

REFERENCES

- Craig, E. J., 1982. Meaning, Use and Privacy, *Mind* 91: 541–64.
Dummett, M., 1978. The Philosophical Basis of Intuitionistic Logic, in *Truth and Other Enigmas*, London: Duckworth.
Kripke, S., 1982. *Wittgenstein on Rules and Private Language*, Oxford: Blackwell.

Inventing logical necessity*

CRISPIN WRIGHT

I INTRODUCTION

1. The topic of this paper is the *objectivity* of logic; specifically, the objectivity of the relation of logical consequence and of the notion of logical proof. Both, as ordinarily conceived, implicitly involve the idea of logical *necessity*: if B is a logical consequence of $\{A_1 \dots A_n\}$, then – according to our ordinary conception – if the latter are all true, B *must* be true; likewise, if a structure constitutes a valid logical proof, then its starting point and its successive operations are such that what eventually results *must* result if exactly those operations are correctly carried out on exactly that starting point. The latter consideration is what, according to our ordinary thinking, essentially distinguishes a proof from an experiment (and even if a similar-sounding claim can be made of certain very well-established experimental routines, we think of the ‘must’ involved as being quite different).

The notion of logical necessity is not in good standing among many contemporary philosophers, but the perception does not seem to be widespread that, for such simple reasons, our intuitive conceptions of logical consequence and of logical proof must fall with it. Perhaps some appropriate refashioning of those notions is possible. But in this paper I shall work within the framework of the intuitive conceptual connections outlined. The strategy will be to make a case for a certain sort of doubt about the objectivity of necessity, and then to let that doubt transfer, via those connections, to the subject matter of logic. If the doubt is sustained, we shall find in favour of those philosophers – most notably, the latter Wittgenstein – who have urged that proof in logic ought not to be viewed as a medium of discovery of a special category of fact, and that logical relations do not stand independent of our cognition of them in the manner of, say, spatial relations among material objects. Wittgenstein’s own distrust of the opposing belief in the objectivity of logic is closely connected to that cluster of

* I am grateful for comments and criticisms received both at the Lyme Regis Thyssen conference and at the Universities of Belfast, Manchester, Pennsylvania and Harvard at which versions of this material were presented.

challenges, against certain intuitive preconceptions about meaning and understanding, to which in (1980) I gave the rather pedestrian title of the 'Rule Following Considerations'. The critique which I shall outline here also has Wittgensteinian roots; but it succeeds or fails, I believe, independently of the ideas on rule following.

2. Let me begin with some brief remarks by way of elaboration of the conception of logic which is at stake. Eventually we shall require a more refined account of what a believer in the 'objectivity' of logic is committed to, but the following rough characterisation of a familiar syndrome of ideas will serve for now. The believer will likely accept each of the following.

(A) There is a special category of truths which could not be conveyed in any language from which was absent (the means for defining) a unary sentential operator equivalent to 'it is logically necessary that...'. That is, some statements just are logically necessary truths; a language which failed to contain the means for affirming their necessity would, in consequence, fail to contain the means for saying everything true that can be said. The task of formal logic is, as far as possible, to codify algebraically this type of truth.

(B) The 'spectator conception' of proof; if, for example, we are applying a decision procedure to some formula of monadic predicate calculus, then the idea would be that we have only a *passive* part to play (cf. Dummett 1959); that what constitutes correct implementation of the procedure at every stage, and its eventual outcome, are predetermined – not causally but conceptually – by the character of the procedure and the identity of the tested formula. There is a similar predetermination even in cases where no effective decision procedure is to hand: of any well-formed formula of predicate calculus with identity, for example, it is, so to speak, stored up in the specification of the system whether an admissible proof of that formula can be constructed within it. What is *possible* in logic is laid down from the outset; and laid down, to repeat, purely conceptually, independently of any neurological or cybernetic considerations. The task of the logician is to unpack the store, and to make an inventory of its contents.

(C) Less figuratively, the logician is a *scientist*, his task one of discovery. His project is to chart the extensions of *logical necessity*, *logical consequence* and *cogent argument* (valid proof). These notions have determinate extension, fixed independently of his investigations, every bit as much as the concept, 'mountain exceeding 20,000 feet in height' has an extension fixed independently of the investigations of the terrestrial geographer. The difference is only that the extensions of the concepts in which the logician is interested could not have been otherwise than as they are.

There cannot be many students of logic who have never felt the tug of these ideas. Nevertheless they should be resisted if the drift of section III

of this paper is correct. First, however, an account is owing of why an attack is still in point: why, in particular, the objectivist syndrome was not routed a third of a century ago by Quine.

II TWO DOGMAS OF EMPIRICISM

3. Quine's famous paper, it will be recalled, mounts a two-stage attack on the notion of an *analytic* statement – a statement whose truth is settled purely by the meanings of its constituent expressions and the way in which they are put together. The first two-thirds of the paper are devoted to arguing that the concept of analyticity eludes satisfactory explanation and so is not, presumably, fully intelligible. Then in the last third of the paper, where the famous holistic picture of language is presented, it is contended that, to the extent that the intention of proponents of the notion of analyticity is clear, there are actually no analytic (or synthetic) statements.

Now, the formulations above involved the notion of necessity rather than that of analyticity. For Quine himself there is nothing at stake in the contrast (1961: 29–30). But, should any reader scruple over the switch, it is easy enough to see that the intuitive notions of logical proof and logical consequence depend upon the notion of analyticity in a manner very similar to the dependence upon necessity noted above. For *B*'s being a logical consequence of $\{A_1 \dots A_n\}$ is standardly taken to involve the analyticity of the conditional: if A_1 , then if A_2 , ... then if A_n , then B . Likewise the status of the structure as a formal proof would standardly be taken to require the analyticity of the *corresponding descriptive conditional*: that conditional, that is, whose antecedent hypothesises correct implementation of a series of specified procedures on a certain initial basis and whose consequent specifies a certain outcome.¹ Accordingly Quine's attack upon analyticity, whether or not it is *eo ipso* an attack upon the notion of logical necessity, threatens similarly destructive consequences for the intuitive set of beliefs in the objectivity of logic. How may the logician be conceived as a scientist aiming to map the domain of a special category of truths if the notions of logical consequence, and of proof, enter into contaminating relations with the very notion which Quine sought to discredit? Quine himself originally seemed to see no threat to logic from his attack, believing that both logic and mathematics would secure an appropriate dignity by placement at the deeply entrenched core of the totality of empirical science. But the threat is there; and I shall try to indicate briefly below why Quine's holistic empiricism is no satisfactory response.

