of pr(∼Bel(BARN)| ∼BARN) and its adherence by the value of pr(Bel(BARN)|BARN). How does the presence of fake barns affect these values? We can easily see that the adherence of the belief will be unaffected: the presence of fake barns in the vicinity should have no influence on what the subject believes when he is in front of a genuine

[3] Nozick himself used the term *sensitivity* as synonymous with *truth-tracking*, and refers to what I'm calling *sensitivity* as *variation*. For my use of the label see, e.g., DeRose (1995: p. 27) and Williamson (2000: p. 148).

[4] Following Nozick, Roush adds to her formulation of the adherence of S's belief that p the condition that the value of pr(Bel(∼p)|p) should be low (Roush 2005: p. 45).

barn. Its sensitivity, by contrast, will be reduced. Surely the fact that our subject is surrounded by fake barns can be expected to increase the probability that he forms the belief that there is a barn in front of him when there isn't one. This will happen whenever he is placed in front of one of the fakes. By reducing its sensitivity, fake barns can in principle prevent the subject's true belief that there's a barn in front of him from attaining the status of knowledge, on the probabilistic tracking account.

We can consider the situation in light of a principle that I call *distribution* (Zalabardo 2012: p. 112), according to which the value of the sensitivity of the subject's belief in BARN (pr(∼Bel(BARN) | ∼BARN)) can be rewritten as follows—with FAKE denoting the proposition that the subject is in front of a fake barn:

$$\text{pr}(\sim\text{Bel(BARN)} \mid \sim\text{BARN} \,\&\, \text{FAKE}) \cdot \text{pr}(\text{FAKE} \mid \sim\text{BARN}) +$$
$$\text{pr}(\sim\text{Bel(BARN)} \mid \sim\text{BARN} \,\&\, \sim\text{FAKE}) \cdot \text{pr}(\sim\text{FAKE} \mid \sim\text{BARN})$$

I take it to be common ground (1) that pr(∼Bel(BARN) | ∼BARN & FAKE) has a low value, since barn facades are so convincing, and (2) that pr(∼Bel(BARN) | ∼BARN & ∼FAKE) has a high value—the subject is highly unlikely to form the belief that he is in front of a barn when he occupies what I propose to label as *transparently barnless locations* (TBLs)—containing neither a real barn nor a barn facade. This suggests that the sensitivity of the belief will depend on the ratio of TBLs to facades. An increase in this ratio will increase the value of pr(∼FAKE | ∼BARN) and reduce the value of pr(FAKE | ∼BARN). And this, in turn, will increase the sensitivity of the belief.

At this point the proposal appears to face a version of Brandom's challenge. The ratio of TBLs to fakes, and hence the sensitivity of the belief, will depend on which geographic area we treat as epistemically significant. In Barn Facade County this ratio will be very low, and the belief will be very insensitive relative to this geographic area. In the state as a whole, however, the ratio will be much higher, thanks to the TBLs to be found in the remaining 99 counties, with no additional fakes to compensate. Hence, relative to the state, the belief will count as sensitive. But the fakes in the rest of the country will lower the ratio again, bringing down once more the sensitivity of the belief. Brandom's challenge to reliabilism can be easily adapted to probabilistic tracking:

> Which is the correct reference class? Is the subject's belief that there's a barn in front of him sensitive or not? I submit that the facts as described do not determine an answer. Relative to each reference class there is a clear answer, but nothing in the way the world is privileges one of those reference classes, and hence picks out one of those answers.

This adaptation of Brandom's challenge focuses on the question, how far out a barn facade (or a TBL) can be located and still be allowed to affect the sensitivity of the belief. Brandom's point would be that different answers to this question will yield different values for the sensitivity of the belief, and hence different verdicts on the status of the belief as knowledge, on the probabilistic-tracking account, but there is no non-arbitrary way of picking one of these answers. I'm going to argue, to the contrary, that the question can receive a perfectly natural, principled answer.

In order to motivate this answer, we need to consider the intuitive reason why barn facades are supposed to undermine the epistemic status of a true belief in BARN. I want to suggest that the reason has to do with the risk they pose of the subject believing BARN in error. Knowledge intuitively requires that this risk should be low, but barn facades increase the risk: if the subject had been placed in front of one of them, he would believe BARN in error.

