

Chapter 37

Perceptual Justification—Two Conceptions Compared



Crispin Wright

37.1 Introduction—Quasi-libertarianism about epistemic value

Without making any judgement about whether any of them may be reducible to or explicable in terms of others, we have, on the face of it, a *plurality* of values of action which rank as ethical in a broad sense of the term: it is, necessarily, a consideration in favour of judging that someone acts well if they act in a way that realises others' needs, furthers others' flourishing, or is fair, or just.... But irrespective of the results of their action—its actual effects—it is also a consideration in favour of their having acted well—or at least better than otherwise,—if they acted out of a *concern* to do well in these respects, even if in the event unsuccessfully. Broadly, there are ethical values of acts and ethical values of agency.

It doesn't seem quite right, however, that all the latter—ethical values of agency—can be accounted for in terms of an agent's concern to achieve the former. There are, on the face of it, values of agency that do not seem quite to fit that bill. It is, for example, also a consideration in favour of judging someone to act well that they act with *integrity*, or *bravely*, or *compassionately* in circumstances where it is clear to all that there is nothing to be done to help the object of compassion.

A central question underlying what follows is whether a similar distinction applies to our *doxastic* agency—to our formation and revision of beliefs and to our general management of our respective systems of belief. Notwithstanding the familiar point that believing/disbelieving are not states that we can freely initiate in whatever sense it is that we can freely initiate actions, we are familiarly very ready to appraise others' management of their beliefs in a vocabulary of Strawsonian reactive attitudes: appraisals such as that someone did well to realise that P, ought to have known that

C. Wright (✉)
University of Stirling, Stirling, UK
e-mail: cjgwright@gmail.com

Q, ought not to continue to hold such and such views, that they have failings of bigotry, superstition, cynicism, or gullibility which—it seems to be implied—they should do something about. It seems that, at least in our intuitive thought, we do employ roughly the same kind of pluralistic value scheme for belief as for agency. We have a plurality of values of belief ranking as ethical in a broad doxastic sense: it is necessarily, a consideration in favour of judging that someone believes well if they believe truly, better yet knowledgeably, better yet in a way that enhances understanding, or if their belief is well grounded, or articulates their recognition of a commitment of others of their beliefs. But irrespective of such actual properties of a belief, it is also a consideration in favour of someone's doing doxastically well¹—or at least better than otherwise,—if they conduct their believings out of a concern to do well in such respects, even if in the event unsuccessfully. A subject is doing doxastically well insofar as they manage their beliefs in a way that reflects concern with these epistemic values, even if in the event they do so, on an occasion, or even regularly unsuccessfully. And, as with action, so with belief, it is arguable that there are values of doxastic agency which cannot be readily construed as derived from the value of concern with what we may think of as primary values of belief. Rationality is not the clearest of notions, but it incorporates elements of proportionality, balance and coherence that are not obviously merely instrumentally valuable in the pursuit of knowledge and truth.

We can term this conception of the structure and variety of the values that govern belief *quasi-libertarian*. I reserve for another occasion the question whether the whole quasi-libertarian way of thinking about believing is *au fond* deeply unsatisfactory. However, my overarching assumption in what follows will be that it is not. And if it is not, any satisfactory account of the values that engage believing must make good sense of an analogue of the distinction between, on the one hand, properly considered, well-motivated and well chosen action and, on the other, action possessed of values which it could have quite independently of the motivation for which it is performed. Such a distinction would seem to be naturally captioned by a contrast between justification and truth. It accordingly goes with quasi-libertarianism that we make sense of how a justified belief can be false.

My question is whether or how we can do this for perceptual beliefs—roughly, beliefs delivered to us just by perceiving. I shall consider first an internalist and then an externalist account of how to return a positive answer.

¹ Martin Davies' useful phrase.

37.2 The Internalist Q-I (Quasi-Inferential) Model of Perceptual Justification and a Challenge to It²

According to any form of perceptual internalism, as I shall here understand the notion, doxastic justification is conferred on a (perceptual) belief by other (narrow) phenomenal states of the subject. I am—*ceteris paribus*—justified in believing that P by having a relevant kind of experience, E: that is, by having it (apparently) sense-experientially seem to me that P—by having things look to me as if P, or sound as if P, or feel as if P... or some combination of appropriate apparently sensory inputs. Justification then proceeds along quasi-inferential lines:

$$\frac{E}{P} \quad \downarrow$$

In forming a (presumptively) perceptually justified belief, I transition from a state of apparent perceptual experience to an appropriate conclusion about the external world. The possibility of the belief's being justified but false is very straightforwardly secured by the defeasibility of the transition.

There are various choice points in spelling out this idea. First, it will be literally an inferential transition that is involved if one thinks of E not as the experience itself but as a propositional expression of the character in relevant respects of that experience. It will, in contrast, be merely a 'quasi-inference' if E is conceived instead as the experience itself.

The former choice is not a good idea if one wishes to extend the capacity for perceptual justification of belief to thinkers who have no ability to entertain propositions about their phenomenal experience. However on the latter account, E will still have to be given conceptual shape before it can engage supportively with P—it will have to be identified as an experience whose phenomenal character is that of its-seeming-that-P and identified as such by the subject if they are rationally to avail themselves of its evidential payload.

A second choice point opens there. We may think of the experience E, in the style of *Mind and World*, as already itself a conceptual event: when a subject undergoes a perceptual experience, certain conceptual capacities are "passively" drawn upon. Just as entertaining an occurrent thought that P not merely presupposes the subject's possession of the conceptual components of P but actually involves exercising them, so too in (presumptively) perceptual experience—in undergoing an experience as of its perceptually seeming that P—the same capacities are, on the *Mind and World* view, drawn upon as in episodically thinking that P.

The alternative is to think of the experience, *qua* event, not as actually drawing on the conceptual capacities of the subject but as conceptual only on the sense of being *conceptualisable*: of having a character which allows of conceptual articulation but whose presence to the mind, fully formed as it were, is antecedent to its being given

² This section is heavily indebted to the work of Jim Pryor.

conceptual shape. It will still be true, even so, that it is only as so shaped that an experience can enter into justificatory relations.