4. How strong, though, is Quine's attack? Let it be true, for example, that

¹ Cf. Wright 1980: 454. We shall consider a simple example below.

analyticity, *syntheticity*, *meaning*, and other cognate concepts interlock in a circle, no one member of which is explicitly definable without recourse to the others. Still, no conclusion seems to be warranted about their intelligibility unless we have reason to suppose that the situation prevents the construction of any coherent model of how an understanding of (any one of) those concepts might be acquired. Yet that supposition will be justified only if we have reason to suppose that the route into each of these concepts has to be by explicit definition. Do we have any reason to suppose so? At least part of the role of the notion of meaning, for example, is as a theoretical concept in the explanation of linguistic behaviour – and no one since Carnap has expected that it should generally be possible to explain theoretical concepts by explicit definition in terms of concepts of a somehow less problematic status. Besides, if the challenge is merely to indicate how the notion of analyticity can possibly be explained, what is wrong with the obvious answer, appropriate to most of our concepts: by example, and by immersion in linguistic practices in which the concept is in play?

Quine, like Socrates, seems to have supposed that the absence of any clear, non-circular definition of a concept somehow calls its propriety into question. The proper response is that it does nothing of the sort, provided there is independent evidence that the concept is teachable and is generally well understood. A sceptic about the intelligibility of a concept does not have to be answered by a rigorous explanation of it; it is enough to supply unmistakable evidence that the concept *is* well understood. And what better evidence could there be than that there is manifest in the community at large a disposition towards non-collusive assent in the application of these concepts – crucially, in our application of them to previously unconsidered cases? (cf. Grice and Strawson 1956.)

Strategically Quine ought, it seems, to have denied that there is any such manifest disposition. But the nearest he came to such a denial was to dwell on the unclear status, in point of analyticity, of certain examples like ‘everything green is extended’ (1961: 32). This consideration is not to the purpose. It is quite consistent with our possessing a genuine concept that in certain cases we hesitate over its application or contradict one another. Quine requires that such cases be typical. But they do not seem to be so. What is evident is that if there is indeed among speakers of English, who have had a certain sort of standard training in logic and philosophy, a disposition towards non-collusive assent in the application of ‘analytic’, ‘logically necessary’, and their kin, then it is folly to deny that we have *some sort of* concept of analyticity, *et al.*, and that any argument which, if allowed, establishes that we do not must therefore contain error.

5. We noted, however, that Quine stops short of suggesting that we have no genuine concept of analyticity at all. The central thrust of the last third

of ‘Two Dogmas’ is rather that such concept as we have has no instances. For Quine, belief in analyticity is thus rather like belief in witchcraft: the central concept is poorly explained, and the practices which are based on the belief that it applies to anything are based on a mistake.

Quine writes, in a famous passage,

It becomes folly to seek a boundary between synthetic statements which hold contingently on experience, and analytic statements, which hold come what may. Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (1961: 43)

What exactly is the argument here? Analytic statements, Quine evidently supposes, whatever other properties they may have, ought to be immune to revision. Whereas, he urges, our total corpus of belief has a kind of holistic elasticity which involves that no statements are immune to revision. But it is hard to see why a defender of analyticity should wish to resist the suggestion that logic, or other disciplines conceived to involve analytic statements, are revisable. What a defender fundamentally wishes to maintain is that the truth of some statements is generated purely by the semantic machinery of the language. He has no legitimate interest in maintaining that we cannot be *in error* in judging a statement to have that status. He can therefore give Quine the claim – however implausible it may seem in certain cases – that any particular statement which we accept as analytic could, in certain circumstances, reasonably be discarded. For to grant the claim need be to grant no more than that our assessment of any particular statement as analytic may always in principle turn out to have been mistaken. Quine, like so many writers on the topic,² has slipped into thinking of analyticity as involving indefeasible certainty. To claim that a statement is analytic, however, is only to make a claim about the *kind* of truth it has – there is no immediate reason why the claimant has to agree that, when statements are analytic, their truth may be known with special sureness.

That we have the practice, then, of very occasionally effecting revision among beliefs formerly regarded as analytic is no argument for Quine’s view of the matter. Indeed it is quite irrelevant to Quine’s claim, properly understood, whether we ever carry out such revisions or not. The crucial question is rather what latitude is left to a rational subject by his experience

² Including the later Wittgenstein in off-guard moments. But contrast, for example (1969: §651).

when he endeavours to mould a system of beliefs adequate for explaining and predicting that experience. A defender of the traditional distinction between analytic and synthetic statements will hold that there are certain beliefs – the analytic ones – which, no matter what the course of the subject's experience, he cannot rationally be constrained to discard; and others – the synthetic ones – which, should his experience take a certain course, he must, rationally, discard.

Quine's claim – a generalisation, just as he says (1961: 41), of what Duhem held concerning scientific hypotheses – is that, such is the mode of functioning of our system of beliefs, experience *cannot* bear on the acceptability of single beliefs in the manner which the defender of the traditional distinction endorses. Only in the context of a theory, with its underlying logic, can experience confirm or disconfirm a particular belief; and, in principle, it is possible that the most fruitful response to a range of experiences may be to modify the underlying logic rather than the non-logical part of the theory.

Philosophy is still, I think, some distance from the stage at which it could be claimed that we have adequately taken the measure of the holism which Quine expressed in his classic paper. Unquestionably there are the most profound implications for the theory of meaning and for our view of the nature of truth if Quine is right. At present, however, my only concern is whether this holism, whatever other insights it may prove to contain, can be made to yield a satisfactory philosophy of logic. I want to suggest that it cannot. The reason is one which I have tried to express elsewhere (1980: 327–30).³

Schematically, Quine's picture is somewhat as follows. Let θ be a theory, and L its underlying logic. Suppose that from θ we can derive, via L , a conditional, $I \rightarrow P$, whose antecedent describes certain initial conditions and whose consequent formulates a prediction relative to those conditions. Now suppose that we suffer a barrage of experience, E , which is *recalcitrant*: that is, I suppose, it inclines us to assent both to I and to the negation of P . Now, since there are no synthetic statements, it cannot be the case that E forces our denial of P ; likewise, since there are no analytic statements, it cannot be the case that we are rationally prevented from pointing the arrow of suspicion, as it were, at L . Thus the theory, plus its underlying logic, plus our 'observational' responses to E confront the recalcitrant experience *en bloc*; and how it is best to respond is to be determined by pragmatic criteria, applied to the belief-systems which respectively result from the variety of responses open to us.

