Social Epistemology, 2015
http://dx.doi.org/10.1080/02691728.2014.990281

Routledge
Taylor & Francis Group

Perspectivism, Deontologism and Epistemic Poverty

Robert Lockie

The epistemic poverty objection is commonly levelled by externalists against deontological conceptions of epistemic justification. This is that an "oughts" based account of epistemic justification together with "ought" implies "can" must lead us to hold to be justified, epistemic agents who are objectively not truth-conducive cognizers. The epistemic poverty objection has led to a common response from deontologists, namely to embrace accounts of bounded (perspectival) rationality—subjective, practical or regulative accounts rather than objective, absolute or theoretical accounts. But the bounds deontological epistemologists and their opponents entertain rarely include cultural limitations. This paper considers neo-Vygotskian (contextualised cognition) arguments that we must consider such cultural limits in defending deontologism, and thus that any deontologically motivated perspectivism must be in part a cultural perspectivism. The dangers of strong relativism are flagged and an attempt is made to steer a middle path between this and a parochial anti-rationalist objectivism.

Keywords: emic; etic; relativism; bounded rationality; epistemic deontologism; epistemic poverty; epistemic perspectivism; regulative-theoretical; competence-performance; decontextualised vs contextualised cognition

Robert Lockie is senior lecturer in the Psychology Department at the University of West London. Major research interests are Epistemology, Free Will, Metaphilosophy, Philosophical Psychology and Philosophy of Mind; though he also has research interests in theoretical psychology, especially in cognitive psychology, neuropsychology and individual differences. He is the author of a number of papers in these areas, including (recently) "Three Recent Frankfurt Cases", forthcoming in *Philosophia*, online first May 2014; "The Regulative and the Theoretical in Epistemology", forthcoming in *Abstracta*; "The Epistemology of Neo-Gettier Epistemology", *South African Journal of Philosophy*, 33(2) April 2014, 247–258. Correspondence to: Robert Lockie, Psychology Department, University of West London, Paragon House, Brentford TW8 9GA, UK. Email: bob.lockie@uwl.ac.uk

© 2015 Taylor & Francis

The Epistemic Poverty Objection and its Standard Response

The deontic conception of internalism in epistemology (henceforth and throughout: just "internalism") involves the idea of cognitive accessibility, of epistemic deontology, and of these being connected via an "ought" implies "can" principle. These three things may be used to create a problem for internalism. Here is Alston, seeking to do just that. Note that his target, "J_{di}" stands for *deontic, internalist, justification*: just the deontic notion of epistemic justification as developed after Clifford, Descartes, Locke, Chisholm, Foley and many others—not least Alston himself.[1] Its chief characteristic is that the person possessed of such a conception of justification has done everything they reasonably could do in the circumstances to be justified, to fulfil their epistemic obligations. Still, he notes:

> J_{di} does not give us what we expect of epistemic justification. The most serious defect is that it does not hook up in the right way with an adequate, truth-conducive ground. I may have done what could reasonably be expected of me in the management and cultivation of my doxastic life, and still hold a belief on outrageously inadequate grounds. (Alston 1985, 67–8)

Alston then gives examples of such cases—of dutiful but helplessly flawed cognizers. These are examples of a familiar type, typical of many in the externalist literature: A tribesman may have been brought up to accept the traditions of his tribe as authoritative, and have never seen anything to call these traditions (of inquiry etc.) into question. (We will return to this "tribesman" example below). Or a person may be intellectually honest and diligent, but just rather dim; or not dim especially, but highly impressionable. A person may lack an education, being vulnerable to all kinds of unfounded hearsay and superstition as a result. In general, one may have discharged one's epistemic duties as responsibly as one is able, but still (blamelessly) be holding one's beliefs on "outrageously inadequate grounds". In a nice turn of phrase, Bonjour (2003, 176) refers to these as situations of "epistemic poverty".[2]

The point here is that if epistemic justification is to be construed deontically, then, *a priori*, "ought" implies "can"—the deontic terms of approbation and disapprobation alike require that I could have done otherwise (and didn't or did respectively). But then, if I could not have reasoned any better than I did; by *modus tollens*, the internalist may not hold me to be unjustified. Yet epistemic poverty cases of this kind are meant to indicate that one may be outrageously unjustified for all that, since one may blamelessly fall foul of Alston's "most serious defect" and "not hook up in the right way with an adequate, truth-conducive ground".

The Standard Response: Distinguish Objective From Subjective Justification

This stock argument—in both ethics and epistemology—has a stock response. This is to draw a distinction between *objective* and *subjective* duty. One is culpable,

blameworthy, irresponsible, should one fail to discharge one's subjective duty; one is not blameworthy, irresponsible, merely in virtue of failing to discharge one's objective duty—which failure may be quite out of one's hands. Owens (2000) notes this distinction in Sidgwick (*cf.* 1907, 413). Plantinga (1993) credits his version of this distinction to Aquinas; though a more immediate source might well be his formative influence, Chisholm, who himself credits Richard Price:

> When it is said, then, that people ought to believe what is evident and that they ought not to believe what is unreasonable, "ought" has its practical use. But when it is said that they ought to believe what is true and ought not to believe what is false, "ought" has its absolute use. (Chisholm 1957, 6–8)

As Alston notes:

> S has fulfilled his *objective* obligation iff S has brought about A. S has fulfilled his *subjective* obligation iff S has done what he believed to be most likely to bring about A. (Alston 1985, 61–2)[3]

The deontic understanding of justification, plus the ought-implies-can principle, leads to the threat of being never able to convict an agent of being unjustified, even when he is radically awry, should he or she merely be unable to apprehend the fact that he or she is mistaken. The solution to this problem for the internalist is to mark a distinction between "absolute" and "practical" justification—between fulfilment of "objective" and "subjective" duties. To escape the epistemic poverty objection, deontic, oughts-based justification must be restricted in its application to the "subjective", "practical" realm. There is another, "objective", "absolute", sense of being justified for which the discharge of duty, the fulfilment of obligations, be we ever so diligent, is not guaranteed to satisfy.[4] So, consider when Goldman says