However, not all barn facades will increase this risk to the same extent. The risk of error posed by a given barn facade can be expected to be an inverse function of the probability that the subject is located in front of it. Barn facades in locations that the subject is highly likely to occupy can be expected to lead to significant increases in the risk of error, while facades in locations that the subject is less likely to occupy will not affect the risk to the same extent. I take it that these considerations underlie the epistemic significance of the political borders in Brandom's example: it is assumed that the subject is more likely to be placed in front of barn facades in his home county than in front of those in other counties in the state, more likely to be placed in front of these than in front of out-of-state facades, etc.

The probability that the subject will occupy a given location will be affected by many factors, but I propose to simplify matters by assuming that it is determined by distance alone—that as the distance between a subject and a facade increases, the probability that the subject is located in front of it will be uniformly reduced. On this assumption, the risk of error posed by a given barn facade will be an inverse function of the distance between the facade and the subject. Close-by facades will bring about substantial increases in the risk that the subject believes BARN in error, more distant facades will bring more moderate increases in the risk of error, and remote facades will produce only insignificant increases in this risk.

These reflections ground a non-arbitrary way of picking an answer to the question, how far out a facade can be located and still be allowed to affect the status of the subject's belief in BARN as knowledge. The answer is this: all facades should be taken into account, but the extent to which they undermine this status should vary with the distance between the subject and the location of the facade—as this distance increases, facades should pose a smaller obstacle to the status of the belief as knowledge. Likewise, all TBLs should be allowed to boost the claim of the belief to this status, but their influence should decrease as the distance from the subject increases. And for remote barnless locations, which the subject is very unlikely to occupy, whether or not they contain a facade should have only a marginal influence on whether the belief is knowledge.

Notice, finally, that this is the answer that we obtain if we treat probabilistic sensitivity as a necessary condition for knowledge. As we saw above, given the subject's doxastic dispositions, the value of pr(∼Bel(BARN)| ∼BARN) will decrease as the value of pr(FAKE | ∼BARN) increases. This value, in turn, will be increased by barn facades and reduced by TBLs. But the influence on this value of an individual facade or TBL will decrease as its distance from the subject increases, on the assumption, once more, that the subject is more likely to occupy nearby locations than more distant ones. As a result, the 999 facades in Barn Facade County will drive the sensitivity of the belief down. The 99 fakeless counties in the state will drive it back up, but are unlike to neutralize the effect of the fakes in Barn Facade County. The fakes in the

remaining 50 states will once again drag the sensitivity of the belief down, but won't compensate for the increase effected by the fakeless counties in the state, etc. And the content of the barnless locations in the rest of the world or the galaxy will only have an insignificant effect on the sensitivity of the belief, given the very low probability that the subject occupies one of these locations. In practice, the 999 fakes in Barn Facade County can be expected to drive the value of pr(∼Bel(BARN)| ∼BARN) so far down that it will remain below the threshold for knowledge whatever happens beyond the county's borders, let alone further away. The advocate of probabilistic tracking can give a negative answer to Brandom's rhetorical question: "Do we need to know about the relative frequencies of barns and facades in the solar system or the galaxy in order to answer questions about the cognitive status of our subject's beliefs?" No, we don't, since the probability that the subject occupies those (barnless) locations is so low that what he would believe if he occupied them can only have a negligible influence on the sensitivity of his belief.

I conclude that, when levelled against probabilistic tracking, Brandom's challenge can receive a plausible answer: if we treat probabilistic sensitivity as a necessary condition for knowledge, each barn facade will affect the status of the belief as knowledge to the extent that it should, given its location. Objective probabilities, deployed in this way, can overcome Bandom's challenge.

3 Generality strikes back

In the previous section I've argued that probabilistic tracking has the resources for dealing with Brandom's version of the problem of generality. However, this victory might be short-lived, since, according to Roush, probabilistic-tracking accounts of knowledge face another version of the problem. According to her, the problem is faced, in some form or another, by any externalist view. The problem is, she tells us,

> saying in general and in a principled way the level of generality or specificity at which facts about the world are to be described when they enter into judgments of whether the subject's belief has the right relation to the facts to be counted as knowledge (Roush 2005: p. 76).

For probabilistic tracking, the problem takes the form of specifying which probability functions we are using. This involves, among other things, specifying "which statements to fix as they are in the actual world and which to allow to take other values" (Roush 2005: p. 75).