Whatever view is taken of these tricky issues, there is a difficulty with the internalist Q-I model, at least as so far sketchily outlined. It is surely way too permissive. Which are the propositional contents, P, that we should count as apt for *perceptual* justification? Obviously a very wide class of contents are justifiable for which perception is but one—albeit an indispensable—component in their justificational architecture. What we had in mind in raising the question of perceptual justification was surely a narrower class of propositional contents for which an episode of perception can supply an, in some sense, *complete* justification. Indeed cases where perception justifies only as augmented by collateral information and background theory themselves invite us to try to extricate the kinds of more purely perceptual nuggets of information that, in such cases, perception specifically contributes to their overall support. It is such nuggets that should constitute the range of ‘P’ as schematised in the Q-I model.

So which, if they are not chimerical, are the contents whose justification can in principle be provided purely by perception, without any contribution from auxiliary information? The question seems to raise something like the challenge to make out that notion of an *observation statement* which, in the wake of the simplistic ideas of the logical positivists about the relationship between empirical scientific theory and sense experience, provoked so much soul-searching in mid-C20 philosophy of science. Still, without setting out on that wild goose chase all over again, perhaps something can be said to corral a range of judgements for which the simplicity of the Q-I model of perceptual justification might seem to have some justifiable appeal.

To begin with, let’s term any judgement *recognitionial* which is expressed by a sentence composed just of one or more demonstratives coupled with predicates and relations of arbitrary degree, and whose competent assessment would normally demand (besides possession of the ingredient concepts and appropriate attentive capacities) no more than ordinary perceptual faculties and routine conditions of observation. ‘That is red’, ‘That is deeper than that’ (said of sounds), ‘That is salty’, etc., all qualify. So, however, may judgements expressing recognition of natural kinds and artefacts: “This is a cedar tree”. “That is a tyre pump” and so on for pineapples, bananas and tomatoes, motorcars, Geiger counters, lawn mowers, trousers, televisions and smartphones. Can we trim the list in a principled way?

Maybe. Let F be any predicate capable of featuring in recognitionial judgements. F’s application to a demonstratively given object can thus be recognised by unaided perception. Now stipulate that F meets two additional conditions.

1. No one counts as understanding F who is unable, even when perceiving normally in normal circumstances, competently to appraise recognitionial judgements involving predication of F. Understanding F is to *require* an associated perceptual-recognitionial competence.

This condition excludes expressions for artefactual concepts like *lawn mower* and *Geiger counter* whose instances, while perceptually recognisable, could in principle take any of indefinitely many various shapes and overall appearance. It is thus no

essential part of understanding such concepts to have the capacity to handle recognitional judgements involving them; you can know exactly what tyre pumps are without knowing anything about the gross appearance examples of such devices typically assume.

The corresponding claim about *pineapple*, *cedar tree*, and indeed the general run of natural kind terms is however less plausible. Intuitively, no one *fully* understands those concepts who is innocent of the respectively distinctive appearances of pineapples and cedars. But in these cases appropriate recognitional capacities, while arguably necessary for understanding, are consistent with misunderstanding. You can know what pineapples characteristically look, smell and taste like while not knowing that they are a kind of fruit—indeed while innocent of the very concept of *fruit*—or while confusing them with a type of look-alike whose DNA is quite different. What distinguishes these concepts from those in the narrower, more purely perceptual class we seek to characterise is that such recognitional capacities are not sufficient for full understanding. We should therefore stipulate a second condition:

2. It's not possible to have the relevant recognitional capacity for Fs yet lack an understanding that something's being an F is not exhausted by the features thereby recognised.

Let us call a recognitional judgement whose ingredient concepts meet both conditions 1 and 2, *purely* recognitional. The stark simplicity of the Q-I model requires that, in a good case, the occurrence of E be, unsupplemented, rationally sufficient for the judgement that P. That in turn requires that the justification of the transition from E to P not depend on collateral information, unrepresented in the model. My suggestion is accordingly that the best case for the Q-I model requires that P be *purely* recognitional in the sense outlined. The two conditions, 1 and 2, collectively require that possession of appropriate recognitional capacities be both necessary and sufficient for understanding a concept of the kind required: the former rules out contingency in the relationship between F's and any characteristic appearance they have—you need to know how Fs characteristically appear to understand what it is to be F—while the latter rules out that, although Fs may have a characteristic appearance, there be more to being F than having that appearance: something which someone who merely masters the appearance may be ignorant of. There is to be no such additional component in qualifying as F.

37.3 A Second Challenge for the Q-I Model: Authenticity-Conditions, Lemmas and Props

To be sure, I have provided no demonstration that there are indeed any purely recognitional judgements, in the aforementioned sense, whose ingredient concepts meet both stipulated conditions. Nor have we considered, even if there are, which if any types of judgement besides such purely recognitional judgements might in turn be apt for immediate perceptual justification after the fashion schematised in the Q-I

model. But let's assume that, at least as a start, the foregoing suggestion is not totally misguided. There is a second challenge for the justificational architecture represented in the model, even when we consider only purely recognitional judgements.

We need a bit of apparatus to bring the challenge to the fore. First, we require a very general notion of a *cognitive project*. Such a project may be thought of as a pair, $\langle P, M \rangle$, consisting of a thesis, that P , and a method M for forming a view about it, for example gathering relevant first-hand evidence, seeking competent testimony, looking and seeing whether P , bringing memory to bear, seeking a contradiction spawned by supposition of P 's negation—anything such that, once having implemented it, and depending on the upshot, one could reasonably form a view on that basis concerning whether P was true or not.

