Reply to John Skorupski

PETER RAILTON
University of Michigan

INTRODUCTION

I am grateful to John Skorupski for his thoughtful response to *Facts, Values, and Norms*. He understands my overall aims well, and so can pose some of the hardest questions I face. I will limit my reply to several of his key points, though even so, there is space here only to begin the response they merit.

FROM SUBJECTIVE TO OBJECTIVE, FROM FACT TO VALUE

After describing my attempt to make sense of normativity through an account of the aspirations and practices of our shared moral life, John Skorupski writes:

> I entirely agree . . . that focusing on these phenomenological questions, of what in our experience and discussion with others actually prompts us to basic normative judgements, is the most helpful way of demystifying philosophical worries about how normativity is even possible. (p. 228)

This characterization identifies a guiding idea throughout in these essays. We humans are a curious species. Since the earliest recorded times, we humans appear to have posed – and felt the need to answer – questions about our place in the world and the legitimacy of what we do. A capacity to articulate and respond to normative concerns is a deep part of us, including the capacity to feel the force of radical doubt. Should we take existing morality as a given when doing moral philosophy, or should we take skepticism seriously? I align with the latter view – history has witnessed the flourishing of many normative systems we now find dubious or even paradigmatically unacceptable. We obviously can't refute the skeptic, or escape the confines of our going world-view. But we can attempt to determine whether morality and its demands can be understood and motivated in independently intelligible, epistemically respectable terms. To me, this suggests beginning with elements that are empirically accessible and minimally controversial normatively.

In that spirit, I attempt to move within a broadly naturalistic framework from the familiarly *subjective* and prosaically *factual* to the *objective* and *normative*. This movement is, as Skorupski notes,

© 2008 Cambridge University Press
doi:10.1017/S0953820808003014
Utilitas Vol. 20, No. 2, July 2008
Printed in the United Kingdom

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

dialogical rather than purely conceptual. Normative concepts and objectivity seem to me irreducible to natural concepts and features of subjects. Instead of an outright reduction, then, I attempt to give a constructive account of morality that shows how subjective beings like ourselves are able to attain a degree of objectivity within a normative realm, and how prosaically factual features of ourselves and our practices can meet the job descriptions of moral properties. If these arguments succeed, the result will be what is sometimes called a *fundamental* or *vindicatory explanation* (for short, a *funding explanation*), namely, an account of how morality can arise and live up to its objective purport in a world of facts and subjectivities.

Skorupski rightly identifies the kinship of this project with Mill's much-criticized attempt to provide a naturalistic reconstruction of morality along empiricist lines. If I am to escape the powerful criticisms leveled at Mill, I must provide an account of the semantic content of moral thought and discourse that does not run afoul of Moore's 'open question' test, explain how a naturalist can capture normativity, and dovetail the account with a realistic account of human psychology and natural and social history. This is an absurdly ambitious set of goals. My more modest hope is that, insofar as I make some progress along these lines, morality will lose some of its mystery, and emerge as a domain of human thought and activity that can survive the reflective scrutiny of reasonable, well-informed people and win their open-eyed allegiance. I think this makes me what Skorupski calls a 'New Constructor', though one who believes that the construction itself will license a significant degree of 'Conservative Holism', since the construction shows why existing morality can be expected to embody a great deal that is true and important. Hence the Neurathian image with which I conclude, duck or boat.

PUTTING INSTRUMENTALISM IN ITS PROPER PLACE

Skorupski's analysis of my view assigns a pivotal role to the instrumentalist conception of practical rationality, which he speculates I fully endorse. He needs to speculate, since my desire to start with minimal normative assumptions led me to shy away from offering any account of rationality as a whole. Instead, I made claims like, 'for all its faults and all that it needs to be developed,' the instrumental conception is 'the clearest idea we have of what it is, at a minimum, to have a reason for action' [6, 47; quoted by Skorupski].[1] This, I now realize, is

[1] Page numbers in the text inserted in square brackets refer to my book, *Facts, Values, and Norms* (Cambridge, 2003); page numbers inserted in parentheses refer to John Skorupski's critical notice (this issue).