On the account of probability that I favour, probabilities arise from the nomological order—they are determined by the laws of nature (Zalabardo 2012: pp. 68–70).[5] Hence in order to obtain the probabilities that figure in my account of knowledge all we need

[5] See Lewis (1986, 1994). The notion of probability presupposed by my account of truth tracking differs from Lewis' in two important respects. First, propositions can have non-trivial probabilities even in a deterministic world. See Hoefer (2007) and Glynn (2010) for proposals as to how to achieve this. Second, propositions about past events can have non-trivial probabilities. We can achieve this by rejecting Lewis' idea that the history of the universe is taken into account in the determination of probabilities. See Hoefer (2007).

to fix is the relevant nomological facts. Let me use the label *generalism* for this position on how to specify the probabilities that figure in the tracking conditions.

Roush rejects generalism. She thinks that the view that "nothing except general facts about things and the way they work is fixed" will have disastrous consequences (Roush 2005: p. 78). I am going to argue that the difficulties that Roush raises for generalism are specific to her version of probabilistic tracking. On the version of the position that I have defended (Zalabardo 2012), generalism doesn't face these difficulties.

The first problem that Roush considers is posed by cases in which the probability of S believing p given p is extremely small simply because the probability that S will form a belief on p is extremely small. She illustrates the point by taking S to be a palaeontologist and p the proposition that dinosaur Bob, which S has been studying, was a vegetarian. Under the right circumstances, we would want to say that S knows p, however, S's belief that p will have a very low level of adherence:

> From the point of view of a world where there was a male vegetarian dinosaur at the appropriate time and other things are constrained to be as they have been since only insofar as p's truth and the way things generally work constrain that, S cannot be said to be likely to form any belief at all about Bob. She might have been a carpenter or taxicab driver or any number of other things instead of a palaeontologist. In fact, S might not have existed at all, or might have existed but died in the car crash she narrowly escaped ten years ago, all of which says that condition IV [adherence] is not fulfilled (Roush 2005: p. 78).

This kind of case certainly poses a problem for Roush, since she treats a high level of adherence as a necessary condition for knowledge by tracking. Given the low level of adherence of the palaeontologist's belief, generalism would force Roush to conclude that the palaeontologist's belief doesn't track the truth. To solve the problem, she proposes to depart from generalism by stipulating "that the probability functions used to evaluate IV assign probability 1 to the claim that either S believes p or S believes –p" (Roush 2005: p. 79).

Another problem that Roush raises for generalism is that it doesn't fix anything "about the actual situation in which the subject came to believe, and her placement in it" (Roush 2005: p. 79). This is a problem because

> if we did not fix something here, there could be an overwhelming number of scenarios considered in which our subject was not even present and so would have a 50–50 chance of forming either belief on the assumption that she formed one or the other belief (Roush 2005: p. 79).

Roush worries that this would have a swamping effect that would prevent us from treating as knowledge beliefs to which we should accord this status. In order to overcome this problem, Roush contends, we need to abandon generalism in favour of a more restrictive specification of the probability functions used in evaluating the tracking conditions.

I agree with Roush that generalism makes her account yield incorrect knowledge verdicts. Hence her version of probabilistic tracking requires a more involved specification of which probabilistic functions are being used to determine the adherence and

sensitivity of beliefs—i.e. a solution to her version of the problem of generality. I want suggest, however, that the difficulties are specific to her version of the view, and they not affect the account that I defended in Zalabardo (2012). The difference between the two accounts that matters for our purposes is that whereas Roush treats high absolute values of probabilistic adherence and sensitivity as necessary conditions for knowledge by tracking, I replace this with treating a high *ratio* of adherence to sensitivity (pr(Bel(p)|p)/pr(Bel(p)| ∼p)), to which I refer as the *tracking ratio* of a belief, as a necessary condition for knowledge by tracking. On my account, what matters for truth tracking is not how likely the subject is to believe p if p is true, or how unlikely she is to believe p if p is false. What matters instead is how much *more* likely she is to believe p if p is true than if it is false. As I argue in Zalabardo (2012: pp. 213–14), this approach can be defended on independent grounds. One consequence of focusing on tracking ratios is that sensitivity plays a more important role than adherence in determining whether a belief tracks the truth, as I argue it should be Zalabardo (2012: pp. 45–49). Low levels of sensitivity (high values for pr(Bel(p)| ∼p)) will of necessity result in low tracking ratios, since pr(Bel(p)|p) is at most 1. Low levels of adherence, by contrast, are always in principle compatible with high tracking ratios, provided that they are accompanied by sufficiently high levels of sensitivity (sufficiently low values for pr(Bel(p)| ∼p)).