It wants remarking, to begin with, that this schematic picture is

³ I hope that I have improved, in this paper, upon a presentation which succeeded only in being suggestive of what I believe, properly formulated, is a conclusive objection.

incomplete. Experience E is recalcitrant for θ -with- L only if it is presupposed that $\theta \vdash_L I \rightarrow P$. This statement – call it W – will have been established by constructing a derivation of $I \rightarrow P$ from θ using L ; and would ordinarily be conceived as, if true, analytic. Clearly, however, the acceptability of W need in no way depend on the acceptability of L ; ratifying the proof of W simply consists in checking that $I \rightarrow P$ does indeed follow from θ by a series of L -accredited steps – no endorsement of the principles of L need be involved in so working with the notion of an L -accredited step. So the very description of E as recalcitrant for θ -with- L presupposes acceptance of a statement which is established by proof, which is analytic if any statements are, and which is independent of L .

The reader may now be expecting to be presented with an argument to the effect that the notion of analyticity is implicit in the very idea of recalcitrance. But, supposing that is true, it would not be a very destructive conclusion against the Quine who holds not that the notion of analyticity is unintelligible but that it requires holistic reconstrual. There is nevertheless a serious difficulty close by. For Quine, presumably, judgments about recalcitrance are in the same boat as the rest: even after the proper description of a particular sequence of experience has been granted, its recalcitrance for a particular theory-with-a-logic will have to be regarded as a *hypothesis*. And if we now ask: under what circumstances is it reasonable, on the Quinean view, to hold such a hypothesis?, it is clear what answer Quine officially must give. He must affirm that among the responses available to us in the original schematic situation is indeed the option of denying that E is recalcitrant for θ -with- L , by way of rejection of W ; and that, as in the case of other available responses, pragmatic considerations should be allowed to determine whether this is a good move. But determine it how? The decisive consideration ought to be, presumably, the degree of further recalcitrance with which the various alternative courses tend to be beset. But once the recalcitrance of experience becomes, in the way noted, a hypothetical matter, the question is transformed into: how often are the various alternative courses beset by sequences of experience which, *according to the best hypothesis*, are recalcitrant? And now, in order to decide whether recalcitrance *is* the best in hypothesis, we have to consider how it tends to fare in pragmatic competition with the alternatives – and the beckoning regress is evident. So the official Quinean answer to the question, when is it reasonable to believe a statement like W , is no answer. The moral is that it cannot be a correct account of the basis of our confidence in statements like W that belief-systems in which they figure enjoy relative success; if that were the right account of the matter, there could be no explaining the requisite notion of 'success'.

In summary, the Quinean methodology is, crudely: where experience is

recalcitrant, make whatever adjustments suffice both to eliminate the recalcitrance and to minimise subsequent recalcitrance. But if the latter part of the injunction amounts to: make whatever adjustments suffice to minimise the occurrence of situations whose recalcitrance is the *best hypothesis*, it is hopelessly impredicative – nothing has been said concerning by what methodology such a ‘hypothesis’ should be judged ‘best’. So the reasonableness, or otherwise, of judgments of recalcitrance must be exempted from appraisal via the Quinean methodology. And that must go for the ingredients in such judgments, including statements like *W*. The Quinean, as noted, might well have been prepared to regard such statements as candidates for analyticity anyway – in his laundered sense of the term. But now it transpires that that is not enough. If we are supremely certain of the truth of at least some such statements, the source of this certainty simply cannot be accounted for by Quine’s generalised holistic model. The very coherence of the model requires an account of a different sort.

The right account is, I believe, the obvious one: such statements, or at least an important sub-class of them, admit of totally convincing *proof*. We must, I suggest, take seriously the idea of proof as a theoretically uncontaminated source of rational belief. One reason why it is easy to overlook the incoherence in Quine’s attempt to ‘Duhemise’ the traditional realm of the *a priori* is because one naturally thinks of proof as conferring no more than a conditional warrant upon its conclusion, from premises for which, for this reason, the ultimate ground cannot itself be proof. But that is just an oversight. If we derive *B* from *A* using classical propositional logic, then *B* may be said to be proved conditionally on our acceptance of *A*; and $A \vdash B$ may be said to have been proved conditionally on our acceptance of classical propositional logic; but $A \vdash_{\text{CPL}} B$ has been proved conditionally on nothing at all. A sequence of operations of the relevant sort, taken as a proof that a certain logic does indeed have the materials to yield a specific conclusion from specified premises, can possess complete phenomenological cogency. No coherent methodology of empirical science can avoid recognising that such judgments, proved unconditionally in this way, play an ineliminable part in our conception of what it is for experience to collide with a body of theory. What, it seems to me, is fundamentally unsatisfying about the philosophy of logic of the global pragmatism in ‘Two Dogmas’ is that it is forced to locate the rationality, or otherwise, of our acceptance of such a judgment quite elsewhere than in the cogency of the operations which constitute its proof – and indeed, if what is said above is correct, winds up giving it no proper location whatever.

III FACTUALITY

6. I wish to carry forward two things from the preceding. First, I shall take it that we do possess some sort of concept of logical necessity (analyticity). Second, the correct account of the basis for the majority of judgments of logical necessity which we are prepared to make must make reference to the utterly convincing, self-contained character of suitable proofs. But the crucial point is that this much luggage is by no means a commitment to the objectivity syndrome adumbrated in section 1. If Quine’s doubt about analytic statements may usefully be compared to the contemporary doubt about the existence of witches, the line of thought to be developed now will suggest a doubt about the reality of logical necessity akin to the doubt which Locke had whether anything is *really* red, from an objective point of view, or the doubt which most of us have whether anything is, in the same sense, *really* funny.

It is evident enough that there are uses of declarative sentences which are not aimed at fact-stating. Promises, rules (‘The King moves one space in any unobstructed direction’), and commands (‘The platoon will be on parade at 6:30 am tomorrow morning’) are obvious examples. But a large and important class of philosophical disputes pivot precisely on whether the declarative sentences in a contentious family are apt for fact-stating, whether there is any genuinely factual, or cognisable, subject matter for them to state. Such disputes arise for a variety of reasons. The Logical Positivists’ conception of literal significance more or less forced them to deny the factuality of anything but reports of immediate observation. Similarly, the Dialectical Materialism of hard-line Marxists pushes them towards the view that all facts about human society and consciousness are ultimately constituted in the economic sub-structure. A third and perhaps more appealing motive is the thought that any genuine fact ought to be available to the cognitive powers of an appropriately endowed being; and that pure cognition cannot ever require the exercise of anything but intellect and reliable sensory faculties. It would follow that sentences involving terms like ‘funny’, ‘boring’, or ‘obscene’, competence with which requires a subject’s capacity for certain sorts of *affective* response, cannot be (purely) fact-stating.