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

a confused way of characterizing what I was after, so I will try in what follows to be clearer. It would be unfortunate if my view *does* commit me to full-blooded instrumentalism, since that is a controversial and, to my mind, implausible view of practical rationality as a whole.[2]

We can think of full-blooded instrumentalism (hereinafter simply *the instrumental conception of rationality* or *instrumentalism*) as the view that, apart from assuring internal coherence in one's beliefs and preferences and assisting in inference and the discovery of facts, the role of reason in action is restricted to identifying and directing us to take the most effective means toward our ends. My mistake was to write as if one could treat instrumentalism as a 'minimal' core element likely to be shared by most fully developed accounts of practical rationality. Instrumentalism is instead a *maximal* conception, since, beyond commending coherence in preferences and accuracy in factual belief, instrumentalism *precludes* rational criticism of an agent's ends. For my part, I want to give a theory of just how such criticism is possible, though I could not assume it at the outset.

I should have detached the means/end principle (M/E) from instrumentalism:

(M/E) *Ceteris paribus*, reasons transmit from ends to the means that promote them.

What I wanted was a principle that 'captures a central normative feature of reason giving, since we can readily see the commending force for an agent of the claim that a given act would advance his ends' [6; quoted in Skorupski]. Unlike instrumentalism, (M/E) does not rule out the existence of non-instrumental reasons or the rational evaluability of ends. Nor is it controversial. Indeed, something like (M/E) is common ground among all the major traditions in practical rationality, Aristotelian, Humean, and Kantian.[3]

Is this because it is a mere truism, rather than a *bona fide* normative principle with which I might begin a construction of morality? Skorupski at one point suggests something like this. He considers the possibility that I might be able to rely upon a more modest principle than instrumentalism, namely, 'if there's reason to advance an aim there's reason to do whatever would advance that aim'. However, he claims, this would-be principle is 'arguably analytic', and so presumably cannot do the substantive normative work I need done (p. 228). (M/E),

[2] I therefore would reject Skorupski's suggestion that, for me, 'the truth-maker for the proposition that I have reason to X is the fact that X-ing would advance my aims' (p. 227).

[3] In fact, I borrow the formulation (M/E) from Thomas Nagel, *The Possibility of Altruism* (Princeton, 1970), p. 33. See also Aristotle, *Nicomachean Ethics*, 1112b13.

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

however, is not a principle bridging the admittedly tiny gap between 'reasons *to advance* an aim' and 'reasons *to do what would advance* an aim'. Rather, (M/E) presents a substantive normative connection between two quite distinct things, *ends* and the *means* requisite to them. It tells us, for example, that an act with nothing to recommend it in itself, like walking in circles carrying dumb-bells, may have rationality conferred upon it if it serves *y* for a worthwhile end, like healthy bones. (M/E), then, is neither tautological nor controversial – a good place for me to start.

But (M/E) alone won't get us far, since it yields no reasons at all until it is combined with ends that we have some reason to seek. A further, appropriately minimal principle is needed, something like (B/D):

(B/D) *Ceteris paribus*, agents have reason to accord their own current beliefs and desires *prima facie* (default, defeasible) action-guiding authority.[4]

Is (B/D) appropriately minimal, or is it too controversial, or too biased toward Humeanism, to serve as part of a minimal starting point? Anti-Humeans often claim that some desires or beliefs are so bizarre, perverse or self-defeating that no rational person would accord them any authority. (B/D) does not preclude this possibility, however, since a default, defeasible reason has only presumptive *ceteris paribus* authority, and thus can be contextually silenced or even self-cancelling.

Moreover, ask what would become of a would-be agent who was disposed *not* to accord even *prima facie* authority to his own beliefs and desires in their own right, who always insisted upon having some independent justification before taking their proffered guidance seriously. In what terms could such an individual provide herself a justification, if not those she finds antecedently credible? Suppose she were to propose as a form of independent justification subjecting her beliefs to a test of coherence or cogency. Well, either she trusts this test and her ability to apply it or she doesn't. If she lacks such trust, she will have no justification; if she is trusting, then this is yet another set of credences of hers, to which she cannot accord authority without an independent justification, launching a regress.