What matters for our purposes here is that by shifting from high absolute levels of adherence and sensitivity to a high tracking ratio as our necessary condition for knowledge by tracking we circumvent the difficulties faced by Roush's position with generalism. Consider first the palaeontologist's belief that dinosaur Bob was a vegetarian. The possibility that she might have been a carpenter, died in a car crash, etc. won't affect the tracking ratio of her belief, since it has the same deflating effect on its denominator as on its numerator. Both values will be lowered, as Roush argues, but there's no reason to think that their ratio will change. A similar situation obtains with respect to the second problem, concerning the swamping effect of scenarios in which the subject is not even present. Once again, the problem disappears as soon as we move from requiring absolute levels of adherence and sensitivity to concentrating on the ratio of the two. For the swamping effect of these scenarios can be expected to affect the denominator of the tracking ratio as much as its numerator. Both will be lowered, in principle to the same degree, e.g. by the possibility that the subject was not even present.

In sum, generalism, combined with Roush's account of truth tracking in terms of absolute values of adherence and sensitivity, yields counterintuitive results. But when we combine generalism with my account of truth tracking in terms of the tracking ratio, the problem disappears. I conclude that the difficulties have to do, not with generalism, but with Roush's specific version of probabilistic tracking. The version I favour has the resources for overcoming this manifestation of the problem of generality.

4 Justificatory internalism

Brandom uses the label *justificatory internalism* for the view that justification is the only knowledge-making feature—the only feature that can turn a true belief into

knowledge. As Brandom explains, justification cannot be taken to require that the belief was formed as a result of a process of inference (Brandom 2000: p. 211). On any plausible version of justificatory internalism, a non-inferential belief will count as justified so long as the subject is able to produce suitable justifications for it.

Reliabilism is in direct conflict with justificatory internalism. According to reliabilism, justification is not the only knowledge-making feature:

> true beliefs can, at least in some cases, amount to genuine knowledge even where the justification condition is not met (in the sense that the candidate knower is unable to produce suitable justifications), provided the beliefs resulted from the exercise of capacities that are *reliable* producers of true beliefs in the circumstances in which they were in fact exercised (Brandom 2000: p. 97).

Brandom refers to this claim as the *Founding Insight* of reliabilist epistemologies. Brandom's attitude towards the Founding Insight has two components. On the one hand, he accepts that there are cases in which a true belief qualifies for the status of knowledge by virtue of the fact that it has been reliably formed, even though the justification condition is not satisfied. At the very least, he thinks that this kind of case is "coherent and intelligible" (Brandom 2000: p. 105). On the other hand, he thinks that cases of this kind are in principle exceptional:

> The proper lesson to draw from the tension involved in the sorts of examples of knowledge to which the Founding Insight draws our attention, I think, is not that those examples are incoherent but that they are in principle exceptional (Brandom 2000: p. 106).

Let me refer to the following as *Brandom's exceptionality claim*:

> Cases in which a true belief qualifies for the status of knowledge by virtue of the fact that it has been reliably formed, even though the justification condition is not satisfied, are in principle exceptional.

Brandom's exceptionality claim can be seen as reducing the appeal of reliabilist accounts of knowledge. Reliabilists will typically want to go beyond the claim of the Founding Insight that reliable formation is the knowledge-making feature in some instances of knowledge. Their claim is that reliable formation is the *only* knowledge-making feature:

REL: Reliable formation is the only feature that can confer on a true belief the status of knowledge.

As Brandom explains, REL doesn't entail that adequate reasons can't confer on a true belief the status of knowledge. What it does entail is that knowledge by adequate reasons should be construed as a special case of knowledge by reliable formation: "Accepting only beliefs one *could* give reasons for […] is, under many circumstances, a reliable technique of belief formation" (Brandom 2000: p. 99).

Brandom's exceptionality claim would offer the justificatory internalist a strategy for resisting REL. She could argue that, given that the cases of knowledge that justificatory internalism can't explain are in principle exceptional, the Founding Insight offers

no support for the adoption of REL. All it supports is the replacement of full-strength justificatory internalism with a mitigated version of the view:

MJI: Justification is the feature that turns most true beliefs into knowledge. There may be also true beliefs that are turned into knowledge by being reliably formed, but these are in principle exceptional.