Whatever their motivation, there are unsettled philosophical issues of this general character in ethics, aesthetics, philosophy of science, and the philosophy of mathematics whose resolution would constitute a tremendous advance in our philosophical understanding. Yet what is really in dispute? How can we characterise the needed distinction between *genuine statements*, declarative sentences apt to have truth or falsity conferred upon them by the properties of a real subject matter, and hence suitable for the expression

of genuine knowledge, and the rest:⁴ those sentences which have all the syntax – the susceptibility to embedding in conditionals and expressions of propositional attitude, etc. – of genuine statements yet which do not play a fact-stating role?

A sound and simple thought is this. If we are concerned with a genuine statement, apt to be rendered true or false by germane aspects of the world, then there must presumably be sense in the comparison between what any particular subject, or group of subjects, takes the truth value of the statement to be and what its truth value actually is. The point can be made vivid by a dilemma. Suppose we are concerned with a type of statement which (in principle) we can come to know. Well, cognition is a relational business: it involves getting one's beliefs, in appropriate ways, into line with the way matters stand. But there is no sense in the idea of securing such a coincidence unless there is a distinction between describing how things stand on one end of it, in the realm of the subjects' opinions, and describing how they stand at the other, in the realm of fact. Suppose on the other hand we are concerned with a type of statement which we cannot (in principle) come to know. Even here we can, presumably, guess at the truth value of the statement, and possibly get it right. So the relevant idea of coincidence must at least make sense. Accordingly we may affirm quite generally that only if:

- (i) it is accepted by such and such a person/group of persons that *P*;
- and (ii) it is the case that *P*;

enjoy an appropriately contrasting content, is it in order to regard *P* as a genuine statement. 'Appropriately contrasting' leaves lots of scope; but, standardly, we should expect space to be made for the possibility of being in position to assert either (i) or (ii) while being in position to deny, or at least being in no position to assert, the other. (It is the first two possibilities, of course, that provide room, respectively, for a subject's *error* about and for his *ignorance* of the facts.)

These ideas ought to seem uncontentious. But they are not toothless. Wittgenstein's leading idea, as I read him, in *Philosophical Investigations*, §§256–61 is exactly that the requisite contrasting content cannot be made out if the subject in question is to be the 'private linguist' and *P* is to range over the 'statements' of his private language.⁵ Moreover the hesitation which it is natural to feel about the factuality of our judgments concerning what is funny, or what is obscene, surely has something to do with our diffidence that we really understand what it would be to be entitled to regard a majority – or even a large – group as strictly mistaken, or ignorant, in their

⁴ What Dummett calls 'quasi-assertions' (1981: 353–63).

⁵ I pursue this idea in my (forthcoming).

opinion on such a matter. Yet these considerations point, evidently, only to a necessary condition for a class of declarative sentences to count as genuine statements. At any rate, the (i)–(ii) contrast is *in use* in many of the areas – theoretical science, pure mathematics, ethics, aesthetics – where factuality is in dispute. Likewise with the topic of our present concern: if *P* depicts *B* as a consequence of $\{A_1 \dots A_n\}$, or avers that such-and-such a structure is a valid proof in such-and-such a system, nothing seems easier than to understand what it would be to be in a position to make any of the four contrasts adverted to. Rather than find so easily in favour of the realist/cognitivist view we should seek a strengthening of the proposed account. Not only the possibility but, in a sense we have to explain, the *propriety* of practising a distinction between judgment and fact is what is at issue.

One useful suggestion, I think, originates in the idea, touched on above, that knowledge is *dispassionate*. Another is the thought that truth is *coercive*: when a statement expresses a matter of fact, assent to it may, in certain circumstances, be commanded of us. Of course, finding something funny may also be, in context, an irresistible response. The intuitive difference – putting the two suggestions together – is that when a statement of fact commands our assent, it does so independently of any emotional or affective response which we have to the matter. If ever a genuinely factual statement is beyond dispute, it is, properly speaking, for the rational, whose intellects, senses, and memories are functioning properly, that it is beyond dispute.

A bold proposal would now take, as the hallmark of the factual, the appropriateness of an ideal of rational consensus. The idea would be that we should think of a class of statements as expressive of genuine matters of fact only where it can be shown that, if perfectly rational beings were permitted to conduct a sufficiently thoroughgoing investigation, the opinions which they formed about the acceptability, or otherwise, of such statements could not but coincide. Genuine truths, on this view, are what perfectly rational beings would agree to be true on the basis of a sufficiently lengthy and painstaking investigation. Failing better motivation, however, the proposal seems over-contentious: again, too easy a resolution is promised of too many of the controversies in which factuality is pivotal – only this time the verdict goes against the realist/cognitivist. Theoretical science for example, could not qualify as factual under the proposal unless the falsity could be demonstrated of the notorious thesis of the underdetermination of theory by empirical data; and it is, at best, highly controversial whether all ethical, or aesthetic, disagreements may in principle be resolved by rational means alone. More generally, the proposal assumes that there is no such thing as our forming a clear conception of a possible state of affairs for supposing which to obtain, or to fail to obtain, even a perfectly rational,

indefinitely extended enquiry might be able to disclose no reason. That is to presuppose the falsity of realism in Dummett's sense. I myself do not regard such a presupposition as an error of substance. But it is certainly a political error: if the right way to draw the distinction between genuine statements and the rest is to eventuate in an objection against realism, the way the distinction is initially drawn should carry an appeal for the realist and anti-realist alike.

The most plausible way, it seems to me, of preserving what seems right about the bold proposal – the idea that fact-stating has somehow to do with an appeal to our rational faculties – while avoiding immediate question-begging against Dummettian realism, is to take the 'tug' on the rational faculties to be exerted not by an idealised enquiry but by the state of information in which we happen to be. A genuine statement will be: not a statement about whose truth value ideally rational investigators could not disagree after a sufficiently thorough-going investigation, but rather a statement about whose assertibility, or otherwise, ideally rational subjects will not disagree *in any particular state of information*. Anti-realistically problematic examples can now qualify: if Goldbach's Conjecture, for example, is utterly 'verification transcendent', then, no matter what their state of information, perfectly rational subjects will presumably agree that there is no basis either for its assertion or for its denial. In contrast, such subjects may disagree – our intuitive feeling is – about whether, say, a particular interview of a leading politician was unintentionally comic without there having to be any suggestion that one enjoys a superior state of information to the other or, if their states of information are the same, that one is being less than ideally rational.