[4] 'Belief' and 'desire' in (B/D) are to be understood in the loose, encompassing manner now common in philosophy, as covering a wide range of conscious and non-conscious cognitive and conative states. Thus, implicit trust and unacknowledged suppositions count as beliefs, and devotion to the moral law and intrinsic valuing count as desires. '*Prima facie*' is used in many different ways by philosophers. Herein it will by stipulation be understood as 'default, defeasible'. Interestingly, something like (B/D) may be true without argument for those who analyze desire as a judgment or appearance of desirability, or belief as a judgment or appearance of credibility. However, I will need to argue for (B/D) because I reject these accounts. See my 'Kant Meets Aristotle Where Reason Meets Appetite', *Argument und Analyse*, ed. C. U. Moulines and Karl-Georg Niebergall (Paderborn, 2002).

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

A similar dilemma confronts an agent who insists that she will accord no desire action-guiding authority until it has been independently certified, perhaps by a test like universalizability. But how is she to view her favorable attitude toward acting only on universalizable desires? Does she defer to it immediately, without subjecting it to an independent test? Then by her own standards she is unjustified. But if, out of a sense that she must be consistent, she does demand an independent certification, she now will find she needs independent justification for according authority to *that* sentiment. And so on.

Thoroughgoing rejection of (B/D) thus threatens a would-be agent with inability to manage the ABCs of forming beliefs and stable intentions. Deliberation, decision and judgment need something to work with, something that doesn't first require deliberation, decision and judgment. This suggests a *weak present aim* principle:

(WPA) *Ceteris paribus*, action guided by deliberation and stable intention (and *a fortiori* deliberatively rational action) is possible only if individuals are disposed to accord *prima facie* action-guiding authority to at least some of their current beliefs and desires.

(B/D) is a strengthening of (WPA) from *some* to *all*. Consider things from the standpoint of a would-be agent facing a question of whether and how to begin deliberation – which beliefs or desires are to be allowed the *ab initio prima facie* action-guiding authority needed to get things underway? Would it make sense to her *at this point* simply to pick some subset of her beliefs and desires and ignore many of things she thinks true or cares about it? In practice, under the pressure of time, we of course always consider only a fraction of what we think or want. But absent such limitations, it seems most reasonable in principle to seek the widest possible reflective equilibrium in light of the whole truth as we see it, and everything that we feel matters.[5] *Then* we can show our discernment by how we go on to assess, weigh, cull, act on or revise these beliefs and desires as we acquire new ideas and undergo new experiences. (B/D) simply holds that rational agents have the disposition that accords with this sensible advice.

Let me, then, replace misleading talk of starting out with instrumentalism 'at a minimum', with an initial commitment to (M/E) and (B/D). They strike me as minimal enough, but contentful enough, to get my constructive account of intrinsic non-moral value and well-being underway, and thus take the first steps toward morality. I will, however, need two more principles before going much further:

[5] In epistemology, this is known as the 'principle of total evidence'.

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

(Exp) *Ceteris paribus*, rational agents accord greater action-guiding authority to those beliefs and desires they form as they gain greater information and wider experience.

(Refl) *Ceteris paribus*, in rational agents the dispositions corresponding to (B/D) and (Exp) tend to survive reflection, even as information and experience grow.

Combining these elements, it is possible to carry out a stepwise approach to intrinsic value, beginning with what people actually believe and care about, and proceeding through the evolution of these attitudes and interests in light of increasing information and wider experience. Each step in this evolution would be normatively intelligible to the agent (assuming (Exp) and (Refl)), and therefore would, *ceteris paribus*, preserve reason-giving force. If this evolution at some point stabilizes, then we have arrived at what would matter to an individual in its own right even when she fully understands its nature and her own.[6] This, I claim, is the best evidence we can have of what would give her life real, and not merely apparent, value.[7] The dialectical strategy here obviously traces back to Hume on taste and Mill on desirability.

As we noted earlier, Skorupski also favors a dialogical approach to explaining the source and nature of normative force, but he asks pointedly whether such an approach could warrant the *realism* about value I advocate – surely dialectics are better suited to a more constructivist understanding of value and objectivity? Here is the beginning of an answer. What appears to underwrite the strategy in Hume or Mill, and to provide the basis for their confidence that the dialectical process will have some definite, stable upshot, is a conception of aesthetic or intrinsic value as arising from a natural *match* or *fit* between individuals and their capacities, on the one hand,

[6] Here I simplify, of course. The full account involves higher-order attitudes as well. Even so, Skorupski has reason to be doubtful. 'The problem is we don't know whether [the ideally informed individual $A+$] would have *any* desires about A's wants as he actually is.' What's wanted from $A+$, however, is not what $A+$ himself desires with respect to what A is to want. Neither is it 'what $A+$ would want *if he had A's good at heart*'. (8) Rather, it is what $A+$ would want to want in the hypothetical circumstance of finding himself in A's shoes, stuck with living the rest of A's life. $A+$ will not be indifferent regarding what he wants to want in such circumstances.