> I shall assume that only *right* epistemic rules make a difference to genuine justifiedness. This point should be equally acceptable to both internalism and externalism. (Goldman 2009, 5–6)

This point will be "equally acceptable to both internalism and externalism" only should this statement be read by each under a different interpretation of "right epistemic rule". For the internalist, this means *subjectively* right; for the externalist, *objectively* right.[5] Argument at cross purposes beckons if we do not keep this in mind.[6]

Bounded and Grounded

The development of this distinction in the epistemic and ethical literature was paralleled by neighbouring ideas in the psychological and economic literature. Following Herbert Simon's work (itself contemporaneous with Chisholm's) it was widely acknowledged that rationality had to be *bounded*. That is, the rational standards that any agent could be held to needed attenuation in the light of processing

constraints and other human limitations (computational restrictions, working memory constraints, informational restrictions, time to reach decision, etc.). An idealised ("reified") notion of rationality would be no theory of (human) rationality at all, it would face problems of computational intractability and "combinatorial explosion"—it would perhaps pertain to Martian rationality, or the rationality of Broad's Mathematical Archangel, but not to that of a human subject in an epistemic quandary. Workers in these literatures differ considerably in how bounded their theories are—with some being more idealised and others more perspectivally attenuated. A basic distinction may be drawn (after Elqayam 2011, 2012) between *bounded* and *grounded* perspectival restrictions on epistemic justification. Roughly, the former pertain to species-wide psychological constraints—on working memory, processing capacities, computational limitations and time constraints—in educated humans who may nevertheless be otherwise considered idealisations or abstractions. The latter factor in individual differences limitations (how dim or bright the rational agent is, their cultural and educational limitations—Alston's "tribesman" point, and the like). The spectre of relativism hovers more obviously over a theory that acknowledges the latter constraints than the former, but neither is an "objective" or "absolute" account. It is the same "ought" implies "can" principle that is used to argue for either kind of constraint, bounded or grounded, thus acknowledging one kind of perspectival attenuation yet not another may be deemed rather arbitrary.

Deontically internalist theories' restrictions on only accepting a justificatory ground that may be accessible to the agent is then best seen not as a definitional primitive. The access restriction is rather derived from the deontic core of the theory. To be necessarily capable of directing cognitive conduct—of providing "rules for the direction of the mind"—an epistemology must not go beyond the resources of the epistemic agent: those resources that are bounded by the agent's own intellectual abilities and such cultural facilities as they have access to some distance out from these. That is, it may travel perhaps into their "Zone of Proximal Development" (Vygotsky 1978); but it must remain within the compass of their bounded and grounded resources in at least some, perhaps extended, sense (*cf.* the Theaetetan aviary example (Plato 1992, 197c–d)).

The dialectic to this point should be familiar to most readers. The challenge is whether the position we have reached is a stable and defensible one. Too few deontic-internalist epistemologists who are tempted to draw the standard perspectivalist restriction appreciate how severe will then be the pressure on them towards quite an extreme relativism. Too few externalist epistemologists who are tempted to repudiate the standard perspectivalist restriction on grounds of a lack of truth-conducivity appreciate how severe then will be the pressure on them towards quite an extreme (and arbitrary) parochialism. The attempt in what follows will be to chart a middle path between these extremes; but before doing this we should introduce a small part of the empirical evidence that informs these debates.

Perspectivism and Real Rather than Armchair Tribesmen

Alston's tribesman case is put forward by him, quite legitimately, only as a philosopher's "toy" example; but as examples of epistemic poverty, toy or not, such cases should be seen as raising large issues. That such armchair constructions are cases where agents may straightforwardly be taken to lack knowledge should be conceded far more readily than that they, with the same facility, should be taken to be irrational. After all, a necessary condition on knowledge is truth, truth is absolute, and if said tribesman is in possession of a falsehood then he is in this respect, straightforwardly epistemically impoverished; and further, regarding the *epistemology* of these epistemic attributions, may in principle be seen by us (if we are in possession of said truth) to be relatively more impoverished than us (though apropos the neo-Bartlettian tradition, here there are some caveats, even as regards knowledge attributions). But attributions of irrationality, whether between cultures, or across individuals within a culture who start from the possession of differing epistemic resources, will be a lot more problematic. Philosophers commonly underestimate how culturally situated, how educationally laden, and how restrictive is their conception of the rationally acceptable variance here. As an example of between-cultures variance, consider Luria's famous "white bear" case (Luria 1976 in Mackintosh 2012). In sensitive, culture-fair testing that was unusually sophisticated for its time, Luria/Vygotsky presented illiterate rural Uzbek peasants in the 1930's with the following:

> In the far north, where it snows, the bears are white.
>
> Nova Zemblya is in the far north, and it is always snowy there.
>
> What colour are the bears in Nova Zemblya?

Typical responses were

> How should I know? I have never been to the north.
>
> Why are you asking me? You have travelled and I have not.
>
> So-and-so said the bears were white, but he is always lying.[7]

The peasants refused/failed to "decontextualise". They didn't get as far as Cheng and Holyoake's (1985) claim that what was lacking in cases of failure in reasoning using conditional syllogisms (within our own culture) was a context that made clear to participants that they should apply "pragmatic reasoning schemata" (rules of *permission*, *obligation* or *authorisation*). The prior contextual failure here was much more radical. The Uzbeks may not even have got as far as what is known (in the literature seeking to defend the rationality of western psychology undergraduates in the face of their difficulties with the conditional syllogism) as an *alternative task construal*—one of L. J. Cohen's (1981) "performance errors".[8] Luria's Uzbek peasants appeared to refuse/fail to accept such problems as legitimate at all. Are we to maintain that conditions of (acknowledged) epistemic poverty withheld

from Luria's Uzbeks the capacity for rationality in regard to extraordinarily basic logical inferences?