Genuine statements, according to the milder proposal, distinctively command a particular response from the rational, *modulo* a state of information. Contraposing, differences of opinion about such statements – that is, one subject holding an opinion which another does not – will have to be traceable back to some breach of ideal rationality or material difference in the subjects' respective states of information. This proposal should, I believe, commend itself to the reader as intuitively correct; but only in a formal sense, since no condition has so far been imposed on what can qualify as an item of *information*, meet to enter into a state thereof. If we are aiming at the adjudication of controversial cases, the proposal is so far, therefore, entirely powerless. The way to give it some cutting edge is to retain the suggestion that what is distinctive of genuine statements must be sought in the range of possible sources of differences of opinion about them, but to attempt to be more specific about the members of that range. The following more detailed proposal is such an attempt.

Where 'facts of the matter' are concerned, differences of opinion can be rendered fully intelligible to a third party, I suggest, only if he can

- (a) identify a material mistake on the side of one of the parties; or
- (b) identify some material ignorance on the side of one of the parties; or
- (c) identify some material prejudice on the side of one of the parties; or
- (d) disclose some material vagueness in the statement used to express the opinion in question.

Mistake is here to be taken to cover any sort of perceptual, recollective, or intellectual malfunction. More needs to be said, obviously, about when it is satisfactory to explain a difference of opinion by placing it in this category, but the following remarks will serve our immediate purpose. Presumably it should count as satisfactory if the mistake is identifiable independently of any view about the disputed opinion, as in the case of an error made in the course of a calculation, for example, or a misreading of a gauge; or, failing such an identification, if aspects of the condition of the subject, or of the circumstances of his judgment, are known about which, against a background of information, say, concerning the physiology of germane modes of the subject's functioning, would make a mistake of one of the relevant sorts likely. But attribution of mistake will *not* count as a satisfactory explanation so long as the sole ground for the attribution is the subject's view of the disputed statement. Similar points apply to *ignorance*. It will be satisfactory so to explain a difference of opinion whenever some material ignorance is identifiable independently of one's view of the disputed statement; or, failing such identification, when there is at least an explanation of why it was likely, or even inevitable, that the subject should be left in ignorance of the status of that statement. But it will not be satisfactory if the sole ground for attributing ignorance to the subject is that he does not hold a certain view of the disputed statement.

Prejudice would be the appropriate form of explanation in situations where the protagonists agree about the material data but disagree about its supportive strength. However this needs a qualification. Intuitively no compromise of rationality is involved if, for example, *X* requires that the probability of getting a certain favourable outcome from a change in some policy be at least 0.75 before he is prepared to implement the change, whereas *Y* is satisfied with a probability of 0.70. The question, whether a certain degree of probabilistic support makes it rational to hold a particular belief, does not everywhere admit of a determinate answer. That, however, cannot *always* be the situation so long as the belief does admit in principle of probabilistic confirmation and disconfirmation. If something factual is at issue, it must be possible for the evidence to assume such a shape that

only irrationality can explain the refusal, or willingness, of someone who knows of the evidence to accept the belief in question. Prejudice, as the notion is here intended, is what is operative when someone assigns an irrationally high or low supportive force to an agreed body of evidence. Rationally permissible differences in personal probability thresholds, as it were, cannot be the explanation in all possible cases of differences of opinion, concerning genuine statements, for which the only evidence is probabilistic and is agreed on both sides.

The relevant point about *vagueness*, as a quite general possible source of differences in opinion about genuine statements, is similar. To wit, for the statements in question to qualify as factual, it is necessary that not *every* possible difference of opinion about their status, in given circumstances, can be put down to vagueness. Descriptions of colour, for example, if *pace* Locke we regard them as factual at all, so qualify only because, though the borderline between various colours is blurred, some things are, for example, determinately pink and others determinately not.

The proposal, then, comes to this. Statements of a certain class are apt for the expression of genuine matters of fact only if there are contexts – in which vagueness, or permissible differences in personal evidence thresholds, are not to the point – in which it is *a priori* that differences of opinion concerning one of the relevant statements can be fully explained only by disclosing (in a manner which observes the constraints sketched above) some material ignorance, error, or prejudice on the part of some or all of the protagonists. And a particular debate concerns a genuine matter of fact only if the statement(s) which express what is in debate satisfy this condition in the context in question. By contrast, members of a class of apparent statements will not count as apt for fact-stating if, whenever such a statement is in dispute, it can never be ruled out *a priori* that an explanation of the dispute should be of some other kind than those just described. Disputes about ‘matters of taste’, for example, *may* be traceable to ignorance, error, or prejudice but their intelligibility never *requires* that they be so: *X* and *Y* may, for example, just find different things agreeable in matters of interior decoration and design – that may just be the whole of the matter. On the other hand the proposal seems to do justice to our preconceptions concerning the fact-stating character of, for example, statements concerning yesterday’s weather, the whereabouts of the cat, or the number of times I have been to Holland. It is, for example, *prima facie* impossible to understand how you and I could have a difference of opinion about the last whose explanation – saving some material misunderstanding like the belief that ‘Holland’ covers all the Low Countries, or our counting ‘trips’ differently – would not make good such a claim as that I have forgotten, or dreamed up, a trip; or that you are in no position to know; or that you are for some reason

pleased to think of me as someone with no experience of travel abroad... and so on.

I commend this proposal only as an outline, in essentials, of (part of) the inchoate notion of factuality which we actually have. And be it noted that it is the truth of the proposal, rather than its status as a (complete and non-circular) analysis, which matters in what follows. If it incorporates at least a necessary condition for a class of statements to qualify as genuine, then it will be reasonable to demand of a factualist about ethics, or aesthetics, or indeed necessity, that difference of opinion about such matters – if genuine, that is, based on no misunderstanding – must, at least in the relevant sort of contexts, be explicable in one of the three ways described. Let us say that a dispute is *Humean*⁶ provided there is no material misunderstanding of any concept involved in the formulation of the object statement, and the source of the dispute is not error, nor ignorance, nor prejudice. The anti-factualist about a given class of statements will hold that, whenever there is a difference of opinion about such a statement, it will always be a possibility, *a priori*, that it is Humean, that is, may be successfully explained without being placed in any of these three categories. The factualist, in contrast, must hold that, unless vagueness, or rationally permissible differences in evidence thresholds, are to be the explanation, it has to be the case that it is in one of those three areas that the explanation lies.