In the presence of Brandom's exceptionality claim, the justificatory internalist will be able to claim that MJI is more faithful than REL to the spirit of the Founding Insight. My main goal in the present section is to argue that Brandom hasn't provided a cogent line of reasoning in support of his exceptionality claim.[6]

Brandom seeks to derive the exceptionality claim from his views about the intrinsic role played in the concept of belief by the ability to produce reasons:

> It is at the very least unclear that we can make sense of a community of believers who, while often holding true beliefs, and generally acquiring them by reliable mechanisms, *never* are in a position to offer reasons for their beliefs (Brandom 2000: p. 107).

By contrast,

> we *can* make sense of a community whose members formed beliefs only when they thought they had justifications for them (Brandom 2000: p. 106).

What Brandom's arguments purport to show in the first instance is that belief for which the subject is not in a position to offer reasons is an in principle exceptional phenomenon.

Hence Brandom's argumentative strategy is this: he wants to establish that the kind of case that motivates the Founding Insight is in principle exceptional. And he seeks to derive this conclusion from a lemma to the effect that something else—belief for which the subject is not able to produce reasons—is in principle exceptional. The cogency of his argument at this point will depend on the relationship between the phenomena to which in-principle exceptionality is attributed by his premise and his conclusion. Both phenomena can be usefully clarified.

Consider first the kind of case that supports the Founding Insight. I want to suggest that in order to support the Founding Insight, a belief has to satisfy the following conditions:

A. We are intuitively inclined to ascribe to it the status of knowledge.
B. It has been reliably formed.
C. It doesn't exhibit the feature that, according to justificatory internalism, turns a true belief into knowledge.

I'm going to refer to beliefs satisfying conditions A–C as *JI-counterexamples*. I am proposing that in order to establish the exceptionality claim, Brandom needs to show that beliefs satisfying conditions A–C are in principle exceptional.

[6] Rejecting Brandom's exceptionality claim removes one obstacle to the adoption of REL, but it doesn't force us to adopt this view. Someone who rejects the exceptionality claim could still hold that there are instances of knowledge that can be explained only by the justification condition. Roush accepts this for deductive reasons (Roush 2005). I accept it for any kind of reasons (Zalabardo 2012).

Consider now the kind of belief that figures in his premise—beliefs for which the subject is not able to produce reasons. If S believes p, we can say that S is *able to produce reasons* in favour of p just in case there is a proposition q satisfying the following conditions:

1. S believes q.
2. S believes that q provides sufficient support for p.

Let's say that S's belief that p is *groundless* when S is not able to produce reasons for p in this sense—there is no proposition q satisfying conditions 1 and 2 with respect to her belief that p. What Brandom's arguments concerning the role of reasons in the concept of belief would show, if successful, is that groundless beliefs are in principle exceptional.

Hence Brandom seeks to derive the conclusion that JI-counterexamples are in principle exceptional from the premise that groundless beliefs are in principle exceptional. In the remainder of this section I'm going to argue that the conclusion doesn't follow from the premise. Even if we concede to Brandom that groundless beliefs are in principle exceptional, we are not entitled to conclude from this that JI-counterexamples—beliefs that acquire the status of knowledge through reliable formation even though they don't satisfy the justification condition—are also in principle exceptional. The reason is that groundless beliefs are not the only kind of JI-counterexample.

The prospects for Brandom's defence of his exceptionality claim clearly turn on the relationship between groundlessness, on the one hand, and the knowledge-making feature postulated by justificatory internalism, on the other. In one direction, the connection is straightforward: groundless beliefs cannot exhibit the knowledge-making feature postulated by justificatory internalism. However, this isn't the direction that matters for Brandom's argument. What Brandom needs to assume is this:

JIG: Beliefs can fail to exhibit the knowledge-making feature postulated by justificatory internalism *only if* they are groundless.

With JIG in place, the result that groundless beliefs are in principle exceptional would entail that beliefs satisfying the knowledge-making feature postulated by justificatory internalism are in principle exceptional, and, a fortiori, that the same goes for JI-counterexamples. But in the absence of JIG, the in-principle exceptionality of groundless beliefs would have no obvious consequences for JI-counterexamples.

