

This has a local peak which can be levelled, transforming it into the canonical argument of degree 1:

$$\frac{A \quad B}{A \& B}$$

This is valid, because it contains no critical subarguments. We have thus provided a justification of the original argument, a proof-theoretic justification of the *third* (and highest) *grade*.

This gives us a comprehensive criterion for the validity of an arbitrary argument, relative to our introduction rules. In fact, our criterion takes no overt account of more than the initial premisses and final conclusion of the argument: what comes between may help to supply us with the effective means of transformation that we need to show the argument valid, although these intervening lines play no explicit role in our criterion. We may therefore convert our criterion for the validity of a rule into one for the validity of a *sequent*; the sequent $A_1, \dots, A_k : B$ will be valid if the one-step argument

$$\frac{A_1 \quad A_2 \quad \dots \quad A_k}{B}$$

is valid. A *rule* of the simplest kind, that involves no discharge of hypotheses (or introduction of new assumptions), may be represented by a schematic sequent, one whose antecedent is composed of formulas rather than sentences and whose consequent is a formula. As before, we may call such a rule valid if we have a means of showing any sequent resulting from the schematic sequent representing it by instantiation to be valid. A rule in general may be represented by a finite number of *base* schematic sequents, and a single *resultant* schematic sequent. (It would be confusing to call these the 'premisses' and 'conclusion' of the rule.) A rule so represented may be said to be valid if we have a means of showing, of any uniform instantiation of the schematic sequents, that if it renders the base sequents valid, it will render the resultant valid also.

The Fundamental Assumption

Is the Fundamental Assumption Plausible?

It is plain that proof-theoretic justification of the third grade is a powerful procedure. It has here been formulated so as to be applicable to any set of logical constants, governed by whatever introduction rules are chosen, provided only that they conform to the mild constraints we laid down. Given the usual introduction rules, it will certainly serve to justify all valid laws of first-order positive logic (the negation-free fragment of intuitionistic logic), a fact that can be verified by confirming that it validates all the standard elimination rules. It is recognisable as a *justification* procedure, however, only to the extent that the fundamental assumption is plausible: that must therefore be the next topic of our enquiry.

Evidently, the plausibility of the fundamental assumption is entirely relative to the logical constant in question and to the set of introduction rules being proposed as governing it. For instance, it would have no plausibility at all if applied to the modal operator ' \diamond ', regarded as subject to the sole introduction rule allowing an inference from A to ' $\diamond A$ '. If the fundamental assumption were taken to hold in this case, the converse inference could be validated, so that the operator ' \diamond ' would become quite nugatory; for, if a canonical derivation of ' $\diamond A$ ' must end by deriving it by means of the sole introduction rule, we must be able to give a canonical derivation of A whenever we can give one of ' $\diamond A$ '. We can therefore consider the fundamental assumption only on a case by case basis.

Disjunction

The problem is in part one of elucidating the 'could have' that occurs in the statement of the assumption. What does it mean to say that, if we are entitled to assert a statement of the form ' A or B ', we *could have* arrived at that position by applying one or other of the or-introduction rules? Plainly, this is untrue if applied to individual speakers.

I may be entitled to assert 'A or B' because I was reliably so informed by someone in a position to know, but if he did not choose to tell me which alternative held good, I could not apply an or-introduction rule to arrive at that conclusion. Of course, the fundamental assumption is not intended to be understood so as to make this a counter-example to it, or else the most doctrinaire intuitionist would be unable to endorse it. We must distinguish between individual possession of a piece of information and our collective possession of it: *my* source may be the testimony of another, but *our* original source must stem from whoever first established the statement as true, say by observation. Here whatever witnesses we trust must be included among 'ourselves'. If an angel reveals that either the citizens of Nineveh will repent or the city will be destroyed, whoever accepts the revelation must include the angel among the community with whose collective information we are concerned. The testimony may derive from someone long dead: we may believe, on the authority of a contemporary chronicler, that Constantine either murdered his son or procured his murder; if so, the chronicler must rate, for this purpose, as a member of the community, which will include the dead as well as the living. The fundamental assumption, as applied to disjunctive statements, must thus be to the effect that whichever member of the community originally established the truth of such a statement could have arrived at its truth by the rule of or-introduction. To make this plausible, we must not only broaden the boundaries of the community to comprise the dead and, in so far as we can communicate with them and trust them, the non-human, but distinguish our past from our present selves; for memory must be treated in analogy with testimony. If, having been an eyewitness, I report that either Gertrude or Diana tore up her invitation card, but confess that I cannot remember which of them it was, my memory is delivering partial information to me, just as another informant might do. My observation at the time was the original source, the information I have retained being only a weakened consequence of that which I originally possessed.