7. This proposal, too, may be in tension with Dummettian realism. Only now, I think, that will constitute an argument, rather than a question begged. Can any genuine statement – by the lights of the proposal – be utterly verification transcendent, so that not even the weakest ground for believing or doubting it can be given? Only if a difference of opinion about such a statement need not be Humean. But that requires that there are circumstances in which it is *a priori* that such a difference of opinion has to be put down to ignorance, error, or prejudice. Now, there might be no difficulty were it that agnosticism was always the rational attitude to take up towards such a statement. For then, since any difference of opinion is always going to involve at least one of the parties not being agnostic, the explanation would presumably always be either a mistaken belief that certain data were available or a prejudicial indifference to the total absence

⁶ At the Lyme Regis Thyssen conference in 1983 at which the original version of this paper was presented, I called such disputes ‘Homeric’. The epithet was meant to convey an allusion to Sir Peter Strawson’s inaugural lecture in which he spoke of the ‘Homeric struggle’ between truth conditional and communication-intention approaches to meaning, urging that the two approaches are in no genuine dispute. But (i) nobody got the allusion; and (ii) it was not in any case quite felicitous, since disputes with no factual subject matter can be real enough, and should therefore be contrasted with cases where the protagonists merely believe they are in dispute.

of data. But agnosticism is not always the favoured realist response to such statements. A good example is provided by Edward Craig's handling of his 'assumption of uniformity' in his contribution to the present volume. The assumption is, as Craig conceives it, one which human beings make quite naturally; it is that

Other members of their species experience inner (that is to say, epistemically private) states which are pretty much like the ones they experience themselves, and that they experience them when the outward circumstances are broadly speaking similar. (p. 175)

It is essential to Craig's argument that this is a genuinely factual assumption, which may be 'frequently correct in particular cases' (p. 176). And Craig evidently thinks that, so far from it being unreasonable to make such an assumption, it is a natural and proper thing to do. Yet to hold that we are indefeasibly within our rights in making such an assumption, in the absence of any possible evidence, and, at the same time, to concede that there is no guarantee that an agnostic need be guilty of any identifiable error, or ignorance, still less a prejudicial response to data, is to bring it about that, by the lights of the proposal, a statement of the assumption fails to qualify as factual. For in that case neither the Craigian nor the agnostic can be brought under any of the three pertinent headings. (The result, I suggest, is that the idea of *consensus* with respect to our description of inner states cannot be elucidated, in the manner Craig suggests, by reference to the possible truth of such an assumption. But I have no space here to engage Craig's argument in further detail.)

There are actually rather a lot of 'assumptions' which seem to be in this category: hypotheses of apparent depth which we are inclined to regard as possessing determinate truth conditions but for which, under sceptical pressure, we seem to be unable to find even the weakest support to corroborate the attraction which they have for us. Familiar examples would be 'there are other consciousnesses besides my own', 'the earth is many millions of years old', 'material objects exist when unperceived', and many more. In each case, of course, we can offer what at first sight appears to be plentiful and powerful supporting evidence; but it is, famously, easy for the sceptic to argue that such evidence is variously question-begging. One issue now is whether it can be coherent to grant so much to the sceptic while retaining the beliefs both that there *are* substantial matters of fact at issue and that it is somehow not improper to retain one's convictions in the teeth of the sceptic's arguments. The threat presented to the Dummettian realist is exactly that the price of saving the factuality of the subject matter in sceptical disputes may turn out to be that the sceptic *wins*: that all our talk about other minds, the past, and the material world turns out to be based

on groundless assumptions which, rationally, deserve agnosticism. Of course the belief that such 'assumptions' are misinterpreted if taken to be factual is a leading theme of Wittgenstein's notes (1969). Wittgenstein has a variety of motivations for the idea, none of them exactly coincident, it seems to me, with the train of thought adumbrated. (I defer fuller discussion of the matter to another occasion (cf. Wright 1985).)

IV HUMEAN LOGICAL NECESSITY

8. No doubt the proposal could do with more refined formulation. But it is not, as it stands, too crude to display a problem with the factualist conception of necessity which Quine rejected and to suggest merit in the attitude of those who would have described themselves as 'conventionalists' when the (unresolved) dispute about the nature of logical necessity was in its heyday during the 1930s and 1940s. The problem is simply that it is unclear, on reflection, why it is not *always* possible to have a Humean difference of opinion about the necessity of a statement generally accepted as necessary; in particular, why someone may not always Humeanly step short of accepting the necessity of such a statement while allowing its truth.

To elaborate briefly by reference to the case of formal proof in logic. Suppose that what is at issue is the *technical* correctness of such a proof – so the sort of disputes which classicists, intuitionists, and relevantists', for example, might want to have about its soundness are not to the point. The issue is to be merely whether what we are presented with is a correct proof in a particular formal system. Now, the concept of logical necessity enters even here, it will be recalled, in so far as the status of such a construction as a proof – rather than, for example, an experiment – depends upon its *essential stability*: it must not be 'logically possible' that the outcome of the proof should vary through successive performances in the way that the outcome of a physical experiment can. Accordingly there will be some description of the successive operations of the proof such that, although the result at each stage is not explicitly given by this description, it is nevertheless necessary, or so we ordinarily think, that if precisely the sequence of described operations is carried out on the starting point of the proof, nothing but the eventual outcome *can* result. To take a trivial example, consider

- (1) $A \rightarrow B \vdash A \rightarrow B$
- (2) $A \vdash A$
- (3) $A \rightarrow B, A \vdash B$ 1,2 MPP
- (4) $A \rightarrow B, A \vdash B \vee C$ 3, vel-I

An appropriate description, given in the form of a conditional, could be something like:

If any proof commences with a pair of assumption-sequents, $A \vdash A$, and $A \rightarrow B \vdash A \rightarrow B$, followed by the *modus ponens* step which those two lines furnish, followed in turn by a step of vel-Introduction on the result with C as the right-hand constituent in the then resulting disjunction, then that disjunction will be $B \vee C$, and will depend on A and $A \rightarrow B$ as assumptions.