Next we need a fairly standard notion of an *underminer* for $\langle P, M \rangle$. Such an underminer is any proposition expressing a condition, U , such that reason to think U obtained would be reason to doubt the significance with respect to P of the upshot of implementing M —any condition whose obtaining should rationally be regarded as giving one reason to think that one's execution of M was compromised in such a way as to some degree disqualify the epistemic significance of its results. Undermining defeat contrasts, as familiarly, with overriding defeat. Should $\langle P, M \rangle$ conclude in an apparent justification for P , and then an underminer for the project emerge as true, one will have reason to weaken one's confidence in the project's upshot, perhaps even to discount it altogether, but not, or not necessarily, to increase one's confidence in the negation of P . An overriding defeater, by contrast, provides exactly that but may not provide any reason to question the *bona fides* of the support, as far as it went, provided by the executed project for P .

Next define an *authenticity-condition*, A , for M with respect to P as an anti-underminer,³ that is, the negation of an underminer for M with respect to P . The authenticity-conditions for M with respect to P thus comprise every circumstance whose obtaining is necessary if the result of implementing M is to amount to a finding in good standing for or against P .

Now ask, what in general should be the attitude of an agent to something that they recognize as being an authenticity condition for a cognitive project they go in for? It depends. Authenticity-conditions divide into two kinds. Some are *lemmas*: the rational attitude to these is that positive evidence of their satisfaction is required before it can be justified to repose trust in the outcome of the project. If the project is, for example, a scientific experiment involving special apparatus and calculations on the basis of the measurable results it generates, ordinary canons of good scientific method will require the experimenter to have antecedent reason to believe the apparatus is up to the necessary standards, and to gather reason to believe that the relevant readings have been taken accurately and that the relevant calculations are error-free. And such reasons will be provided, for example, by the reputation of the manufacturers, and by routine due diligence being exercised with the taking of readings and checking of calculations.

³ The term is Jim Pryor's (2012).

However we don't always regard authenticity conditions as lemmas. Sometimes we instead regard them as what we may term *props*. If an authenticity condition, A, is a *prop* for a project $\langle P, M \rangle$, it will generally be reckoned to be acceptable to, as we say, take its satisfaction for granted—indeed to give the matter no thought in one's conduct of the project. That, for example, there has been no conspiracy on the part of the manufacturers of the physics lab apparatus to promote misbegotten schoolroom experimentation, that one has not unwittingly ingested an hallucinogen in one's morning coffee which induces impressions of cognitive and perceptual lucidity while one mistakes the final readings and commits egregious errors of computation,—such assumptions would normally be regarded as props: as things whose failure would undermine the project but of which normal standards of epistemic responsibility require no special policing or consideration.

The authenticity conditions even of simple projects like looking at one's watch to determine the time of day are ramified. Proper functioning of the instrument will require multiple conditions of mechanical detail, there will have to have been a sometime correct setting, and no subsequent mis-setting, of the hands. Then there are conditions governing the reliability of one's observation of the watch face, recollection of the significance of the positioning of the hands on a watch, and general cognitive lucidity.

Where is this going? Well it should be obvious on reflection that one could not rationally, in any feasible cognitive project, treat *every one* of its authenticity conditions as a lemma. The reason is simply that the required validation of a lemma, qua lemma, will constitute a discrete cognitive project with its own authenticity conditions. So when A is regarded as a lemma, the options are either that one nevertheless merely takes it for granted—a compromise of epistemic responsibility—or one faces the obligation to launch a secondary cognitive project to plug the gap, as it were. But that secondary project will in turn throw up its own authenticity conditions. If they are all supposed to be lemmas too, a tertiary group of required cognitive projects will then be spawned, with in turn their own authenticity conditions, all of which are by hypothesis lemmas.... We must conclude that any feasible conception of epistemic responsibility must allow a rational subject who implements any cognitive project to treat some of its authenticity conditions as props—to act on the unevicenced assumption that they are satisfied. What can be said about which?

There is nothing in the preceding reasoning to require that, for any particular project, $\langle P, M \rangle$, the status of any of its authenticity conditions as lemma or prop respectively is absolute, fixed purely by the proposition and method concerned. It is to be expected that the matter will depend on the investment of the agent in getting a reliable result—in general, how careful you need to be in any project, cognitive or otherwise, depends on how much it matters whether you get a good or a bad outcome, (and perhaps also on standing, a-rational dispositions of risk aversion/tolerance.) However one plausible interpretation of the paradoxes of scepticism, is, I believe, that some very general—metaphysically “heavyweight”⁴—theses are both authenticity conditions for huge sweeps of cognitive projects and essentially beyond

⁴ The term originates in this context in the debate between Dretske (2005) and Hawthorne (2005).

the reach of independent enquiry everywhere, so are necessarily props if the projects they underwrite are to be reckoned rational. I'll come back to this perspective.

How does any of this bear on the internalist Q-I model of perceptual justification? Let P be a purely recognitional judgement and let E be an episode of sense-experience apt to justify P according to the model. Let A be an authenticity condition for the project of investigating P by the method of seeking an appropriate such experience. There will be, even in so basic a project—essentially just positioning oneself in a situation where P's truth or falsity can be expected to “shine through” in one's experience—a plethora of authenticity-conditions for the project. They will include the appropriateness of one's position, one's adequate cognitive focus on the issue at hand, appropriately directed perceptual attention, the proper functioning of one's perceptual apparatus, and the suitability of the prevailing circumstances to allow the F-ness, or otherwise, of the demonstrated object to manifest —“shine through”—in one's occurrent experience.

Do these reflections give us any reason to complicate the structure of the Q-I model? Certainly in any situation where there is some initial reason to doubt any of the relevant conditions, it would be inappropriate to represent a justification for P or its negation as grounded simply in the subject's awareness of the occurrence, or non-occurrence respectively, of E. But where there is no such extant reason to doubt, it is plausible that our actual practice is to treat the kind of authenticity conditions listed as props. All we have to do to get knowledge in such cases—so fully justified belief—is to “look and see”. And we are for example perfectly happy to grant children fully justified beliefs of the relevant kind simply on the basis of their experiences, even if they have no capacity to understand, let alone check up on, the authenticity conditions involved. So if the question is, what, in actual practice, besides their sense experience, do we require a subject to take account of in making justified purely recognitional judgements, the answer in general is “Nothing”. We are indeed, at a pre-theoretical level, Pryorian Dogmatists: the having of (apparently) perceptual evidence, conceived as an internal psychological state, is treated as providing default sufficient justification for basic perceptual beliefs. In such cases the authenticity conditions of perceptual evidence are, absent other relevant information, all treated as props.