Well, presumably we should accept that some kind of epistemic error/failure/limit is indeed revealed here. (Though note that in the political context of the time, to be as terse as the above on being questioned by educated Russian outsiders might have been no *error* at all[9]). Question: do we hold this to be a *competence* failure—logical incompetence concerning an inference principle as basic as *modus ponens?* Or do we rather say that some kind of "front-end" performance difficulty is present, yet with no clear evidence of logical incompetence—a logical competence that may still be defended as universal? If we accept this forced-choice, surely we must say that we don't know (though with the "rationalist" suspicion that it may be a more-immoveable-than-usual, culturally determined performance error); but I would suggest that we should also consider whether such cases stretch the limits of the performance-competence distinction itself. Problem cases such as these (and they are common in the cultural psychology and anthropology literature) are unlikely to be dissolved by a project to make what is already an apparently "culture fair" front-end even more "culture fair"—or indeed, (*per impossibile*) "culture free". Rather, the aspiration to what must always and inevitably be merely a *peripheral* cultural fairness is here found to be a vain one. Further, by pursuing the project of fractioning out cultural, perspectival, differences—as contextualised performance noise against a backdrop of "pure" decontextualised logical competence—one may throw out precisely what it is one wishes to investigate (*cf.* Bartlett 1932).[10]

But should we then judge, à la Alston, that epistemic poverty has made the Uzbek peasants straightforwardly irrational? Some psychologists (especially the defenders of two-system theories)[11] who, when faced by the rationalist ("Panglossian") apologia that participants much closer to home are not illogical/irrational but are merely construing the reasoning problem differently, counter that there are rational norms too, concerning the appropriate heuristics and algorithms employed by participants for *task construal* (e.g. Stanovich, West, and Toplak 2011). So, such psychologists offer as response, the options either that participants are irrational for failing to solve the problem that they have construed correctly, or that they are saved from this charge of irrationality only to be convicted of prior irrationality at the level of *task construal*—say, for failing to decontextualise appropriately (*ibid.*, Stanovich and West 2000, Adler 2000). Perhaps that is an acceptable line to take with one such as Cohen (1981), who tries to defend, with narrow reflective equilibrium, the rationality of those who fail e.g. Kahneman–Tversky or Wason tasks yet have been granted access to the full range of our decontextualised cultural and educational training to adulthood (such psychologists' undergraduate participants).[12] Do we though, apply such an approach cross-culturally—to participants such as Luria's? One hopes not. There has been a *huge* amount of prior training for us to have reached the stage at which we can interpret, say, an IQ test as intended, and there needs to be more for us to interpret a Bayes theorem task or Wason task or rational decision theory task as intended. Decontextualised contexts

of this nature are, in a sense, as contextualised as any the Uzbek peasant or the Kpelle tribesman trades in—they are *contexts of decontext*—hugely powerful and extremely acculturated ways of gearing up our cognition. Because we have been so exhaustively inculcated, literally from birth, into such decontextualised contexts, we fail to appreciate how extraordinarily immersive and intensive our acquisition of this context (of decontext) has been (*cf.* Flynn 2011[13])—or how recent its arrival in human cultural and intellectual history. Though it wouldn't be as easy as for Socrates' exchange with the slave boy, we could, in principle if not in fact, give Luria's peasants a wholesale modern schooling, training them to a performance criterion whereby they decontextualise somewhat as ourselves (or, as Luria/Vygotsky emphasised, as their Soviet-educated children)—but what would that prove? We want to know whether *these peasants here* are rational beings, not whether a genetically identical but culturally-cognitively altered counterpart would be. To know the answer to that, one would need to see the problems which confront them in their actual lives, understand something of the resources they bring to bear upon these problems, and decide matters accordingly. That would require a *lot of knowledge* about them, a *deep engagement* with them—an understanding of their epistemic perspective at least partially from within (an "emic" rather more than an "etic" approach[14]). As the tradition from Vygotsky and Luria puts it: it would require insight into their *cultural construction of the world*. Absent such knowledge, I assume, at least in principle, that they may be rational in the ways that it is open to them to be rational, no less than you or I may be rational in the ways that it is open to us to be rational.

The Danger of Parochialism

An issue confronting any epistemologist is how much of what he or she has to say about normative epistemology, and especially about rationality, will generalise cross-culturally. In common with most epistemologists, I do not have a specialised training in anthropology or cultural psychology; and a measure of disciplinary humility is appropriate. But the approach to epistemic justification I endorse—namely, a thin deontology, bounded by the acceptance of an "ought" implies "can" principle together with a recognition of human cognitive and (here) cultural limitations, an approach which thereby emphasises the regulative rather than the theoretical desideratum of epistemic adequacy—is clearly vulnerable to counterattack until or unless my more theoretical opponents are made aware of how vexed these issues actually are, and how fraught is the "objective", "absolute" position on rationality that they are committed to defending. *All* accounts of rationality are and must be must be bounded accounts—to at least some degree. The stance Oaksford and Chater adopt (here, of Bayesian reasoning) applies not just to Bayesianism, but to other areas in which a conception of "rationality as objective optimisation" is operative:

> Spontaneous, real-world risky decisions, even of moderate complexity, are not made using Bayesian inference, because they could not be. Since the mind/brain is a limited information processor, the processes of risky decision-making cannot be based upon optimal, algorithmic procedures. This means that the only rationality to which we can aspire, as individual decision-makers, is one bounded by our limited computational resources. In consequence, the observation that we do not behave in accordance with Bayes's theorem could not impugn our rationality. Our rationality could be questioned only if we were capable of using the optimal strategy but failed to do so. Thinking otherwise is akin to condemning us because we do not fly even though we do not possess wings. (Oaksford and Chater, 1992, 226–7)

The issue is not then, one of bounded versus unbounded rationality, but of how bounded, and which bounds. I believe there can be no principled grounds for including species-wide biological constraints (working memory, processing speed, etc.) yet not cultural constraints (which will also be mediated and expressed biologically). But this position—endorsing an emic somewhat more than an etic approach to rationality—must then be defended in turn against charges of (a) all-out relativism (whether cultural or subjectivist) and (b) that by offering a ready-made exculpation for so many putative errors of rationality it undermines the project of normative rationality *tout court*. It also has to be stated (c) where it stands on the notion of a project for "objective" rationality at all.