[7] Skorupski asks, appropriately, just what sort of value this is. Is it well-being? Or something else? And is it the value that morality should take into account when assessing her situation or counting her welfare? These questions require us to distinguish the value of a life for the person who leads it (which most closely corresponds to well-being) from the values realized through a life (which most closely corresponds to what, from the standpoint of the agent, gives her life significance and worth). Morality takes both into account, though in different ways and with respect to different attitudes. Speaking roughly, questions of justice and benevolence concern the first, while questions of personal respect and self-respect concern the second. For further discussion, see my 'Cost-Benefit Analysis as a Source of Information about Welfare', *Valuing Health Risks, Costs, and Benefits for Environmental Decision Making*, ed. P. Brett Hammond and Rob Coppock (Washington, D.C., 1990).

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

and certain activities, experiences, relationships, and forms of life, on the other. Hume, for example, speaks of certain qualities as 'fitted by nature' to bring down an appreciative response when attended to by a sensitive individual.[8] In a similar vein, Mill speaks of discovering what is genuinely desirable in life – what is 'grounded on the permanent interests of man as a progressive being' – through the experience of 'experiments in living'.[9] In both cases, the inherently compelling, intrinsically rewarding character of 'finding fit' affords a real-world translation of the otherwise mysterious 'normative force' of intrinsic value into actual lived experience.[10] If we understand objective value in realist terms – as a matter of such relations of fit – we have more than a dialogue. We have a funding explanation of how the dialogue, and associated practices, can be both a source of genuine knowledge and the discovery of a domain with robust normative allure.

Skorupski is right to highlight the ways in which starting out with a minimalist conception of rationality shapes my work. My hope is to start the search for a funding explanation 'close to the ground' of everyday life, not with an ambitious theory of the demands of reason but with our experience of what makes our lives go well or ill, and our sense of the reasonableness of acting so as to promote or protect the things that matter. (M/E), (B/D), (Exp) and (Refl), taken together, embody this minimalism. Do they amount to an objectionable instrumentalism? It is hard to see that they do. For although they provide an initial core of reasons centered on an agent's present beliefs and current aims, the outer edge of this core is not the *horizon* of the agent's potential reasons. Rather, it is the jumping-off point for an agent to begin experimentation with a wider world of facts, values and reasons. Surprisingly, perhaps, this minimal, personal approach can bring us to embrace impartial, non-hypothetical norms for belief, action and appreciation. These norms gain their standing for us in part *because* they grant no privilege to anyone's particular standpoint or preferences, including one's own. Only by taking up this wider, less subjective point of view can we speak in the name of *value* rather than *liking*, and

[8] David Hume, 'Of the Standard of Taste', *David Hume: Essays Moral, Political, and Literary*, ed. E. F. Miller (Indianapolis, 1987), p. 235. He also writes that the sentiment of beauty 'marks a certain conformity or relation between the object and the organs or faculties of the mind' (p. 230).

[9] John Stuart Mill, 'On Liberty', *John Stuart Mill: Utilitarianism and Other Essays*, ed. M. Warnock (New York, 1962).

[10] Recent empirical work on 'optimal experience' draws upon a similar idea of the inherently compelling, intrinsically rewarding experience of 'finding fit' between activities and fully utilized human capacities. See *Optimal Experience*, ed. M. Csikszentmihaly and I. S. Csikszentmihaly (Cambridge, 1988). Similar work has been done in neurophysiology and comparative psychology concerning universals in aesthetic appreciation. See the special issue on music of *Nature Neuroscience* 6 (2003).

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

express the sort of respect that we feel value is due, whether it be realized in our own lives or the lives of others.[11]

PUTTING INTERNALISM IN ITS PROPER PLACE

In the case of the means/end principle, then, the interesting question about theories of practical reason is not *whether* they incorporate it, but *where* and *in what role*. The same is true with the much-disputed question of 'internalism'. A moral theory that nowhere incorporated an 'internal connection' with actual human concerns or motives would be hopelessly alienated. Where moral theories differ is in *where* and *in what role* internalism shows up.