That seems to be our actual practice. It is a different question whether this is the way we *ought* to think about the matter, of course—whether, indeed, we ought to think of the justifiers of the judgements concerned as internal states at all, and whether, if we do, we ought to think of the fashion in which they justify as satisfactorily captured by the simplicity of the Q-I model. On the latter point, the opposed *conservatism*, which I have defended in other work,⁵ holds that sense-experiential evidence justifies perceptual beliefs only within the setting of some kind of positive warrant for certain of the authenticity conditions involved. And, at least at the level of fully reflective justification, there is, as it seems to me, a powerful argument in favour of something in this direction.

⁵ Replace with: See Wright (2007) and (2012) at pp. 451–76.

The argument flows from the phenomenon of *Prop* \rightarrow *Lemma inflation*, something which emerges when we move from consideration of what we consider it takes for a subject to *form* perceptually justified belief and consider instead what a subject needs to be able to do in order responsibly to *lay claim* to possession of such justification. Simply, at the level of full reflection, it is very difficult to sustain the idea that anything is merely a prop. Suppose Q is merely a prop for the evidential bearing of experience E on P; and suppose a subject has E. If they have no other relevant information, and so in particular have no evidence counting against Q, they are, according to the Q-I model, justified in accepting P. But if the subject announces—*lays claim to*—possession of that justification, we can ask whether they take themselves to be in position to rely on Q. And the answer from someone who is precisely and knowingly assigning to E the bearing on P which Q underwrites, that “I do not need to have a view about that; it is enough that I have no reason to doubt Q” will always seem manifestly irrational.

This phenomenon is not confined to authenticity conditions for the perceptual justification of (purely) recognitional judgments but is pervasive. Here are two further examples to ‘pump intuition’:

Time:

- A: “Do you know the time?”
 B: “Yes, it’s 5.30 pm”
 A: “How do you know — you’re not wearing a watch.”
 B: “The clock in the tower over there reads 5.30 pm”
 A: “So you are taking it that your visual apparatus is functioning properly at the moment in conditions conducive to its effective operation?”
 B: **“No, I am making no such assumption. I do not need to have a view about that; it is enough that I have no reason to doubt it”**

Promotion:

- A: “How did Jones’ sudden promotion land with his immediate colleagues?”
 B: “They were jealous and annoyed”
 A: “How do you know — it’s not the kind of reaction people usually make obvious.”
 B: “Well, there were subtle and some not so subtle behavioural signs, including a very acidic conversation I overheard in the Common Room”
 A: “So you are assuming that the people’s reactive psychological states tend to be evidenced by what they say and do in unguarded moments?”
 B: **“No, I am making no such assumption. I do not need to have a view about that; it is enough that I have no reason to doubt it”.**

In these examples, B’s final response seems manifestly absurd and irrational. The proper responses to A’s penultimate questions is rather something like,

“Well yes, of course”, or “Well yes, don’t we all?”

At the level of *reflective claims* to the accomplishment of justification, then, it seems that if one is to offer a proper representation of the full justificational architecture

of perceptual judgements,—indeed of any judgements,—one must offer some kind of acknowledgement of the role of various props. I am not now going to argue but I suggest that some of these have to be at the level of generality of Hawthorne's (2005) and Dretske's (2005) “metaphysical heavyweights”. These, like all authenticity-conditions, will occupy, from a reflective standpoint, the role of preconditions for the idea that the outcome of a project is so much as germane to the particular judgement under investigation. Moreover, unless one turns to scepticism, one is required to believe that it is rationally acceptable to take such conditions for granted. Accepting them is something to which a fully rational reflective subject is committed in claiming justification for treating their evidence as corroborative of the judgement concerned.

Assuming it is correct, how should we react to this thought? The conclusion seems to be that at the level of reflection—which is of course the level of the discussion Descartes himself initiated, the level at which the problem of scepticism arises in the first place—all our quotidian claims to knowledge and justification are conditional on the obtaining of conditions that have not been corroborated by enquiry, and some of which, indeed, are beyond corroboration by enquiry, since they will feature as authenticity conditions for any further relevant enquiry. A kind of second order scepticism may seem to be suggested: we have proper justification for no such quotidian claims, but can only invest in the hope that the ultimate authenticity-conditions on which they rest are met.

Moore's notorious ‘proof’ encapsulates the problem. If the best possible evidence for any claim about the material world is drawn, as internalism has it, from the world of phenomenal experience, then the evidential significance that the Q-I model will assign to Moore's visual and bodily experiences as of handedness will rest on the authenticity-condition that there is so much as an external material world in the first place. But the only route that Moore's ‘proof’ offers to corroboration of that condition proceeds through the premise that “Here is a hand”. At the level of reflection, and on the assumptions of the Q-I model, the proof is accordingly question-begging. Its conclusion is something you need to acknowledge in advance in order to assign the desired evidential significance to the phenomenal experiences that ground its premise.

However there is a different perspective. It is that the lesson of the unavoidability of the acknowledgment, at the reflective level, of props in every enquiry should be not any form of scepticism—of which it is of the essence to see us as coming short of justificatory standards which we are right to admire and to wish to meet—but rather that the ideal of an enquiry in which all authenticity conditions are fully recognized and corroborated in their turn is simply an incoherent ideal, comparable to the idea of an ideal queue in which everyone is at the front, or an ideal competitive game in which there are only winners. Rather, it is of the essence of queuing that some are placed behind others, and of the essence of competitive games that if someone wins, others lose—and of the essence of full reflective enquiry to rest on assumptions that certain conditions obtain whose obtaining has not been checked out and, in some cases, is beyond the reach of certification by enquiry. This is nothing to lament but something to know. It is a Socratic lesson. Such lessons are one of the traditional benefits of Philosophy.