There can be no doubt that Brandom's reasoning relies on JIG. But is JIG true? I'm going to argue that this question needs to be answered in the negative—on any plausible construal of justificatory internalism. Being able to provide reasons for p, in our sense, is a necessary condition for the knowledge-making feature postulated by justificatory internalism. However, and this is the crucial point, it's nowhere near sufficient. The existence of a proposition q satisfying conditions 1 and 2 with respect to S's (true) belief that p can't be taken as conferring on the belief the status of knowledge. q would also have to satisfy other substantial conditions.

First, in normal circumstances, q would have to be true.[7] Suppose that I see my neighbour in her garden drinking from a glass, I form the true belief that the liquid in the glass is water and my only reason for this is my belief that she is abstemious. Suppose, however that I'm wrong in this: my neighbour drinks like a fish; this is the first glass of water she's drank in years. In this case I am able to offer a reason for my belief that she is drinking water, but because my reason is false I shouldn't count as satisfying the justification condition for knowledge.

Second, it would have to be true that q provides sufficient support for p. Suppose that I believe truly that the postman has come because my neighbour tells me that he has. I believe that my neighbour is truthful and reliable, but she is in fact a compulsive liar. She's told me that the postman has come only because she mistakenly believes that he hasn't come. Once again I'm able to offer a reason for my belief—that my neighbour says so—and in this case my reason is true, but I shouldn't count as satisfying the justification condition for knowledge, since my belief that my reason provides adequate support for my belief that the postman has come is false.

Third, my belief in my reason should have some positive epistemic status. Suppose again that I believe truly that the liquid in my neighbour's glass is water, and that, as before, my reason for this is my belief that my neighbour is abstemious. Now change the case a bit and assume that she is indeed abstemious, but that I only believe this because that's what the tea leaves at the bottom of my cup appear to indicate. Now I have a true reason for believing that there's water in her glass, and this reason provides adequate support for my belief, but I shouldn't count as satisfying the justification condition for knowledge because my belief in my reason lacks positive epistemic status.

Forth, although there is some controversy about this, I maintain that my belief that my reason provides adequate support for my belief should also have some kind of positive epistemic status.[8] Suppose I believe truly that the substance is an acid because I've seen the litmus paper turn red, but I have no information about the behaviour of litmus paper from any reputable sources, and I only believe that red litmus paper indicates acidity because that's what the tea leaves tell me. Now I have a true reason that truly supports my belief, and my belief in the reason has positive epistemic status, but my belief that the substance is an acid shouldn't count as satisfying the justification condition, because my belief that my reason provides adequate support lacks positive epistemic status.

In sum, I would argue that S's belief that p satisfies the justification condition for knowledge just in case there is a proposition q satisfying the following conditions:

1. S believes q.
2. S believes that q provides sufficient support for p.
3. q is true.
4. q provides sufficient support for p.
5. S's belief that q has positive epistemic status.

[7] But see Warfield (2005) for possible exceptions to this principle.

[8] For criticisms of this condition, see Alston (1980: p. 144). This condition is sometimes taken to generate an infinite regress. See, e.g. Boghossian (2001). I have argued elsewhere that this regress is not unavoidable (Zalabardo 2011: p. 127).

6. S's belief that q provides sufficient support for p has positive epistemic status.[9]

I can now state the problem that I want to raise for Brandom's argument in support of his exceptionality claim—that JI-counterexamples are in principle exceptional. The main premise of the argument is that beliefs with respect to which no proposition satisfies conditions 1 and 2 (groundless beliefs) are in principle exceptional. However, showing that JI-counterexamples are in principle exceptional would involve showing that beliefs with respect to which no proposition satisfies conditions 1–6 are in principle exceptional. But the former does not entail the latter. Even if, as Brandom claims to have shown, for most beliefs there must be a proposition satisfying conditions 1 and 2, this places no limit on how many beliefs there might be for which no proposition satisfies conditions 1–6. Brandom has failed to show that the instances of knowledge that can be explained by reliabilism but not by justificatory internalism are in principle exceptional. Hence the justificatory internalist cannot legitimately invoke Brandom's exceptionality claim to resist the transition from the Founding Insight to REL.