As regards (c) there are a range of legitimate projects for investigating more objective concepts of rationality, projects distinct from both the objective investigation of knowledge and the kind of highly perspectivist approach to rationality that I advocate. Although this paper is not concerned with such projects, they are legitimate areas of inquiry—especially the more psychological projects, and especially in their more recent (less irrationalist) versions. The psychological literature on judgement, reasoning and choice is clearly of importance, both positively and negatively—say, in taxonomising the syndromes through which objectively sub-optimal reasoning often manifests itself across a range of highly decontextualised tasks in our culture (what formerly was called the heuristics and biases literature).

There are likewise legitimate non-psychological projects to derive theorems in logic, statistics, rational decision theory, etc., and apply them intelligently to real-world problems. Someone who is in a context whereby he has access to such algorithms yet doesn't apply them, will objectively be likely to do poorly, and moreover will be at least likely (subjectively, perspectivally) to be irrational for not applying them. There is no suggestion that permitting a Dutch book to be made against oneself, or applying a non-Bayesian algorithm to a clearly Bayesian task, or playing non-expected-value poker against a poker-bot is anything but objectively sub-optimal in terms of desired outcome. We are, however, all of us, unavoidably, deploying objectively sub-optimal algorithms in much of what we do—because we are all limited in our epistemic resources. Even a very idealised, objectified, rational agent is still, *qua* rational agent, very strictly bounded—and grounded. A "reified" rational agent, idealised beyond all bounds, is no rational agent at all (Kant 1933, B8, the "light dove").

Seeing things in terms of individual differences expands on this point. Clearly those with higher general intelligence will often (Stanovich and West 2000) though not always (Stanovich, West, and Toplak 2011) do better across a range of problem situations than those who are less intelligent; as often (depending on the task) will those who are better educated; or have greater "need for cognition" (Kahneman 2003)—at least for the tasks with which this literature is concerned. But rationality is very substantially a *metacognitive* concept. It pertains not simply, or chiefly, to a characterisation of the power and efficiency of the first-order, dedicated cognitive resources that we bring to the task (Stanovich and West's "algorithmic-level efficiency") but rather how we *use* what we have got—the interplay between the algorithmic-level and the level of their "reflective mind": to say nothing of the interplay between the latter and that mind's extended grounding in culture. In this sense, a less intelligent or less educated or more culturally impoverished agent, may yet be more rational than a more intelligent or more educated or more culturally privileged agent—should the former use what she has got to work with better than the latter does relative to that with which she works. We use what we have got, they use what they have got, and they, and we, are rational (in this, subjective, regulative, sense) if we use what we have got well, diligently, and to the limits of our ability. There will be limits to how objectively irrational a subject can be before he or she is too damaged to be considered a part of the community of rational agents at all; but these limits can't, surely, be so restrictive as to require the community of rational agents to include all and only those with a modern university education, or written language, or grasp of statistics, or basic arithmetic, or any level of formal, abstract, education at all. *Most* humans who have ever lived have been rational animals. There is a decent, principled anti-relativism; and then there is a narrow, self-regarding parochialism.

The Spectre of Relativism

This leads to issue b) the question as to whether such a perspectival conception of epistemic justification must lead to the eclipse of the normative project *tout court*. I think that clearly it does not. One basic point has been stated above: someone who has access to an objectively good algorithm, and that it is an objectively good algorithm, yet who doesn't apply it, will objectively be likely to do poorly, and is at least highly likely (subjectively, perspectivally) to be irrational for not applying it. People are, after all, quite often irrational. The fact that there may be perspectival, exculpatory, factors available to them to absolve them from the charge of irrationality doesn't mean there routinely are such factors—hence, that there is such exculpation—it doesn't mean epistemic agents have a "get out of jail free" card: it just means that in any given case we must look at the details. Concerning a) I have argued against relativism *as applied to truth* elsewhere and at length (Lockie 2003a, 2003b, 2006). As regards the commitments of this kind of deontologism, a series of points can be made in summary here.

Firstly, as noted, there is no suggestion that, in the unlovely parlance of our times, the result of the contextualised reasoning/refusal to reason of the peasant/tribesman may be compared with the normatively correct decontextualised response and be deemed "equally valid". Objectively, where such inferences are wrong they are invalid in both the strict logical sense and any jejune metaphysical sense that lies beyond that. No less objectively incorrect would be a failure to draw an inference when enough information exists to do so,[15] as (*salve* the political interpretation) Luria's peasants did. In Bruner's famous phrase, it was possible to have *gone beyond the information given*—to work out that the bears in Nova Zemblya were white. Luria's peasants couldn't have done this, hence subjectively they weren't irrational; but "one could": objectively, that potentiality was there, in the data. Similarly, in, say, the equivalent decision-theoretic cases: where such reasoning is incorrect, nothing can make this anything other than objectively sub-optimal in terms of the desired outcome—and likewise for Bayesian cases and others in which this literature trades. Truth is absolute, without qualification, and where such cases involve one agent making an inference that is not truth-preserving/not truth-conducive and another agent making an inference that is, there is no suggestion that, objectively, the two sets of reasoning are on a par. But it is regulative epistemology we are doing here, not theoretical. Simply: absolutism is acknowledged at the level of the latter, and not at the level of the former. The regulative question is whether the agent has done all she can with the resources at her disposal, not whether, with resources she does not command, some quite other agent might have done better. (Note however, that *whether* she has done all she can, bounded and grounded by the perspective she is in, is itself an absolute, trans-perspectival matter of fact—not something of which she is the arbiter).

Secondly knowledge, the other great normative-epistemic concept, is externalist and not to be understood in this fashion (see Lockie 2004, 2008). The two desiderata carve up our enterprise between them: this, regulative, project of articulating a deeply perspectival account of rationality should not be considered in isolation from this other, theoretical, project.