37.4 Externalism and False but Justified Perceptual Belief—The R-S Model⁶

Currently influential externalist conceptions of perceptual knowledge have tended to play down the very idea of epistemic justification that fails to eventuate in knowledge. Belief that falls short of knowledge is stigmatized as “botched knowledge”, and the sense in which such belief can be ‘justified’ can thus only be a compromised sense—a kind of mitigation of a substandard result.⁷ The idea of justified but false belief now becomes a kind of oxymoron: false belief cannot, on such views be justified in the same sense or as fully as the justification conferred on a belief by circumstances that make it knowledgeable.

It seems to me this stigmatisation of justified belief has more to do with the current fashion for ‘Knowledge First’ approaches to the philosophical elucidation of epistemic states and their various possible merits than with externalism in epistemology, that there is no evident reason why externalists should be hostile to the idea of false but fully justified belief. To be sure, it is not clear how best to characterise externalism as such. The obverse of the internalism we have hitherto been considering would be the idea that the justification of a belief is conferred on it by factors which are at least in part external to the mind of the believer—by whether it is caused, for example, in a way that will tend to generate such a belief only in circumstances which will make it true.

To fix ideas, consider what we may term a *brute reliabilist* version of externalism about knowledge: according to brute reliabilism, a belief is knowledgeable if and only if true and generated by a reliable method. Such a view does actually straightforwardly generate a notion of justification consistent with the possibility of a false belief’s being justified: a false belief may be regarded as justified if and only if generated by a generally *but not exceptionlessly* reliable method. To the objection to such an account that this is no genuine notion of *justification*—that it, for instance, makes no relevant connection with epistemic agency of the kind needed to underwrite something analogous to the notion of the justification of a moral agent who acts out of excellent motives to, in the circumstances, an unluckily negative effect—a theorist attracted to brute reliabilism may rejoin that such a connection is easily supplied: it will be enough that a subject manages their beliefs with a concern for the reliability of their methods at the forefront of their managerial decisions.

Such a notion of justification is in keeping with the idea of justification as *failed* (not “botched”) knowledge. To be justified in a false belief is to have pursued what one rightly regards as generally a knowledge-acquisitive method but to have been let down thereby in the case in point.

⁶ This section is significantly influenced by the work of Alan Millar. See Millar (2009, 2011).

⁷ Williamson (2000) at p. 47.

Arguably, however, this is not enough to get us something recognisable as a notion of *perceptual* justification, properly so termed. To see this, consider the myth⁸ of chicken-sexing. Suppose Fred can very reliably sort day-old chicks coming past him on a slow moving belt into males and females, although they look pretty much indistinguishable to ordinary observers and Fred himself can offer very little by way of explanation of how he is doing it. Suppose it turns out that he cannot do it if he is blindfolded. So it seems sight is somehow involved.

Fred is pursuing what it seems is, for him, a generally reliable method to distinguish the sexes of the chicks. His method is to look and see. Only he does not, in a sense, know how he is doing it. We might suppose that he has some kind of *gestalt* concept of a distinction that he finds it hard to articulate or explain to others. If that were the case, we can imagine him giving us samples—a small group of males and another of females—and saying things like, “There. Can you see the difference now?” But suppose it is not like that. Rather, he just looks and an opinion comes to him—as it were, pops into his head without his having any sense of why.⁹

Two questions arise. First, is Fred’s belief, “This one is a male” *justified*? Second, is he getting *perceptual* justification?

On the first, the answer that suggests itself is that if Fred’s belief is justified, it is so only and purely on the basis of his track-record of success—we would not regard it as justified if it comes early in the process of accumulating that record, and nor should he. Still, if he has been, as we are supposing, egregiously successful over a sustained period, he and we are surely entitled to regard his on-going verdicts as at least inductively justified.

So should we regard his inductively justified beliefs about the chicks, formed as they apparently are just by looking at them, as *perceptually* justified? I think we should want to deny that they are. The example forces a distinction, therefore, between

- (a) a belief formed just by perceptual means, but justified by independent considerations, and
- (b) a belief that is justified by perception.

To be clear what is at stake here, we need to pay attention to “just by perceptual means” and to set aside a different distinction which could casually be expressed in a similar way: the case where the aetiology of a justified belief—“That is F”—factors into two components: the apparent perception of an instance of a purely recognitional concept G, and background information that if an item is G, it is F. Were Fred discriminating the male chicks by relying on some *gestalt* concept G which he has learned is reliably correlated with their being male, that would call for our earlier distinction between a purely recognitional judgement—That’s a G—and a belief based in part on the former but additionally supported by collateral

⁸ In fact, I understand that there are a number of different techniques used by trained chicken sexers, depending on the breed of chicks concerned. Not all involve purely visual cues, but those which do are fairly easily picked up, “once you know what you are looking for”.

⁹ Somewhat after the fashion involved in what is probably another philosophical myth — that of Blind Sight.

information. But we can stipulate that that is not what Fred is doing. Simply, he looks at a chick and a highly reliable opinion is somehow gifted to him. How is this not perceptual justification? What is missing?

The question this gets at concerns how closely within a broadly externalist epistemology of justification can an analogue be fashioned of justification as conceived on the internalist Q-I model, whereby scope is provided for conscientious epistemic responsiveness to the phenomenology of one's perceptual inputs. Intuitively, such responsiveness is something you can 'do' better and worse, can exercise care with, can perform badly because of inattention to detail, etc. Can an externalist account of perceptual knowledge and justification, which dispenses with the idea of phenomenal states as the ultimate justifiers of perceptual claims, offer anything to provide scope for justification of this kind—scope for, as it were, doxastic-managerial justification?