None of this, however, concedes enough. This is evident from the fact that, without further concessions, the plausibility of the fundamental assumption will depend heavily on what we take to be the primitive predicates of the language. Should we construe 'is a child' as a disjunction of 'is a boy' and 'is a girl', or should we construe 'is a boy' as equivalent by definition to 'is a male child', and similarly for 'is a girl'? Under the former alternative, but not the latter, my seeing that a child was playing on the lawn, without being able to tell whether it was a boy or a girl, would appear a counter-example to the fundamental assumption; but, obviously, the choice is spurious. If we are to

distinguish at all between defined and primitive expressions of natural language, the distinction must lie between those an understanding of which is characteristically mediated by knowing a verbal equivalent and those for which this is not so. On this criterion, 'child', 'boy', and 'girl' must all rank as primitive: they belong to a circle of expressions an understanding of any of which demands, but does not consist in, a knowledge of equivalences between each of them and expressions constructed from the others.

This requires an extension of our conception of 'boundary rules'. These were intended to take account of inferential connections between non-logical expressions and were restricted to inferences from atomic premisses to an atomic conclusion. Unless we are prepared to consider deductions as being carried out in a highly regimented version of natural language, in which the primitive predicates have been cut down to a minimum as in an axiomatised mathematical theory, we shall have to extend the notion of a boundary rule to allow the conclusion to be complex. When the conclusion is an open sentence, this will cause no difficulty, since the fundamental assumption will not be applied to it. When it is a closed sentence, however, we are left with an apparent counter-example to the fundamental assumption: if I know that there is a child playing on the lawn, I thereby know that either a boy or a girl is playing there, perhaps without knowing which, even though it is my own observation that constitutes the source of my knowledge. Likewise, if a boundary rule in the extended sense permits an inference from 'That is a child over there' to 'That is either a boy or a girl over there', the disjunctive conclusion was not arrived at by 'or'-introduction, and may well not have been able to be on the basis of the observation actually made.

Manifold other examples are independent of any linguistic question. Hardy may simply not have been able to hear whether Nelson said, 'Kismet, Hardy' or 'Kiss me, Hardy', though he heard him say one or the other: once we have the concept of disjunction, our perceptions themselves may assume an irremediably disjunctive form.

To interpret the fundamental assumption, then, we have to invoke the sense of 'could have' which was used earlier to characterise what may be called the minimal undeniable concession to realism demanded by the existence of deductive inference. The proof of the Königsberg bridge theorem provides an effective means so to carry out simultaneous observations to check whether the traveller crosses every bridge and to check whether he crosses any bridge more than once as to ensure that a positive result for the former will be accompanied by a positive result for the latter. We treat this as warranting us in asserting that some bridge was crossed at least twice, given that he was observed

to have crossed them all, even though we cannot now observe him to have done so or recall specific observations to that effect. This can be explained only in terms of a certain conception of the condition for an assertion to be correct: namely, that a sufficient condition for its correctness is that there exist effective means by which, at the relevant time, someone appropriately situated *could have* converted observations that were actually made into a verification of the statement asserted. The resulting notion of truth, possessed by any statement that meets this minimal condition for being capable of being correctly asserted, is very far from being that of full-fledged realism. Full-fledged realism does not merely regard a statement as true if it could as a matter of fact have been verified at the relevant time by an appropriately situated observer and hold it to be determinate whether or not he could have done so: it goes much further, maintaining that truth attaches to statements that we have not and could not have verified. Nevertheless, even this spare notion of truth, far too lean to satisfy the realist, debars us from making the coarsest identification of a statement's being true with its having been verified: if we understood the concept of truth in so coarse a manner as that, we should be unable to countenance deductive argument at all.