Thirdly, Theaetetan *peritrope* limits on relativism are fully acknowledged: no system, even one as perspectivist as this, is without *any* absolute constraints and commitments. I am not here articulating any specific rational commitments (any "fixed bridgehead") but one obvious Theaetetan point is that any system which argues for a perspectivally bounded, deontological, account of justification on the basis of "ought" implies "can" cannot allow that this latter principle be intra-perspectival in turn; and this is a very substantial commitment. There will doubtless be other "absolutist" commitments, though I suspect these will largely resist armchair elucidation. Perhaps there is a "fixed bridgehead" of canonical rational principles (non-contradiction, say, though from the above, possibly not *modus ponens*); or perhaps there is only a neo-Lucretian variable bridgehead (though I suspect the latter need not be anything like as weak a position on rationality as some philosophers have supposed). I am here taking a stance on the broad-brush conceptual issues. I don't have to take a stance on the details here, and I suspect that when it

comes to such details, even where there are philosophical contributions to be made, they will require the philosopher to work quite close to the data. So, the reader is invited to read Cummins (1996) for an impressive argument to the claim that deontic reasoning is innate. Cummins marshals evidence from cognitive psychology, anthropology, primatology, developmental psychology and neuroscience to triangulate to the conclusion that we have an innate deontic reasoning "module". In only the first of these disciplines do I have something approaching an insider's grasp of the literature; but I find her arguments for innateness per se fairly plausible. I am likely committed to a claim as to the universal nature of deontic reasoning at some level—with innateness then being a candidate constructive explanation why. I am less tempted by the "modularity" aspect of her innateness hypothesis. My own arguments are not empirical however, and I remain a spectator in several areas of the scientific debates.

Fourthly, a reminder: the normative epistemology of these matters is what concerns us, not the cognitive (or other) psychology. Seeing rationality in deontic terms involves rationality as a *metacognitive* notion. We are asking whether the agent has reasoned as she ought. The reasoning we are appraising for rationality may perhaps employ Cummins' "deontic modules", just as it may employ Goldman's "J rules" or Cheng and Holyoak's "pragmatic reasoning schemata", or any of several explanatory posits for why, as the cognitive psychological literature indicates, deontic contents/contexts so enhance performance on, e.g. the Wason task, or other tasks. So, the first-order reasoning may be deontic in its nature (and there is good quality empirical evidence that some of it is); but it may not—and need not, to nevertheless be deontically *appraisable* at the higher-order (metacognitive) level. The first-order reasoning we are appraising for rationality may be any kind of reasoning that the agent *owns*, that belongs to her, that she is *responsible for*. This deontological approach does not then require that the agent's first order cognition be deontological, it requires that at the higher order level, important aspects of the agent's cognitive conduct be *deontically appraisable*: that the *agent* be deontologically appraisable. Even highly contextualised humans—tribesmen, peasants, etc.—*own* their cognition and are responsible for crucial aspects of it (including some of it that is non-deontic in any first-order sense). They own it and are responsible for it because (in certain, significant, respects) they have *control* over it. That is something I suggest that all humans have in common—at least, all that that may be seen as rational or irrational, namely: *Responsibility-relevant access and control over our cognition*.

It seems likely that this higher-order deontic appraisability, even though it can apply to non-deontic reasoning at the first-order level, nevertheless could not exist were there not the concrete cognitive capacities to implement deontic reasoning at some (e.g. metacognitive) level. But that wouldn't require first-order deontic faculties (Cummins' "innate deontic modules" say) to be reflexively applied at the meta-level—though it would be compatible with this as one possibility. An account of rationality answering to the regulative desideratum and phrased in terms of metacognitive deontic appraisability, might, at the executive, agential, level, not be best understood in terms of a faculty psychology at all. It (this execu-

tive metacognitive capacity, the presence of which is here envisaged as being required for us to be held epistemically responsible) might be emergent out of the holistic capabilities of the mind. Or, it might instead be a dedicated, ontologised, central, executive, processor—rather, that is, than a first-order module that is now (here) working on extrinsic tasks and now (here) turned in on itself. From the armchair, I do not have to take a stance here.

Fifthly, and lastly, a counterattack: we have to ask as to who the relativist is here. Is he the epistemic deontologist, acknowledging he is answering entirely to the regulative desideratum, and thereby embracing perspectivism; or is he the psychologist/philosopher/cognitive scientist who (*qua* theory of rationality) represents the objective, absolutist, opposition view? Picture a caricature (not entirely a caricature) "first-wave" ("irrationalist") version of the latter, castigating psychology undergraduates or blue-collar workers (we have not gone as far as peasants or tribesmen) for not decontextualising as ruthlessly as only a tiny (recent, culturally situated) cohort in the history of the world ever have been equipped to decontextualise. I suggest we should see rationality (and irrationality) as a species-specific human capability, potentially attainable by all—or nearly all. I hold that rationality is not situated in one arbitrarily specific way-station located between a humanly unbounded God's-eye reification and the perspective of our ancestors living in caves: a specific way-station to be found currently only in a micro-thin stratum of the statistically/logically educated élite—presumably then to be replaced by a new élite within one or a few scientific generations.

Highly contextualised humans (Alston's generic "tribesman", illiterate agrarian peasants, hunter-gatherers) faced and face enormously complex, open-ended challenges across a great range of startlingly different situations: challenges involving frequent life-and-death decisions in the harshest of environments. The idea that after surviving such challenges to adulthood they were and are not, as a population, rational; or at least not as rational as us in contemplating our supermarket and smart-phone decisions, is, on the face of it, not an attractive position to have to defend. Such agents may not be able to decontextualise a Bayes-theorem taxicab problem or a Wason task or a "culture-fair" version of Linda the feminist bank-teller as well as only a small percentage of us can after two decades of narrow, specific training utterly dedicated from the crib onwards to a very specific *context of decontextualisation*. But all humans who genuinely are epistemic agents own their cognition, exert control over their cognition, and so are capable of reasoning as best they can (and that means, with amazing ingenuity) to address the problems which confront them; thus, I contend, they possess our species' capacity for rationality and irrationality alike.