There is a possible approach that promises something of the sort. It is broadly in keeping with a Knowledge First direction, although it will identify perceptual justification as something which an agent may possess in both knowledgeable and failed-knowledgeable cases. Alan Millar makes this proposal:

X has perceptual knowledge that P if X has formed a true belief that P as the product of an exercise of an ability to tell, of circumstances in which P is perceptibly so, that they are such.¹⁰

Bearing in mind our earlier reflections on the kinds of concepts apt to feature in purely recognitional judgements, I suggest we understand "circumstances in which P is perceptibly so" as circumstances such that anyone perceptually well-functioning who understands P and is appropriately perceptually attentive is thereby in position to make a knowledgeable judgement that P.

So we should ask now why the above formulation fails to apply, if it does, to Fred and his spectacularly reliable discrimination of the sexes of chicks. The answer, I take it is because, although Fred is perceiving the chicks, and his discriminative ability depends on that, and although he is not apparently processing what he sees in the light of some collateral information about how male chicks look, there is nothing in the description we gave of what he is doing to justify the claim that (for him) the gender of a passing chick is something that can be perceptually apparent: we have so far no basis for the claim that there is a way male, respectively, female chicks look and that Fred is tracking that. It may instead be, for example, that Fred has some hitherto undocumented gender-detective faculty, which switches off when he is blindfolded or in the dark.

I want to suggest, though I am not sure how to prove, that the most conservative addition we can make to the Fred example to bring it under Millar's rubric is to introduce a kind of psychological state which I propose to call a *registration-state*: a state

¹⁰ This formulation is not a direct quotation but expresses a key idea recurrent in Millar (2009, 2011, 2019). As Millar is well aware, the idea needs some tinkering to cope with cases, like that of the legendary Barn Façade county, where a subject luckily forms a true belief that P—"That is a barn"—on the basis of exercise of a perceptual recognitional ability to discriminate P-circumstances, in a situation which, unknown to them, is replete with convincing fakes. One option would be to make the notion of "an ability to tell" environment-relative.

which pre-doxastically registers an appearance of P-circumstances. A registration state will be a content-bearing state. It will carry the content of an appearance that P and is what one enters into when one's relevant sensory faculties have done their job and one is primed to form the appropriately corresponding belief. Perceptual knowledge, on this account, is true belief formed by doxastic uptake of the content of a soundly generated registration-state. The range of possible registration states is limited by the range of ways things can possibly appear.

I want to suggest that without any play with such a notion, it is not possible to make the distinction we called for—to distinguish between true beliefs reliably generated purely perceptually, and true beliefs reliably generated by means essentially involving perception but whose workings are otherwise unknown (as with Fred.) Moreover it is only by invoking some notion of a pre-doxastic content-bearing state, induced by the operation of one's perceptual faculties, that it is possible to make room for the aptness and proportionality of the beliefs one forms in response.

Bear in mind, however, that we are seeking an *externalist* conception of perceptual justification. So how should registration states be conceived if their role is to be compatible with externalism? Well, not as the internalist conceives of perceptual evidence. For the internalist, perceptual evidence is constituted by a state purely of the subject. It is a species of a genus that also includes dreams, hallucinations and maybe some vivid imaginings. For an externalist, by contrast, a registration state has to be a *world-involving* dyadic state, involving real relations between the subject and external circumstances. It follows that the externalist who pursues a notion of justified perceptual belief along these lines must go in for *Metaphysical Disjunctivism*—any, as it were, environmentally-disconnected phenomenal state of the subject, however convincing it may appear as a putative registration state, has to be reckoned as of a different ontological category.

As is familiar, however, nothing about epistemology follows directly from metaphysical disjunctivism. Let GOOD be a registration state carrying the content that P and BAD an hallucinatory counterpart, phenomenologically indistinguishable from GOOD. Then the metaphysical disjunctivist is so far perfectly free to allow that GOOD and BAD *equally rationalise* the belief that P. It is an extra step to insist that only a genuine registration state—a real world relational state—can justify an appropriately corresponding belief. For *Epistemological Disjunctivism*, only GOOD is knowledge-conducive, and notions of rationalisation, justification etc. should be so fashioned that only potentially knowledge-conducive states can rationalise or justify.

One immediate consequence of the position at which we have arrived is that we can now fashion, from an externalist point of view, one perfectly good notion of *fully justified but false* perceptual belief: belief that is justified in just the same sense, and to just the same degree, and for just the same reason as a knowledgeable counterpart. For a registration state can *misrepresent* the situation it registers. And it can do so even when there is no relevant shortcoming of the subject's perceptual apparatus. This is because even concepts apt to feature in purely recognitional judgements as characterised earlier, can be associated with misleading public appearances. Colour concepts are presumably in this group by the suggested criteria. *Red*, for example, is plausibly such that to know what it is for something to be red requires knowledge of

how red things distinctively look—there is no such thing as knowing what it is for something to be red yet having no inkling of the distinctive range of appearances being red covers. And conversely, there is no such thing as having a perceptual recognitional capacity to recognize red things as such and to conceptualise that recognition in judgement while failing to grasp *red*. However that the possession conditions of the concept have these features does not exclude the possibility of things appearing red—distant rapidly receding stars, white walls bathed in red lighting, etc., etc.,—which are not. Registration states may accurately register *misleading public appearances*. Someone who undergoes such a state should surely be regarded as epistemically fully justified, absent other information, in judging falsely that the item concerned is red.

There is motive to make a similar concession for at least some cases where misrepresentation is owing to shortcomings in the perceptual functioning of the subject. If one's perception of a certain scene misrepresents certain details, but is nevertheless sufficiently successful to count as portraying the scene, it will do no violence, it seems to me, to the spirit of R-S externalism to allow false beliefs that draw on misrepresented aspects to count, *ceteris paribus*, as no less justified than true beliefs that draw on veridically represented ones.