It is the 'could have' that occurs here to which the fundamental assumption appeals. If I pass from saying, 'A child ran out of the house', to saying 'Either a boy or a girl ran out of the house', it may be said that I *could have* arrived at the latter statement by or-introduction—not on the basis that I was in fact in a position to assert either of the two disjoined statements 'A boy ran out of the house' and 'A girl ran out of the house', but on the strength of the fact that, given that I was entitled to make the statement I did make, I had an effective means available to me for putting myself in a position to make one or other of the disjoined statements. This explanation will force us to adopt a laxer criterion for the validity of an argument. The criterion given earlier was that we had a means to transform any supplementation of an instance of the argument into a valid canonical argument *with the same initial premisses and final conclusion*. If, however, there occurs in the given argument an application of a boundary rule of the extended kind, leading, say, from atomic premisses to a disjunction of atomic statements, the resultant canonical argument will not have the *same* initial premisses as the supplementation but will have new ones; the rationale will be that, if we can or could establish the initial premisses of the supplementation, then we can establish, or could have established, the new ones.

We are concerned with proof-theoretic justification, on the basis of rules of inference regarded as self-justifying inasmuch as they serve

to determine the meanings of the logical constants they govern. The claim of the rules to be self-justifying itself depends upon the fundamental assumption, since, if it does not hold, the introduction rules cannot together exhaust the canonical means of establishing a statement with the logical constant in question as principal operator, and hence cannot suffice, in the framework of a verificationist meaning-theory, to fix the meaning of that constant. The fundamental assumption is even more essential to the claim of our procedure to *justify* other laws. Unsurprisingly, however, what underpin the fundamental assumption are considerations that are not themselves proof-theoretic but are in a broad sense semantic: we are driven to invoke some notion of *truth*, and so have not achieved a *purely* proof-theoretic justification procedure. Now, at first sight, the fundamental assumption, as applied to the connective 'or', cannot hold good for classical logic: for a classical logician, we know a priori that $\lceil A \text{ or not } A \rceil$ is true, although we may not know, or have any means of discovering, which of the disjuncts is true. But the principle that, if a disjunctive statement is true, one of the two disjuncts must be true, although not holding in certain semantic theories, is of very general validity; if the fundamental assumption, applied to 'or', is to reduce to this, it appears largely banal. It will reduce to this only if it is held that any true statement can be recognised as true by one suitably placed and, if necessary, with sufficient powers. A realist may believe that our powers are too restricted for us to be able to recognise the truth of every true statement, however well placed we are to do so; he must therefore interpret the fundamental assumption, applied to 'or', as meaning that one entitled to assert a disjunction could have recognised one or other disjunct as true if ideally placed to do so and endowed with the requisite powers of observation and intellect. But would not the content of the fundamental assumption then dwindle almost to nothing?

A response to this may be that, in interpreting the fundamental assumption, we have to construe the critical modality 'could have' in whatever way our meaning-theory makes appropriate to the validity of rules of inference. A realist believes that a valid rule is required to preserve a property of truth which may attach to a statement independently of *our* capacity to recognise that it attaches. If he is not to render his own position untenable, he must make this a principle of his meaning-theory: he must hold it to be integral to our understanding of our language that we conceive of our statements as determinately true or false, independently of our capacity to recognise them as such. He will therefore accept the fundamental assumption as holding for disjunction on a lax interpretation of 'could have' such as that suggested above. A verificationist will interpret it much more strictly,

holding the criterion for valid inference to be that someone in a position to verify the premisses was then in a position to verify the conclusion, by the means available to anyone with no more than ordinary powers; or even more strictly, in accordance with the principle sketched above, that there exists an effective procedure by which someone in a position to verify the premisses could at that time have verified the conclusion. The choice between these interpretations is not a matter for logic but for the theory of meaning, just as the choice of a semantic theory is not a matter for logic but for the theory of meaning. Given a semantic theory, logic can determine whether a given formalisation is sound or complete; but whether or not the semantic theory is correct it is not for logic to say. Similarly, logic can determine, for a given set of introduction rules, whether some other set of logical laws can be justified by reference to them; but whether or not the fundamental assumption genuinely holds for those introduction rules, or, if so, under what interpretation, or whether that interpretation is the appropriate one, it is, again, not for logic to say. Proof-theoretic justifications form an interesting alternative to justifications in terms of semantic theories. Neither is autonomous, however: both depend on the defensibility of the meaning-theory within which each finds its proper habitat.