Notes

[1] See Plantinga (1993, Ch. 1) for a masterly review. Of course the deontic conception of internalism is not the only conception of internalism, but it is the conception that is operative in this paper.

[2] Bonjour (2003) was so taken by the power of the epistemic poverty objection that it led him to retrench from his earlier (1985) deontological conception of internalism—something I consider regrettable. The alternative (stock response) option which this section is about to articulate, is by far the preferable option here.

[3] This talk of "objective", "absolute" obligations, duties and the like is highly infelicitous. Since we cannot be held responsible for acting or failing to act on these "obligations"—cannot be *obliged* to act on them, etc.— they are then not *obligations, duties, oughts*, in anything like the ordinary sense: We see an example of apparently deontic terms being stretched far beyond what the interconnected deontic conceptual framework actually permits. The objective, absolute, contrast to Chisholm's subjective, practical, terms of requirement, would then be better not put in terms of "duties", "obligations", "oughts" and the like at all.

[4] This stock objection has to be taken seriously: major confusion beckons if it is not. Thus, Owens (2000) recognises the existence of this epistemic poverty argument against internalism, and acknowledges the presence of the stock internalist response: "Some adherents of the juridical [deontic] theory of responsibility … [might hold that] if our agency is the only thing we control then our responsibility cannot extend beyond our actions and the (foreseeable) consequences of our actions: we are responsible for our beliefs only in so far as we could, and should, influence them by means of our agency" (2000, 83). However, Owens then responds: "But … we cannot shrug off our epistemic responsibilities simply by highlighting the difficulties of controlling our beliefs by means of our actions". Talk of "shrugging off" here suggests a view of the internalist response as some kind of entirely optional (and reprehensible) evasion of responsibility, rather than a move forced on us by the power of the epistemic poverty objection. Owens demonstrates the origin of this dismissal of the problem of our epistemic responsibilities when he follows: "We believers are held to epistemic norms even when there is no action which we should perform with a view to implementing them. Irrational belief is not excused just because there is nothing the agent can *do* about it" (2000, 83).

[5] Goldman subsequently considers the application of his interpretation of this principle to a specimen "rightness criterion" from the internalist camp—and he chooses as his specimen Richard Foley. He objects that the subjectivism of a "Foley Rationality" approach "makes Foley's approach ill suited to the objectivist, nonrelativistic spirit of our entire framework" (Goldman 2009, 28). But this section has argued that such a subjectivism (such an "internalist perspectivism") is and must be a feature of any clear-headed internalist account. What then becomes of Goldman's claim that "this point should be equally acceptable to both internalism and externalism"? Against an internalist of Foley's stamp, I'd suggest that this comes worryingly close to begging the question.

[6] Externalists also have to draw this distinction—or something like—because of a standard objection to their position (that it doesn't offer us guidance, a "doxastic decision procedure"). *Cf.* Lockie (2014), Bales (1971), Brink (1986), Smith (1988).

[7] Note that this case is similar to many others, including some that are almost equally famous in the cultural psychology and anthropology literature; and although this type of example is emphasised by one particular "contextualist" (neo-Vygotskian) tradition, nevertheless, these data per se are not in dispute by other traditions in the literature.

[8] Be aware: some (a few) in the detailed psychological reasoning literature have started to use the Chomskyan competence/performance distinction in a far more restrictive sense than Cohen—operationalising "performance error" quite strictly in terms of momentary and *random* errors of attention, working memory, distraction etc. (notably Stanovich and West 2000). I'd rather they had not started to do this, since it removes from us a useful umbrella term ("the Uzbeks aren't logically *incompetent*, they've an intact logical *competence* but performance errors stop them from realising this in our decontextualised,

school-task tests"). Still, I flag for the reader the fact that I am using the competence-performance distinction in this earlier and broader way. Others still use it thus too.

[9] Such a political interpretation will surely not explain these data in their entirety; but this is not meant as a throwaway aside—methodologically, excluding plausible alternative explanations like this is a problem that has bedevilled work of this kind—even before one gets to the fraught issues of accepting the data and disputing the possible different theoretical interpretations of such cases. For example, in probably the next most famous case to Luria's in this literature (Cole's "Black Deer/Spider" example) Cole's Kpelle participants had a tradition of riddles that closely resembled the syllogisms they were being tested on—riddles that nevertheless would tend to lead them to a very different interpretation of the demand characteristics of the problem being set. Such issues expand into large and still disputed themes running through the history of anthropology and [cross-]cultural psychology. A point worth emphasising is that explanations of the responses of highly *contextualised* humans is precisely likely to permit of more nuanced, difficult to interpret, variability by *context* than are the responses of standardly educated decontextualised humans.

[10] "Because normative models are tools of rationality and because these tools undergo cultural change and revision, there is no idealised human 'rational competence' that has remained fixed throughout history" (Stanovich 2011, 269—of Cohen's competence-performance distinction). This constitutes one reason among others why epistemic rationality should be seen as a species-specific capacity. Rationality is intrinsically "culture-laden" and only humans have culture—at least, to any meaningful extent.