37.5 Justifications and Excuses

The foregoing are cases where the subject is on the GOOD side of the epistemic disjunctivist's disjunction but where either external things give off appearances—real public, inter-subjectively appreciable appearances—of being other than they are, or where a limited degree of inaccuracy afflicts the subject's pre-doxastic uptake of those appearances. Where an R-S externalist and the Q-I model internalist must disagree, it appears, concerns the justificational situation of a subject on the BAD side—say, the victim in the New Evil Demon Scenario (NEDS—Cohen (1984)). For the internalist, that situation is exactly as it would be if the subject were perceiving normally. Not so for the R-S externalist. For that theorist, the NEDS victim simply has *no evidence*, no basis for a view about anything in their local environment. They simply undergo no (apparently perceptual) potentially justifying states.

This consequence is apt to seem a crux. It is going to be a consequence of any form of externalism about perceptual justification. Many find it a highly rebarbative consequence, indeed as defeating the chances that an externalist account can do justice to the notion of epistemic justification that we actually have. Surely, it may be felt, if your experience has been as it has and shows every sign of sustaining the continuities and coherence it has exhibited hitherto, you will be justified in viewing the world in general exactly as you do, whatever the actual source of your experience. To be sure, in this scenario you will lack genuinely *perceptual* justification, since you will be having no genuine perceptions. But you will be no less justified, and the phenomenal states of mind that will do the justifying will do so purely on the basis

of the phenomenal character and apparent content that they would share with the possible perceptual states they ape.

In the literature, many have professed an ‘intuition’ along these lines—an intuition that a subject in NEDS can indeed be fully justified in forming exactly the beliefs they would form: in particular, that there is no respect in which they are less than fully epistemically responsible in forming those beliefs. But what has become the standard externalist reply to this “intuition” is that it confuses being justified with having an *excuse*. (Williamson (forthcoming) seems to have been the originator of this response.) The thinker in NEDS cannot be *reproached* for forming what they take to be perceptual beliefs in the way they do. They are excusable for doing so. But they are doing the wrong thing nevertheless.

Calling either of these conflicting accounts an “intuition”, with the overtones that word carries of some kind of positive philosophical evidential status, is tendentious. What’s at stake in the justification/excuse contrast here? Why should we fashion these notions one way or the other? Is it possible to adjudicate the dispute?

Here’s one natural consideration. When an agent, engaged in some project, performs sub-optimally, it’s good to respect a distinction between factors contributing to that sub-optimal performance that lie beyond their control and others which they could have corrected for. It is if the explanation of their sub-optimal performance rests with factors of the second category that we enter the proper territory of possible excuses—“She could have noticed that she left the front door unlocked, but was distracted by her need to hurry to the hospital”. A subject in NEDS is certainly performing sub-optimally in the supposedly perceptual beliefs they form, but the explanation of that can reside entirely in factors beyond their control or anything they are in position to know. So we should not use the notion of an *excuse* to mitigate a negative assessment of their performance.

Maybe that consideration draws only on aspects of our ordinary understanding of “excuse”. The locus classicus for consideration of such questions is of course Austin (1957). In any case, though, even if it would not square with our usual understanding of the term to describe the NEDS victim as merely excusable, that point does not commit us to regarding them as epistemically *justified* instead, let alone as justified to the same degree as they would be were their experience phenomenologically indistinguishable from what it actually is but genuinely perceptual. Maybe the hapless NEDS subject is neither justified *nor* has any excuse but, as it were, slips between these notions as usually understood!

The situation seems thoroughly unsatisfactory. Is there some insight we are missing that will drive a verdict in one direction or the other? I don’t think so. Rather, if we pursue the discussion a little further, I think we can begin to see the outline of a reason why the dispute is actually unnegotiable. To be sure, we have at least to start with aspects of our pre-theoretical understanding of the notions involved. So I suggest that it should be common ground that *justification* for X-ing in circumstances C not merely exonerates one who X’s in C from blame if things turn out badly, but provides a *mandate*—an ‘ought’ or a permission—for X-ing in C. By contrast a (mere) *excuse* for X-ing in C merely exonerates. If we say A was justified in X-ing in C, we imply that X-ing was something A ought to have done in C. or at least that

it was perfectly permissible for A to X in C. If we say, by contrast, that A, though unjustified, had a good excuse for X-ing in C, we imply that while X-ing was nevertheless a wrong or impermissible thing to do in C, the pressures on A were such that they—and perhaps most people—could be expected to X in the circumstances and that it would be harsh to subject them to the negative reactive attitudes that doing the wrong thing or transgressing what is permissible would otherwise normally incur.

If that seems roughly right, then it reminds us that these distinctions presuppose background values. *Oughts* are imposed by pressures to further certain values; *permissions* are generated by there being no compromise of relevant values involved in the permitted act. And *excuses* are considerations that mitigate failures to do what one ought, or mitigate actions that exceed what is permitted. So the question for the internalist who wishes to regard the NED scenario victim as justified in their beliefs is: what values would be frustrated if the victim *failed* to respond to their apparent perceptual states as if they were genuine? It is only if there are such values that there is any question of such a pattern of response being justified. The externalist, by contrast, owes an explanation of the values that are frustrated by such a pattern of response: there must be such values if the invocation of the idea of excuse is to be appropriate—you don't need an excuse for doing something which frustrates no value.

How should the protagonists answer these respective questions? The internalist can answer: the value that would be frustrated by a NEDS subject's failure to respond to their apparently perceptual experience in the way proposed would be a value of *rationality*: the maximisation of coherence of the apparent contents of ones experience with ones beliefs about the world. The externalist likewise can answer that the value that would be frustrated by a NEDS subject's indulgence in the patterns of belief formation that would be appropriate in GOOD is that so responding takes them systematically away from the truth about their world; the value concerned is the maximisation of true belief concerning their environment.

And both of course are right. Rational coherence, it hardly needs saying, and truth are indeed epistemic values that we actually have. We certainly value the avoidance of error and getting things right. And we certainly value overall coherence and integration within our system of belief and presumed evidence. The tragedy of the situation of the NEDS victim is that there is nothing they can do to further either of these values without systematically frustrating the attainment of the other.