I'd like to end this section by emphasizing the limited scope of my argument against Brandom's exceptionality claim. Brandom's ultimate concern in the debate between internalism and reliabilism is to resist drawing from the Founding Insight a conclusion to which he refers as a *recentering of epistemology*:

> that [...] the concept of reliability of belief-forming processes can simply *replace* the concept of having good reasons for belief—that *all* the explanatory work for which we have been accustomed to call on the latter can be performed as well or better by the former (Brandom 2000: p. 100).

Brandom presents his exceptionality claim as a reason for resisting this recentering of epistemology. If the exceptionality claim were correct, there would be one explanatory task that reliability couldn't perform as well or better than the concept of having good reasons—the task of distinguishing knowledge from mere true belief. I have argued that this line of resistance against the recentering of epistemology is unsuccessful.

However, Brandom is very clear that he doesn't see this as his strongest line of reasoning against the recentering of epistemology. He attaches more importance to another line of reasoning concerning a different explanatory task—distinguishing belief "from the sorts of reliable indication exhibited by reliably indicating artifacts such as measuring instruments" (Brandom 2000: p. 110). For this task, Brandom thinks, reasons are indispensable, since "an essential element of that distinction is the potential role as both premise and conclusion in reasoning (both theoretical and practical) that beliefs play" (Brandom 2000: p. 109).

I have argued elsewhere that the concept of having good reasons is not required for this explanatory task either (Zalabardo 2016). However, my main concern here is to emphasize that the argument in this section doesn't affect this line of reasoning. As far as my argument here goes, Brandom's preferred reason for resisting the recentering of epistemology stands unchallenged. My target has been restricted to his other line of reasoning—concerning the role of the concept of having good reasons in the task of distinguishing knowledge from mere true belief. I have argued that Brandom has not

[9] For a more detailed account of inferential knowledge along these lines, see Zalabardo (2012: pp. 87–98).

offered cogent support for his claim that reliability cannot replace good reasons as our main tool for this this particular task—the central task of the analysis of knowledge.

5 Belief, truth and knowledge

Reliabilist and truth-tracking analyses of knowledge offer representationalist accounts of knowledge ascriptions that aspire to respect naturalistic principles. They seek to explain knowledge ascriptions as expressing the presence in a belief of a natural property. As Brandom suggests, many advocates of these analyses see them as applying to the concept of knowledge the methodological approach with which they also expect to deal with the concepts of belief and truth (Brandom 2000: p. 111). According to this programme, we would explain belief ascriptions by identifying the natural property that a subject has to exemplify in order to have a belief with a certain content. We would then explain truth ascriptions by identifying the natural property that a belief would have to exemplify in order to count as true. Reliabilist and truth-tracking epistemologies would provide the next step in this programme of representationalist naturalization. Brandom's anti-reliabilist arguments seek to establish that even on the assumption that the approach works for belief and truth we have to reject the idea that it will work for knowledge.

I wouldn't dispute that many advocates of reliabilist and truth-tracking epistemologies subscribe to this overall programme of representationalist naturalization.[10] What I want to suggest, in closing, is that endorsement of these analyses of knowledge doesn't carry a commitment to a parallel treatment of the concepts of belief and truth. I am claiming that the following is in principle a perfectly consistent combination of views: a representationalist account of knowledge along the lines of reliabilism or truth tracking coupled with a pragmatist account of belief and truth.

A position answering to this description emerges as a possibility from our discussion in the previous section. On this position, belief and truth would receive Brandom's reason-based pragmatist explanation, while knowledge would be explained in terms of reliability. The position that I favour exhibits a similar combination of views. On the one hand, belief and truth are explicated in terms of the rules that define the practice of alethic assessment, and the associated interpretative practice of using the sentences of our language to index the mental states of others (Zalabardo 2016). This approach shares the pragmatist methodology of Brandom's inferentialism, but, unlike in Brandom's approach, the rules that are presented as defining the practice don't invoke the role of beliefs as potential premises and conclusions. On the other hand, the difference between knowledge and mere true belief receives a representationalist explication—using objective conditional probabilities to specify the conditions that a true belief has to satisfy in order to qualify for the status of knowledge. Objective probabilities are deployed in this account in two parallel but separate ways to explicate non-inferential and inferential knowledge. Non-inferential knowledge is explicated in terms of probabilistic tracking, developing the ideas sketched in Sects. 2 and 3 of the present paper. Inferential knowledge is explicated in terms of the concept of adequate

[10] For the idea of non-representationalist naturalism, see Price (2013).

evidence, but this concept is also elucidated in terms of conditional probabilities. Unlike other naturalist accounts of knowledge, this proposal doesn't seek to reduce inferential knowledge to non-inferential knowledge—a true belief can have the status of knowledge because the subject possesses adequate evidence in its support even if the belief doesn't track the truth (Zalabardo 2012: pp. 63–66). In this way the concept of having good reasons is still assigned an essential role in the explication of knowledge.