[11] There are important, and evolving differences between theories positing two Systems/Processes/Types/Minds, a discussion of which would take us beyond the scope of this paper. Psychologists are increasingly moving towards theoretically less committed, ontologically lighter characterisations of this distinction. What used to be called two *process* theory, later two *system* theory, is often now characterised more neutrally—say, after Evans (2009) as "Type 1" and "Type 2" processing (I shall follow this usage, which is becoming common). In this paper I remain agnostic on whether we should conceive of there being two Systems/Processes/Types/Minds in any sense (referring the reader to Evans (2010) and Evans and Stanovich (2013) for a discussion of the differences between these). In light of the concerns of this paper, it is worth noting that the latter state that "it is no longer the case that Type 2 thinking is regarded as abstract and context-free"; and "although the evidence supports the view that engagement of Type 2 reasoning is often necessary for the kind of abstract and elaborated forms of reasoning needed to solve typical laboratory tasks, such a conclusion does not make decontextualization a defining characteristic of Type 2 processing" (Evans and Stanovich 2013, 228, 229). Evans (e.g. Evans 2014, Evans and Stanovich 2013) identifies Type 2 cognition with cognition that engages working memory. In this he diverges from his recent collaborator Stanovich, who emphasises the human capacity for cognitive decoupling—though they claim that their core criteria are "fully compatible" (*ibid*: 226). From this, as an aside, two separate points are worthy of note. First, that Stanovich's position will thereby commit him to distinguish quite strongly "decoupled" from "decontextualised". Second, in light of remarks about decontextualisation and the Flynn effect below, that Evans' criterion for Type Two processing, "working memory" is, in the literature strongly bound up with fluid general intelligence g(F) (Conway et al. 2011)—with, for example, each of these quite commonly operationalised by Ravens matrices. Setting these two points aside for now, were one to embrace some version of a dual process distinction it would be likely that some of the differences between contextualised and decontextualised thought would occur with both of these operative at the level of Type Two thinking (*cf.* Evans and Stanovich *ibid*). However, some of these differences might be due to decontextualised thinkers applying Type Two thinking (or *overriding* their Type One thinking—*cf.* Evans, *passim*) while contextua-

lised thinkers applied only Type One thinking (something closer to the conception of matters found in the older heuristics and biases literature). The devil here will be in the detail, but to repeat: I remain agnostic on the issue of two Systems/Processes/Types/Minds anyway: neither embracing nor eschewing any version of this intriguing family of theories.

[12] In the case of such (e.g. undergraduate) participants, Stanovich and West's (2000) "understanding-acceptance principle" may be invoked to come to the rescue of the objectivist: "that more reflective and engaged reasoners are likely to affirm the appropriate model for a particular situation" (Stanovich and West 2000, 651—after Slovic and Tvesky). Make of that principle what you will; but in any event such agents will clearly be those who have the normatively appropriate response at least within their zone of proximal development.

[13] Cohort effects (the famous "Flynn effect") in the IQ literature permit of many explanations, and may well derive from several quite different sources of variance. However, presently the nutritional hypothesis and the cognitive stimulation hypothesis are widely seen as the leading contenders for an explanation (Neisser 1998: more first class protein and video games leading to IQ gains). I think that cultural/educational changes bringing about increasing *decontextualisation* and thereby leading to more rapid task construal is a far more plausible hypothesis for much of this variance (given how powerful the effect size is, and how short the time-scales involved); though the cognitive stimulation hypothesis may interpenetrate this, as an explanation. Early attempts by strongly nativist g theorists to dismiss the Flynn effect as an artefact predicted that it would be more marked in crystallised than in fluid intelligence (since the former is more obviously culturally loaded, and this tradition did not like to think of g as culturally situated at all). Counterintuitively, *exactly* the reverse was discovered (Flynn 1987, 1994). I think this is actually quite strong supporting evidence for the decontextualisation hypothesis—since fluid IQ (e.g. as operationalised by Ravens matrices) is precisely the most decontextualised IQ. "Decontextualised" does not mean "unacculturated"—see remarks about *context of decontext*, above. Our cultures of decontextualised thought over, say, the last two and a half millennia, are cultural-cognitive achievements of staggering depth, richness, power and scope. Their presumed potentiality in our genome does not alter the fact that they are expressed, at the most fundamental level, through and by profoundly immersive acculturation programs of great intensity and duration—such programs represent the difference between you and the Kpelle tribesman vis a vis modus ponens and the Ravens Matrices alike.

[14] The terms "emic" and "etic" originated in linguistics, spread to anthropology, and are now used across a number of disciplines. Roughly, *emic* accounts aim to uncover and then deploy in explanation the concepts that are meaningful to the members of a given society (here: their concepts of the rational norms). That is, accounts employing intra-perspectival concepts, descriptions and analyses: the subjective or insider's view. Note, the insider's view is not necessarily something that insiders are seen as being the privileged arbiters of (an outside, objective, observer—say, the anthropologist—may be needed to uncover the nature of these notions). A commitment to emic approaches then, should not in itself be seen as committing the researcher to relativism: though some strongly and —*per imposibile*—exclusively emic approaches ("*cultural* psychology") are relativistic, many are not. *Etic* accounts deploy extrinsic concepts and categories—explanatory concepts that have meaning for scientific observers without concern as to these notions tracking categories or mapping on to terms that are meaningful for, or found within, the conceptual scheme of the participants being observed ("*cross-cultural* psychology"). That is, etic accounts offer us extra-perspectival explanatory resources: objective, absolute explanans'. Most researchers acknowledge a need to employ both emic and etic explanans', with these to be used in an explanatory dialectic (something I would endorse). For

an example of a clearly non-relativistic yet explicitly emic account of cultural issues in the neighbouring area of intelligence research consider Sternberg (2004). Sternberg is also an example of two other things: an account that exemplifies the *dialectic* mentioned between emic and etic accounts, and the influence of Vygotsky.

[15] A word of caution: actually, in doing such research, partitioning out performance errors (and other issues that would not impugn agents' rational competence) yields a range of genuinely vexed, in fact genuinely *philosophical* problems, the solutions to which go way beyond mere skill in experimental design. That is, there are problems of avoiding incorrect task construal (e.g. Gricean issues of conversational implicature—and much else). There is the problem that many of the tasks these literatures trade in may have more than one normative solution. That is, many previously seen-as-incorrect responses may nevertheless be defensible in terms of "ecological rationality" (Gigerenzer and colleagues, Cosmides and Tooby, *passim*) or may otherwise involve psychologists wrongly imposing as the "correct" (normative) solution to a problem an answer which involved argument claims is not correct (e.g. neo-Bayesian/information-theory arguments that previously-considered-erroneous responses to the Wason task are in fact normatively appropriate—*cf.* Oaksford and Chater, *passim*). There are also issues of *ecological validity* more generally (the latter in Orne's (1962) sense rather than Gigerenzer's—*cf.* Hertwig, *passim*). "Rationalist" psychologists, for all that they are occasionally mocked as "Panglossians" have much to say in opposition to their "irrationalist" opponents here. These are deep and detailed issues, and I don't have to take a stand in these debates, though my rationalist sympathies should be fairly obvious to the reader.