Expressivists, for example, locate the internal connection in the content of moral judgment, which they see as functioning to express a pro-attitude in the mind of the speaker. I believe this puts internalism in the wrong place. The expressivist does have a very straightforward explanation of the 'action-guiding force' of my moral judgments *on me*, since the state of mind they express is motivating. But Kant points out that a mark of a normative judgment, as opposed to an expression of personal preference, is that it purports to give a reason to others, indeed, to provide a reason for all of us alike. What ground would I have for thinking that my moral judgment makes the same normative claim upon others as it does on me, if its normative force for me resides in the motivating force of the pro-attitude my judgment expresses? Perhaps my pro-attitude favors only claims that can engage others motivationally on equal terms, and perhaps I am appropriately disciplined in forming this attitude. But none of this is essential to moral judgment on the expressivist account. We seem to be missing an adequate account of why moral judgments have the authority they do within our social lives.

Skorupski advocates a cognitivist approach according to which moral judgments express 'normative propositions', which have an internal connection with practical reasons, such that 'one cannot consistently accept [such a proposition] and yet deny that it affords a reason' (p. 227). This, of course, is trivially true if the reason is a *moral* reason, so it must be a reason with normative force in its own right. The question now becomes, is *motivation* internal to the recognition of such a reason – and if so, whose? If its motivation were altogether

[11] Here my views converge with those of Thomas Nagel in *The Possibility of Altruism*, ch. 12, and *The View from Nowhere* (Oxford: Oxford University Press, 1986), ch. 8. Of course, I believe this objective standpoint is best characterized by a valoric consequentialist understanding of the touchstone of morality; Nagel's interpretation is more Kantian.

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

external, the worry would arise that the reason in question cannot really be practical. If its motivation is internal on the part of the speaker, then an *a priori* necessary connection has been forged between *making a cognitive judgment about the existence of a practical reason*, i.e., asserting a proposition with standard truth conditions, and *being in a motivational state*. This leaves us with the puzzle about what such truth conditions could be that expressivism was devised to solve. Fortunately, there is a third possibility: the internal connection can belong to the truth conditions for the judgment if the motivation in question is internal to the *subject* of the judgment. That is, the '*X* has a reason to *F*' can be true just in case *X* has some motive to *F*.

This is the sort of internalism I favor. It helps explain the Kantian symmetry of normative judgment mentioned above: when I judge that *X* has a reason to *F*, I am committed to there being an appropriate motive for *X* to *F*, whether *X* is myself or someone else.

This sort of 'existence internalism' recognizes two species of reasons: *subjective* and *objective*. Subjective reasons support directive judgments of what an agent *ought* to do or would be subject to rational criticism for failing to do, so they must lie within the agent's epistemic ken, and must have the power to move him to thought and action more or less as he is; otherwise, we cannot expect him to see the rationale of the *ought* or hold him responsible for failing to feel its action-guiding force. Objective reasons also have an internal connection to the agent's motivational structure, not as it actually is, but as it would be under an idealization of his information, imagination, experience, acumen and understanding. Objective reasons support evaluative judgments of what would *be best* for the agent – what is rational for him 'in the extended sense' as I called it in the book [cf. 20ff.]. Failure to act on one's objective reasons can occur with no fault in one's epistemic agency or deliberative rationality. How is the man in Bernard Williams' example who is about to drink a toxic cocktail to know that someone has used the gin bottle in the cupboard to store gasoline?[12] The man has good, motivating reason in the objective sense not to drink it, but he is not irrational for failing to act accordingly. Are objective reasons 'explanatorily irrelevant' as some have alleged, since they lie outside the individual's ken and thus cannot explain why he acted as he did? No. Reasons unknown to the agent can shape and color his experience, and thereby steer one in the direction of discovering objective reasons, as this man is about to discover.

Skorupski describes a 'bipolarity' in my view: there is a potential gap between a person's good and what he has reason to pursue. This he thinks is at odds with the most natural way of thinking, according to

[12] Bernard Williams, 'Internal and External Reasons', in *Moral Luck* (Cambridge, 1981), pp. 102f.

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

which it is necessarily true that, 'there is reason for *A* to pursue his good, whatever that might be'.