Needless to say, a pragmatist account of belief and truth will apply to all regions of discourse—not only to ascriptions of the belief, say, that the sun is shining and to ascriptions of truth to this belief, but also to ascriptions of beliefs of the form S knows that p and to ascriptions of truth to beliefs of this form. In this way, knowledge, no less than any other concept, will receive a pragmatist construal at one remove. However, this is in principle compatible with a characterisation of the specific rules that govern epistemic discourse according to which the truth of beliefs ascribing knowledge is inextricably linked with the truth of beliefs ascribing certain natural properties. Acknowledging this link would be tantamount to embracing a naturalist representationalist account of knowledge—i.e. of epistemic beliefs and their truth—within a global pragmatist approach to belief and truth in general.

6 Conclusion

In this paper I have focused on some aspects of Brandom's discussion of naturalist accounts of knowledge. First I have taken issue with his attempt to show that a version of the generality problem invalidates all attempts to construe the concept of knowledge in terms of objective probabilities. I have argued that probabilistic-tracking accounts have the resources for overcoming the obstacle. Then I have considered Brandom's argument for the claim that cases in which a belief acquires the status of knowledge by being reliably formed, even though it doesn't satisfy the justification condition, are in principle exceptional. I have argued that Brandom's reasoning fails, since the justification condition can be easily violated in cases besides those that he treats as exceptional. Finally, I have argued, more briefly, that adopting the kind of naturalist representationalism about knowledge that emerges unscathed from this debate doesn't force us to embrace a form of naturalist representationalism about belief and truth. Hence my defence of naturalist representationalism about knowledge from Brandom's attacks puts no pressure on pragmatist accounts of belief and truth—either on Brandom's inferentialism or on the non-reasons-based version of the position that I favour.

Acknowledgements I am indebted to two anonymous referees for this journal.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

Alston, W. P. (1980). Level-confusions in epistemology. In P. French, T. Uehling Jr., & H. K. Wettstein (Eds.), *Midwest studies in philosophy: Studies in epistemology*. Minneapolis: University of Minnesota Press.

Boghossian, P. (2001). How are objective epistemic reasons possible? *Philosophical Studies*, *106*, 340–80.

Brandom, R. (1994). *Making it explicit*. Cambridge, MA: Harvard University Press.

Brandom, R. (2000). *Articulating reasons: An introduction to inferentialism*. Cambridge, MA: Harvard University Press.

Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, *89*, 1–29.

DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, *104*, 1–52.

Feldman, R. (1985). Reliability and justification. *The Monist*, *68*, 159–74.

Glynn, L. (2010). Deterministic chance. *British Journal for the Philosophy of Science*, *61*, 51–80.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, *73*, 771–91.

Goldman, A. I. (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel.

Goldman, A. I. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard University Press.

Hoefer, C. (2007). The third way on objective probability: A sceptic's guide to objective chance. *Mind*, *116*, 549–96.

Lewis, D. (1986). A subjectivist's guide to objective chance. In *Philosophical papers*, Vol. II. Oxford: Oxford University Press.

Lewis, D. (1994). Humean supervenience debugged. *Mind*, *103*, 473–90.

Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.

Price, H. (2004). Naturalism without representationalism. In M. De Caro & D. Macarthur (Eds.), *Naturalism in question*. Harvard: Harvard University Press.

Price, H. (2013). *Expressivism, pragmatism and representationalism*. Cambridge: Cambridge University Press.

Roush, S. (2005). *Tracking truth*. Oxford: Oxford University Press.

Warfield, T. A. (2005). Knowledge from falsehood. *Philosophical Perspectives*, *19*, 405–16.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Zalabardo, J. L. (2011). Boghossian on inferential knowlege. *Analytic Philosophy*, *52*, 124–39.

Zalabardo, J. L. (2012). *Scepticism and reliable belief*. Oxford: Oxford University Press.

Zalabardo, J. L. (2016). Empiricist pragmatism. *Philosophical Issues*, *26*, 441–61.