References

Adler, J. 2000. Three fallacies. *Behavioral and Brain Sciences* 23: 665–6.

Alston, W. 1985. Concepts of epistemic justification. *Monist* 68: 57–89.

Bales, R. E. 1971. Act-utilitarianism: Account of right-making characteristics or decision-making procedure? *American Philosophical Quarterly* 8 (3): 257–65.

Bartlett, F. C. 1932. *Remembering*. Cambridge: Cambridge University Press.

Bonjour, L. 1985. *The structure of empirical knowledge*. Cambridge, MA: Harvard University Press.

———. 2003. Reply to Sosa. In *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*, edited by L. Bonjour, and E. Sosa, Ch. 10, pp. 173–200. Oxford: Blackwell.

Brink, D. O. 1986. Utilitarian morality and the personal point of view. *Journal of Philosophy* 83 (8): 417–38.

Cheng, P., and K. Holyoake. 1985. Pragmatic reasoning schemas. *Cognitive Psychology* 17: 391–416.

Chisholm, R. 1957. *Perceiving: A philosophical study*. Ithaca: Cornell.

Cohen, L. J. 1981. Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences* 4: 317–70.

Conway, A., Getz, S., Macnamara, B., and Engel de Abreu, P. 2011. Working memory and intelligence. In *Cambridge handbook of intelligence*, edited by R. J. Sternberg, and S. B. Kaufman, pp. 394–418. New York, NY: Cambridge University Press.

Cummins, D. 1996. Evidence for the innateness of deontic reasoning. *Mind & Language* 11: 160–90.

Elqayam, S. 2011. Grounded rationality: A relativist framework for normative rationality. In *The science of reason: A Festschrift in honour of Jonathan St. B.T. Evans*, edited by K. Manktelow, D. Over, and S. Elqayam, pp. 397–420. Hove: Psychology Press.

———. 2012. Grounded rationality: Descriptivism in epistemic context. *Synthese* 189: 39–49.

Evans, J. St. B. T. 2009. How many dual-process theories do we need? One, two or many? In *In two minds: Dual processes and beyond*, edited by J. Evans, and K. Frankish, pp. 33–54. Oxford: Oxford University Press.

Evans, J. St. B. T. 2010. *Thinking twice: Two minds in one brain*. Oxford University Press.

———. 2014. Two minds rationality. *Thinking & Reasoning*. 20 (2): 129–46.

Evans, J. St. B. T., and Stanovich, K. E. 2013. Dual-process theories of higher cognition: Advancing the debate. *Perspectives in Psychological Science* 8: 223–41.

Flynn, J. 1987. Massive IQ gains in 14 nations: What IQ tests really measure. *Psychological Bulletin* 101 (2): 171–91.

———. 1994. IQ gains over time. In *The encyclopedia of human intelligence*, edited by R. Sternberg, pp. 617–23. New York, NY: Macmillan.

———. 2011. Secular changes in intelligence. In *Cambridge handbook of intelligence*, edited by R. J. Sternberg and S. B. Kaufman, pp. 647–65. New York, NY: Cambridge University Press.

Goldman, A. 2009. Internalism, externalism and the architecture of justification. *Journal of Philosophy* 106 (6): 1–30.

Kahneman, D. 2003. A perspective on judgement & choice: Mapping bounded rationality. *American Psychologist* 58 (9): 697–720.

Kant, I. 1933. *Critique of pure reason*. Translated by I. Kemp-Smith. London: Macmillan.

Lockie, R. 2003a. Relativism and reflexivity. *International Journal of Philosophical Studies* 11 (3): 319–39.

———. 2003b. Transcendental Arguments Against Eliminativism. *British Journal for the Philosophy of Science* 54 (4): 569–89.

———. 2004. Knowledge, provenance and psychological explanation. *Philosophy* 79: 421–33.

——— 2006. Response to Anders Tolland's "Iterated non-refutation: Robert Lockie on relativism". *International Journal of Philosophical Studies* 14 (2): 245–54.

———. 2008. Problems for virtue theories in epistemology. *Philosophical Studies* 138 (2): 169–91.

———. 2014. The regulative and the theoretical in epistemology. *Abstracta* 8 (1): 3–14. ISSN 1807–9792.

Luria, A. 1976. *Cognitive development, its cultural and social foundations*. Cambridge, MA: Harvard University Press.

Neisser, U. ed. 1998. *The rising curve: Long-term gains in IQ and related measures*. Washington, DC: American Psychological Association.

Oaksford, M., and N. Chater. 1992. Bounded rationality in taking risks and drawing inferences. *Theory & Psychology* 2: 225–30.

Orne, M. 1962. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist* 17: 776–83.

Owens, D. 2000. *Reason without Freedom*. London: Routledge.

Plantinga, A. 1993. *Warrant: The current debate*. Oxford: Oxford University Press.

Plato. 1992. *Theaetetus*. Translated by M. J. Levett (revised Burnyeat). Indianapolis, IN: Hackett.

Sidgwick, H. 1907. Appendix to *The Methods of Ethics*. London: Macmillan. Revised from first publication in *Mind* 13, 51, (1888).

Smith, H. M. 1988. Making moral decisions. *Nous* 22: 89–108.

Stanovich, K. 2011. Normative models in psychology are here to stay. *Behavioral and Brain Sciences* 34: 268–9.

Stanovich, K., and R. West. 2000. Individual differences in reasoning: Implications for the rationality debate?' & 'Authors' response. *Behavioral and Brain Sciences* 23: 645–65 & 701–726.

Stanovich, K., R. West, and M. Toplak. 2011. Intelligence and rationality. In *The Cambridge handbook of intelligence*, edited by R. J. Sternberg, and S. B. Kaufman, pp. 784–826. New York, NY: Cambridge University Press.

Sternberg, R. 2004. Culture and intelligence. *American Psychologist* 59 (5): 325–38.

Vygotsky, L. 1978. *Mind in society*. Cambridge, MA: Harvard University Press.