I am not sure how to understand this claim. It certainly can't be true of a person's subjective reasons. Skorupski acknowledges this in the case of moral reasons: 'my moral obligations are something I can tell for myself I have, so that moral obligation, unlike optimality of consequences, must be information-relative' (p. 225). It must, therefore, be a thesis about objective reasons. Here it looks promising. If objective reasons are those the agent would recognize or appreciate under ideal circumstances, and if I am right that a person's good is what *best explains* his ideal preferences (with regard to the course of his own life), then there will be a metaphysical connection between one's objective reasons and one's good.

This would seem to remove any worry about bipolarity, but Skorupski argues that the worry arises again at the level of the modality of the connection. It should, he believes, be *conceptual* not metaphysical – it belongs to the concept of a person's good that he has (objective) reason to pursue it, whatever it is.

Is it a bipolar disorder to deny this, as I do? Although many find an analysis of a person's good into what he has reason to pursue plausible, do those of us who think that the concept of a person's good is *sui generis*, or explanatorily prior to the concept of a person's reasons for action, make a *conceptual* mistake? It seems to me, for example, that a deeply pious person would not be conceptually confused if she agrees that what is best for her is to live well, but at the same time insists her *reason* to promote her own well-being is contingent on whether doing so would serve God's will or celebrating his glory. Similarly, I do not believe Kantians have abused the concept of well-being in arguing that acting on *any* subjective maxim, including the maxim 'Promote my own well-being', is rational only if it could be willed as a universal law.

REASONS FOR BELIEF

I now am able to respond to Skorupski's questions concerning the relation between *reasons to believe that P* properly so-called versus *reasons to bring it about that one believes that P* of whatever sort. He believes these are two distinct concepts, as do I. He further believes that it would be 'immoderate' and at odds with common sense to allow reasons of the second kind to supplant reasons of the first kind. Am I immoderate in this way? It certainly sounds that way when I write, 'questions of what it would be rational to believe belong to practical rather than theoretical reason' [35]. Wouldn't this allow wishful thinking to replace epistemic norms, so long as it proved more satisfying?

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

Consider the distinction between logical arguments and human inferences. The first are formal patterns that say nothing about how anyone should think; the second are psychological processes that sometimes do and sometimes do not conform to logical patterns. Let's take Skorupski's example of global climate change. Suppose that I observe that winters are getting warmer, see photos illustrating the polar ice cover in retreat, and hear that the dominant view among scientists is now that global warming is the root cause. Now I ask myself, what should I make of this? I might consult the principles of logic, and draw any of an infinite number of deductive consequences from these premises. None of these consequences, however, will get me to the conclusion that global warming actually *is* taking place – any such inference would be content-extending rather than deductive. Alternatively, I might have confidence in scientific opinion, especially when it shows wide convergence, and think that the best explanation of why winters are warmer, the ice cover is shrinking and scientists are largely of one mind is that global warming really is occurring. In the case of drawing out logical implications, I defer to some very definite norms and rules about which there is little or no controversy; unless I slip up, my conclusions will have at least as much credibility as my premises. In the case of my 'inference to the best explanation', I defer to vaguer, much more controversial norms, and the conclusion I draw far outstrips the content of my premises, and therefore is significantly less probable.

It sounds, then, as if all the normative advantage lies with narrowly logical reasoning. If belief by its nature 'aims at truth', why would I go in for risky, content-extending reasoning when deduction would yield an indefinite supply of truths? The answer seems clear enough – in deciding what to believe, we care about things other than truth: relevance, informativeness, explanatoriness, unification, economy, comprehensibility, understanding, predictive power, fruitfulness, and so on.

As Quine and others argued, even in the paradigmatically respectable domain of natural science, theory choice is massively underdetermined by 'narrowly evidential' or confirmatory criteria. Yet we hardly wish that scientists would just stick to the evidence and its implications. What sort of reasoning is it, then, that leads scientists and common-sense thinkers alike to think it appropriate to rely on non-confirmationalist criteria in moving from modest evidence to ambitious conclusions? *Means / end reasoning relative to the wider goals of inquiry*, whether scientific or common-sense. And that was my point. We reason not a priori, from epistemic first principles, but a posteriori, in light of the track record of various regulative norms for inquiry. Inevitably, in doing so, we rely upon some of the very norms whose benefits and costs

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

we seek to assess. So once again we are in the realm of dialectics rather than rational demonstration. In effect, we reach working conclusions about the expected value of various ways of regulating inquiry and according epistemic authority. It is in light of such conclusions that we think it rational to prefer everyday and scientific practices that take non-confirmationalist criteria into account and permit the formation of general, predictive theories, as opposed to a restriction of inquiry to 'narrowly epistemic' norms.