The general principles invoked by this response are sound, but it skirts one critical fact. It is exceedingly plausible that, on a verificationist meaning-theory, the correct logic will be intuitionistic; and we have noted that the standard introduction rules for 'and', 'or', 'if', and the two quantifiers will validate every intuitionistically valid rule involving these constants, where, by the nature of the case, we need to appeal only to those introduction rules governing the logical constants involved in the general formulation of the rule in question. On a realist meaning-theory, however, the correct logic will be classical; and there will be many classically valid laws involving those logical constants that cannot be validated by appeal to the introduction rules governing them, such as those expressed by the classically valid schemata

$$(A \rightarrow B) \vee (B \rightarrow A)$$

$$(A \rightarrow B) \vee A$$

$$((A \rightarrow B) \rightarrow A) \rightarrow A$$

This difference cannot be explained simply in terms of divergent interpretations of the fundamental assumption. The realist will indeed profess to accept the fundamental assumption, applied to disjunction, as exemplified in instances of the law of excluded middle, for example, provided that the assumption is interpreted in terms of his notion

of truth as verifiability by an ideal observer. But the validity of laws like those cited above cannot be established by simply applying the fundamental assumption, first to sentences of the given form, and then to the subsentences resulting from its application, so as to obtain supplementations whose initial premisses are all atomic. The realist does not believe that even the ideal observer could establish every true complex statement built up by the binary sentential connectives from atomic premisses he had verified by observation; rather, he would need to invoke negations of atomic statements as well. The realist's basic principle, that, for every atomic statement, an ideal observer could verify either it or its negation, will *not* result from applying our fundamental assumption, however interpreted, to a negation-free statement; it is a distinct hypothesis that cannot be incorporated into the proof-theoretic justification procedure.

The fundamental assumption is capable of a great range of interpretations, according to how strictly we construe its critical phrase 'could have'. When applied to disjunction, it will be obviously false under the strictest possible understanding of that phrase, since no conventions can bar incomplete testimony or defective memories. Under a variety of interpretations that equate 'could have been verified' with 'is true', it will almost always be sound. A notorious exception will be any semantic theory for quantum-mechanical statements that denies the applicability to them of classical logic, but treats quantum logic as the strongest logic that holds good for them. The critical question, however, is not whether the fundamental assumption holds, but whether it is sufficient to ground the validity of all logical laws accepted as valid. Those laws remain invariant under considerable variation in the interpretation of the fundamental assumption, because it will still serve to validate them by the proof-theoretic justification procedure, without the need for further assumptions. When the strong realist interpretation is adopted, however, the situation changes: not all laws can any longer be validated by proof-theoretic means, because their validity depends not only on the fundamental assumption but on the further assumption of bivalence.

The intuitionistic theory of elementary arithmetic with only bounded quantification coincides with the classical theory, since all statements are decidable: so we may say that, for the intuitionist too, classical logic holds good in this limited domain. This conceals the fact that, while every instance of a classically valid schema is true for the intuitionist, it will not be a logical truth for him but will be only an arithmetical theorem, if it is not intuitionistically valid; it holds in virtue of the specific meanings of the arithmetical primitives, and not just of the logical constants. Should a thoroughgoing realist, who be-