Imagine now a dispute between X who, viewing the structure, accepts it as a proof, and so accepts the necessity of such a conditional as the above, and Y who merely regards the structure as experimentally corroborating the conditional, whose truth he regards as enormously probable. If there is to be a genuinely factual issue whether that structure is a proof, such a dispute cannot always be Humean: there have to be circumstances in which we can say in advance that the only way of rendering the dispute intelligible is if one or each of the parties has committed some specific error, or is ignorant in some material way, or is guilty of prejudice. Moreover, since there is no germane vagueness in the statement at issue and since what is at issue between X and Y is clearly not to be put down to permissibly differing evidence thresholds, we may take it that the envisaged dispute between X and Y takes place in just such circumstances. Given, then, that the factualist will want to regard X as right, how should Y be handled?

A natural thought is that Y 's very response to the proof – his treating it as, in effect, a sort of experiment – betrays a misunderstanding of the notion of necessity: that if Y genuinely understands that notion, he *must* see that the necessity of an appropriate conditional is apt to be demonstrated by such a structure. Accordingly, there is no need to view the case as a genuine difference of opinion at all. But there are two obvious drawbacks to this. First, a parallel thought is likely to be available *whenever* it looks as though Humean disputes are always going to be possible in some area, and cherished factualist preconceptions are consequently under threat: were it to appear, for example, that Humean disputes about fundamental moral precepts are always possible, the response will always be available that the agnostic can have no proper grasp of moral notions. So the charge of misunderstanding can be admissible only subject to certain controls; otherwise it becomes a two-edged sword, leaving no clear distinction between responsible forms of factualism and mere crankishness. Second, we can imagine Y elaborating his position in a way which makes it evident enough that he understands pretty well what is at stake. He may grant, for example, that he cannot imagine what it would be like for a structure to seem to him to meet the specification of the antecedent of the conditional and yet have a different outcome. He may grant that this marks an

interesting and important contrast with other experiments, where a detailed description of counter-factual outcomes, or even a cine film simulation of them, might be possible. But he sees, he insists, no cause to project aspects of our imaginative powers onto reality, or to dignify them as apprehension of what *must*, or *cannot*, be the case. After all, it surely *is* imaginable that we might somehow, sometime want to describe some structure as indeed a counter-example to the relevant conditional; and that we shall then find it extraordinary that we could have been so blind before....⁷

Given that he explains himself along these lines, the charge that Y misunderstands the relevant notion of necessity looks far-fetched. Indeed there seems to be no cause to reproach his understanding of any relevant concept, since he does, after all, show himself aware of the character of the ingredient steps in the construction and of what statement it confirms. The dispute hinges, rather, on the proper interpretation of the nature of this confirmation.

Granted, then, that the difference of opinion is genuine, it must not, for the factualist, be Humean. So Y must be guilty of error, ignorance, or prejudice. But the last of these possibilities is not to the point, for it requires that the data be agreed; and if X were to admit that the data are just as Y describes – that they are constituted by the empirical features of the construction, plus the considerations about imaginability which Y acknowledged – then he would be obliged to answer Y 's question: what reason do we have to believe that such data are indicative of a genuinely objective *genre* of 'necessities' and 'possibilities'? And it is unclear what response X has to that. Yet making Y out to be in error, or ignorant of relevant data, looks to be a no more promising prospect. Y 's perception of the construction is presumably in order, and he need be guilty of no technical error in working over it. Likewise we may suppose it impossible fairly to interpret him as being ignorant of any relevant consideration, unless it be the necessity of the conditional, and so the status of the construction as a proof, itself. So X 's case demands *either* that the facts acknowledged by Y actually constitute his – Y 's – recognition of the proof, without his realising it; *or* that Y is here 'proof blind', as it were – that more is indeed involved than Y acknowledges, the extra being precisely intellection of a logical necessity. But, either way, the difficulties are evident. The former line demands an answer to Y 's doubts about the warrant for so 'dignifying' features of our imagination. And the latter again faces the reproach that it exemplifies a manoeuvre which is *always* available to factualism, resort to which, without proper controls, will merely erode the very content of the issue.

⁷ Cf. Wittgenstein (1964: 3d edn, III §87).

Above, I imposed the condition that it cannot count as satisfactory to attribute a subject's view of a particular statement to a mistake, or ignorance, if the sole ground for doing so is that he takes the view he does. It should now be clearer why this constraint is needed. Unless there is more to be said, we are powerless to defend the distinction between the operations of any genuinely cognitive faculty, affording non-inferentially based knowledge, and the working of something which, like the 'sense' of humour, we do not wish to regard as genuinely cognitive at all. It may be, in certain cases, that there is no *identifying* a subject's mistake/ignorance except by describing his view of the disputed statement, but the *ground* for so describing him cannot stop there. What the dispute between *X* and *Y* brings out is that there is a disturbing parallel (at least it ought to disturb the factualist) between judgments of logical necessity and judgment about what is amusing. In both cases, disputants may be in agreement about all features of a situation *except* whether it establishes a logical necessity or is amusing; and all the cards may be on the table – no further consideration need be available which, once apprised of it, would bring the disputants into agreement.

The parallel is not, of course, decisive: there is still space for the claim that the status of the construction as a proof, for example, is a further feature of it, over and above its empirical features, sensitivity to which calls for the operation of a special faculty. But then there has also to be space for the corresponding, highly unattractive claim about amusingness. One thing, accordingly, is clear: the postulation of a special intellectual faculty, sensitive to logical necessity, cannot be justified merely by our propensity – if we have it – non-collusively to agree in our ratification of new proofs, or, more generally, in our judgments about necessity in novel cases. What counts is not the propensity towards non-collusive agreement, however widespread, but whether Humean disputes about necessity are always possible. A universally shared sense of humour would not make issues concerning what is funny any more factual – assuming that they would, on proper analysis, prove not to be so. It is a coherent and competitive view to hold both that the notion of logical necessity is genuine and that the anti-factualist spirit of 'Two Dogmas of Empiricism' should be endorsed.

Why does the suggestion that the sense of humour is a cognitive faculty seem so outlandish? At bottom, it is because our conception of ourselves and of our knowledge-acquiring powers is broadly *naturalistic*. We are content to regard something as a cognitive faculty only so long as there is the promise of a proper explanation of its physiological basis, an account which enables us to see its output as the product of physical interaction between our bodies and the environment. There simply is no such promise in the case of amusingness; we do not have the slightest idea what the amusingness of a

situation could physically consist in (contrast redness and the emission of light waves of such-and-such frequency) on the basis of which an account might be built to match the joint achievement of physiologists of the visual system and physicists of light. It is the background supplied by such an account which enables us, ultimately, to substantiate attributions of error, or ignorance, in the manner called for by the earlier constraints.