The question, how should we apply the contrast between excuse and justification in such a situation, is thus an instance of the question, how should we handle that contrast in a situation of conflicting values. Consider an ethical analogue: a situation, say, where my carrying out a promise conflicts with the interest of the promisee. Suppose you think it clear that the latter should trump the former. Then you are liable to regard my attempt to justify what I do by highlighting my promise as the provision of a mere (though perhaps a sympathetic) excuse. On the other hand, if you regard the keeping of promises as a sacred duty, even when doing so will work out to the disadvantage of the promisee, then you will regard my failure to do so as quite unjustified, though my concern for the welfare of the promisee perhaps does me credit. If we can generalise from this example, it appears that, whether one

describes the NEDS victim as justified or merely excusable depends on the relative weighting assigned to values of epistemic product—truth and knowledge—on the one hand and values of doxastic management—rational integration and explanatory coherence—on the other. But of course it is the relative importance of these values that is what is fundamentally at issue in the whole internalist-externalist confrontation. The performance of the NEDS victim will be merely excusable if the first takes priority. For then, by the more exigent value, the subject is performing as they ought not and the overall coherence of the results is merely a mitigating factor. But we will regard the NEDS victim as justified if we take the view that epistemic ‘ought’ implies epistemic ‘can’ and that the very intelligibility of the NED scenario implies that, however important truth and knowledge may be, we can at an ultimate metaphysical level take no responsibility for securing them and that, as good Pyrrhonians, we should therefore aim to control only what we can control—the overall coherence of the way in which we manage our system of belief.

If this is right, then the thought-experimental predicament of the NEDS victim can give us no independent leverage whereby to adjudicate between internalist and externalist conceptions of perceptual justification. Rather what we should think about how the contrast between justification and excuse applies in NEDS depends on antecedent views about the relative ranking of the values prioritised by those conceptions.

37.6 Stocktaking: Internalism, Externalism, and Perceptual Scepticism

Each of the Two Conceptions reviewed recognizes a role in perceptual justification for a predoxastic, content-bearing state: for the internalist, a ‘narrow’ state of being appeared to that P; for the externalist a registration state, that constitutes a pre-doxastic engagement with external circumstance, a taking in of at least the publically apparent “layout of reality.”

What if any differential impacts do internalism, as here understood, and externalism have on the challenge posed by Perceptual Scepticism? I think the answer in the end is “None”, and that the—at least historically—widespread idea that externalism gives us some special leverage with scepticism, is an illusion. Moreover the reason why this is so has nothing to do with the concessions to the idea of false but justified belief that motivate my canvassing the possibility of R-S externalism.

A qualification: there are sceptical paradoxes that differ in significant detail, some of which (e.g. Descartes’ style of argument, when naturally formulated as an explicit deduction) arise only on internalist assumptions. One such is any style of sceptical argument that starts from the premise that Descartes’ evidential situation is the same whether he is indeed sitting in his armchair in front of a blazing fire or merely dreaming that he is. Certainly any externalist conception of perceptual justification, including R-S externalism, will reject that premise. Once the premise is rejected and the question, what evidence we have for claims about our material environment, is

made to turn on the actual, contingent character of the aetiology of our apparently perceptual experience, it ceases to be anything that armchair philosophical argument can be competent to engage.

It plays differently, though, if what is at issue is a *second-order* question, arising at the level of reflection: the question of our *entitlement to claim* perceptual knowledge and justification. Once it is granted that for a subject in BAD, things can consistently and coherently seem as though all is well, both internalist and externalist must recognize that the claim that we are routinely receiving perceptual knowledge and justification every day, notwithstanding their differing conceptions of what that amounts to, will rest on ‘Not-BAD’ as at least a prop. And that consideration will be open to prop → lemma inflation, thus enabling a sceptical challenge whose rhetoric is exactly the inflation:

So you are taking it that what you regard as genuine sensory engagement with the external world is indeed just that, and not some ghastly delusion?

As before, the honest answer is “Yes, of course”. At which point the options are

- (i) to provide some kind of rational backing for that answer—perhaps something like: “It is of the essence of enquiry to make “assumptions” of this general character”—or
- (ii) to offer reasons why no backing is required, or
- (iii) to stop doing philosophy.¹¹

References

- Austin, J.L. 1957. A plea for excuses. *Proceedings of the Aristotelian Society* 57: 1–30.
- Cohen, S. 1984. Justification and truth. *Philosophical Studies* 46 (3): 279–295.
- Dretske, F. 2005. The case against closure. In *Contemporary debates in epistemology*, ed. M. Steup and E. Sosa, 13–25. London: Blackwell.
- Hawthorne, J. 2005. The case for closure. In *Contemporary debates in epistemology*, ed. M. Steup and E. Sosa, 26–43. London: Blackwell.
- Millar, A. 2009. What is it that cognitive abilities are abilities to do? *Acta Analytica* 24 (4): 223–236.
- Millar, A. 2011. How visual perception yields reasons for belief. *Philosophical Issues* 21 (1): 332–351.
- Millar, A. 2019. *Knowing by perceiving*. Oxford University Press.
- Pryor, J. 2012. When warrant transmits. In *Mind, meaning and knowledge: Themes from the philosophy of Crispin Wright*, ed. A. Coliva, 269–330. Oxford: Oxford University Press.
- Williamson, T. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. 2007. The perils of dogmatism. In *Themes from G. E. Moore: New essays in epistemology and ethics*, ed. S. Nuccetelli, and G. Seay, 25–48. Oxford: Oxford University Press.

¹¹ An early draft of this material was presented at the conference on The Epistemology of Perception held at the University of Pavia in June 2019. I benefited greatly from comments on that occasion and on the occasion of a presentation at Stirling’s Work In Progress seminar in November 2021 which, as well of course as the discussion at the PEER conference itself, have led to significant improvements in the version published here.

Wright, C. 2012. Replies. In *Mind, meaning and knowledge: Themes from the philosophy of Crispin Wright*, ed. A. Coliva, 451–486. Oxford: Oxford University Press.