Although this reasoning is hypothetical, it leads us to adopt non-hypothetical norms of inference and inquiry. Mathematics, for example, would be impossible if mathematicians did not impose upon themselves non-hypothetical standards of logical consistency and validity; the spectacular growth of science would never have occurred if scientists did not impose upon themselves non-hypothetical norms of experimental procedure and quantitative statistical evaluation of results. In this way we are led *back* from human inference and its ambitions to logical argument, which we take as a *model* for our reasoning, thereby giving logic a normative standing in relation to belief and inference that it lacks in itself.[13] The normativity comes from us, as we discipline ourselves and others to follow logical patterns because we see the necessity of this if we are to achieve our most elementary theoretical and practical ends.

Thus within a hypothetical framework we develop a working distinction between *reasons to believe that P* proper and other sorts of *reasons to bring it about that one believes that P*. Interest-driven bias and self-fulfilling belief are always a danger, and erroneous expectation can be costly. A lot – indeed, everything – is at stake. Most of us never face grand choices of strategies of inquiry, posed as such. Rather, we face an endless succession of small-scale versions of them in deciding what or whom to trust, what conclusions to draw, when to venture a prediction or explanation, and so on. For all of us, there is a massive disproportion between the rich, indeed infinite content of our beliefs and the very finite amount of evidence and argument we can offer for it. On the other hand, we readily appreciate how badly hobbled we'd be as agents and inquirers if we stuck to the narrowest interpretation of reasons to believe. So we take risks with the truth for the sake of wider cognitive and practical goals, making our peace with a more tolerant, more ambitious epistemic regime. Recognizing

[13] Wittgenstein, I believe, takes just this sort of line on the nature of the logical *must*. He credits the basic idea to Ramsey. See L. Wittgenstein, *Philosophical Investigations*, trans. G. E. M. Anscombe (London, 1953), § 38. For discussion, see my 'Wittgenstein on the Normativity of Logic', *New Essays on the A Priori*, ed. P. Boghossian and C. Peacocke (Oxford, 2000).

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press

these epistemic compromises does not destroy belief, or lead us to cease trusting one another or ourselves. *That* would indeed be an immoderate conclusion.

PUTTING GARY IN HIS PROPER PLACE

This brings us to Gary. In my example, Gary is somewhat obnoxiously asking us to face up to the fact that, at bottom, following norms is an *activity*, *a matter of choice*. It involves us in substantive commitments about what matters and how things are, not just analysis of our own conceptual scheme. Thus, he argues, even when the norms to which we are committed are themselves non-hypothetical, our rationale in following them will always be to some extent hypothetical.

He has a point. What might appear *a priori* to be a rational requirement, such as deductive closure in belief, we may simply disregard in order to get on with our ambitious theorizing and lives. An *a priori* principle like the principle of sufficient reason falls by the wayside if it stands in the way of accepting a promising new science of the atom. Where Gary goes wrong is in thinking that his point calls into question the basic framework of practical and theoretical reason. Far from challenging the normative authority of principles like (M/E), (B/D), (Exp), logical consistency, etc., his argument and his self-understanding depend upon these principles for their own intelligibility and normative force. In arguing with us, he, too, accords authority to these non-hypothetical norms, for he sees himself as giving us reasons to accept his claims that we cannot rationally ignore simply because that would be convenient for us.

It seems to me a way of avoiding confusion to separate the question of what *authority* we attribute to a norm from the question of *rationale* we can give for that.[14] If the rationale is for following a non-hypothetical norm, then *that* is the sort of authority reason tells us to accord it. This is an application of (M/E), and so a piece of practical reasoning. But we'd be flying in the face of the very reasoning we are trying to follow – and so violating (M/E) – if we took this to mean that practical reason is to *supplant* theoretical reason. On the contrary, it *funds* it.

prailton@umich.edu

[14] This echoes the time-honored distinction between a *condition* for a requirement and a *conditional requirement*.

https://doi.org/10.1017/S0953820808003014 Published online by Cambridge University Press