It is a great question whether naturalism affords the materials for an overall coherent epistemology. But the factualist about logical necessity had better believe that it does not. For everything said in the previous paragraph about amusingness applies to necessity also.

9. Readers familiar with another discussion of mine of this issue (1980: chapter 23) will recall the prominence there given, in the attempt to arrive at a general description of the domain of the factual, to the role of the seemingly ubiquitous possibility of sceptical doubt. Of course the sceptic as an actual human agent – one who resolutely seeks to tailor his corpus of beliefs to those consistent with sceptical standards of justification – is presumably a fiction. But the prototypical sceptical *routine* – the play with the inconclusiveness of available data, the fallibility of our capacities, etc. – looks to be a possible manoeuvre everywhere. This reflection need not be a reason for modifying the account proposed above: a difference of opinion generated by sceptical doubt might prove to be best described as involving 'prejudicial assessment of data', for example.⁸ The important point for the anti-factualist about necessity is rather that *Y*'s position in the above dispute must not be best described, in essential respects, merely as a form of scepticism. If it were, the factualist could evade the need to give proper substance to a preferred description of *Y* as in error, or ignorant, presenting him instead as a familiar animal, occupying a stance which is available everywhere and has no bearing on questions of factuality.

The matter needs a more detailed discussion than I have space remaining to attempt here.⁹ Still, one initial consideration suggests that *Y*'s position is not happily described as sceptical. Traditional forms of scepticism – about other minds, or the past, or generalisations inductively arrived at, for instance – never dispute that there is such a thing as getting the truth values of the relevant class of statements right. Their essential claim is rather that, for all the controls which we have at our disposal, success will be a *fluke*: no ground is, or can be, possessed for the reliability of those controls. It is granted that there is an objective subject matter at issue: the challenge of

⁸ This description would require, to stress, that *Y*, as sceptic, concurs with *X* about the data germane to *X*'s judgment that the construction demonstrates necessity. If that is so, the data would have to be empirical features of the construction plus the considerations about imaginability which *Y* acknowledged – so, it might plausibly be urged, *Y*'s scepticism would be well conceived. Cf. Wright (1980: 464).

⁹ The issue is taken up by Edward Craig in his (1984).

the traditional sceptic is that we make good our belief that our epistemic capacities are up to the task of securing reliable beliefs about it. It is clear enough that this is not what *Y* was depicted as saying. *Y* did not grant *X* that *some* descriptive conditional is necessary, and then seek to cast doubt on the effectiveness of our controls – principally, careful attention to the promptings of intuition and meticulous checking of constructions – on attempts to winkle out the right one. Rather, he sought to be persuaded that, in order to do justice to the construction, there was any cause to invoke a special notion of ‘necessary’ truth. And it is, so far, quite unclear whether, or how, he ought to be persuaded of that. The issue is not the presumed reliability of certain capacities of ours, acknowledged as genuinely epistemic on both sides, but rather whether we are concerned with a genuine epistemic capacity at all.

Two concluding remarks. First, I have recently quite often encountered in conversation the impression that the ‘rule following considerations’ somehow dispose of this class of questions, teaching us that all our judgments – in ethics, aesthetics, pure mathematics, empirical science, and any other field of human expression you care to consider – are ineliminably conditioned by basic human reactive propensities; that if our nature has a part to play in judgments about what is funny, for example, it has a comparable part to play everywhere. It would be a major task to unpick all the knots of confusion in this notion. But one thing I hope to have made plausible is that a framework remains for discussion of issues to do with factuality which may be utilised even after the global lessons of Wittgenstein’s ideas about following a rule have been fully digested. Second, although I believe that the framework described provides a context for many of our raw intuitions on these questions, I cannot pretend to certainty that further work will not disclose that there is no real substance to them. *Quietism*¹⁰ may yet win the day: it may prove impossible to give clear content of the distinction between genuinely fact-stating and non-fact-stating declarative discourse, and Wittgenstein’s stress in the *Philosophical Investigations* on the essential multiplicity of language games may prove to be profoundly insightful in just this respect (1953: section 28).

Actually, such a result would itself motivate a kind of anti-factualism, though one defined by its rejection of the basis of the factalist position rather than by its opposition to it. At any rate, the issues are wide open; and if philosophy can legitimately aspire to yield illuminating comparisons between our modes of thought and speech and the way things really are, they are important issues.

¹⁰ The term is Simon Blackburn’s (1984: 146).

REFERENCES

- Blackburn, S., 1984. *Spreading the Word*, Oxford: Clarendon Press.
- Craig, E. J., 1984. Arithmetic and Fact, in *Exercises in Analysis*, ed. I. Hacking, Cambridge: University Press.
- Craig, E. J., 1986. Privacy and Rule-Following, this volume.
- Dummett, M., 1959. Wittgenstein’s Philosophy of Mathematics, *Philosophical Review* 68: 324–48.
- Dummett, M., 1981. *Frege: Philosophy of Language*, 2nd edn, London: Duckworth.
- Grice, H. and Strawson, P., 1956. In Defence of a Dogma, *Philosophical Review* 65: 141–58.
- Quine, W., 1961. Two Dogmas of Empiricism, reprinted in *From a Logical Point of View*, 2nd edn, New York: Harper and Row.
- Wittgenstein, L., 1953. *Philosophical Investigations*, ed. G. E. M. Anscombe & R. Rhees, Oxford: Blackwell.
- Wittgenstein, L., 1964. *Remarks on the Foundations of Mathematics*, 3rd edn, ed. G. H. von Wright, R. Rhees and G. E. M. Anscombe, Oxford: Blackwell.
- Wittgenstein, L., 1969. *On Certainty*, ed. G. E. M. Anscombe and G. H. von Wright, Oxford: Blackwell.
- Wright, C., 1980. *Wittgenstein on the Foundations of Mathematics*, London: Duckworth; Cambridge Mass.: Harvard University Press.
- Wright, C., 1985. Facts, Scepticism and Certainty, *Proceedings of the British Academy*, December 1985.
- Wright, C., (forthcoming.) Does Wittgenstein have a cogent argument against Private Language? Investigations §§258–61, in *Subject, Thought and Context*, ed. J. McDowell and P. Pettit, Oxford: Clarendon Press.