believes classical logic to be valid for statements of all kinds, say the same? The question appears absurd, because classical logic for him rests not on the particular meanings our statements happen to have but on the kind of meaning we can give to any statement we can frame: he may hold, for instance, that we can grasp only those propositions for which we can conceive of an ideal observer for whom they would be decidable. The question is, however, whether he must regard such a limitation on what we are capable of understanding or of expressing in our language as a *logical constraint*. The question cannot be answered until we can distinguish by some precise criterion logical notions and principles from non-logical ones. We are free to choose where to draw this line. If, like, Frege, we make 'topic-neutrality' our criterion, then classical logic will remain a strictly logical theory in its entirety; the same will hold if we treat as a logical constant or device any that serves to form complex sentences from simpler ones. It could, alternatively, be proposed to recognise as logical only those such operators and operations as could be completely characterised by self-justifying logical laws—that is, under the proposal we are considering, by introduction rules, under our proof-theoretic definition of the validity of general rules of inference. On such a criterion, the classical operators would not be *purely* logical constants. That is certainly not, in itself, a ground for rejecting classical logic: no edict requires us to use 'or' as a logical constant in this strict sense. It merely gives a sharper and a better grounded principle than we are accustomed to employ for distinguishing what properly belongs to logic from what does not.

The Conditional

The fundamental assumption, when applied to 'if', makes 'If **A**, then **B**' assertible only when an enthymematic logical entailment holds between **A** and **B**, that is, when **A** in combination with arbitrarily many additional assertible premisses logically entails **B**; the additional premisses may include **B** itself, or may include 'Not **A**'. This is unquestionably a *conceivable* meaning for 'if'; but it is not the meaning we ordinarily attach to it, nor that which is attached to it in intuitionistic mathematics. If we do not presume bivalence, we cannot capture the intuitive meaning of 'if' truth-functionally, that is to say, in terms of the truth or falsity of antecedent and consequent, and hence not by any combination of 'and', 'or', and 'not'. 'If **A**, then **B**' says less than 'Either not **A** or **B**' and more than 'Not both **A** and not **B**'. Its fundamental meaning is more naturally regarded as comprised in the elimination rule (*modus ponens*) than in the introduction rule. 'If **A**, then **B**'

ranks as assertible whenever we have ground to be confident that we shall be entitled to assert **B** on any occasion on which we are entitled to assert **A**. Plainly, this happens far more often than when an enthymematic entailment holds. Given such an entailment, we may transform any proof of **A** into a proof of **B** by simply appending to the proof of **A** a proof of **B** from the hypothesis **A**; but intuitionists allow us to assert 'If **A**, then **B**' whenever we have an effective method of transforming any proof of **A** into a proof of **B**, however complex the process of transformation. Outside mathematics, indicative conditionals, when not expressions of intention, are most often asserted on the basis, in whole or part, of experience, as when someone says, 'If you do business with him, he will find some way of cheating you'.

The falsity of the fundamental assumption, applied to 'if', does not necessarily invalidate the proof-theoretic justification procedure, however. We originally admitted, as occurring within deductive proofs of the kind with which we are concerned, boundary rules allowing the inference of an atomic conclusion from atomic premisses: these were, of necessity, left unspecified. Our original intention was that the boundary rules should be deductively valid. If we now include among them principles of non-deductive (and therefore fallible) inference, this will have the effect that a 'valid' argument, even if canonical, may have true initial premisses but a false final conclusion. It will obviously not affect the justification procedure, however, as a means of determining the validity of logical laws. Under the original, restricted notion of boundary rules, such non-deductive principles would correspond to those conditionals we should be willing to assert whose consequents are atomic sentences (closed or open) and whose antecedents are conjunctions of such sentences; let us call these 'basic conditionals'. Evidently, we frequently make conditional assertions whose antecedents or consequents are highly complex. It thus appears that, if we admit only those non-deductive boundary rules that have atomic conclusions, the legitimacy of the justification procedure will depend on how plausible it is that all such conditionals could be derived by logical deduction, finishing with an application of 'if'-introduction, from basic conditionals.

This hypothesis, unfortunately, cannot be sustained. The recalcitrant case is that of a disjunctive consequent. If you tell me, 'If you ask him for a loan, he will either refuse or make an outright gift to you', because you have never known him do anything else, you presumably know which he did on each occasion that you know about; but, since you do not know on what principle he elects to do one or the other, you are not in a position to make a more specific prediction.

We have, however, extended the scope of deductive boundary rules to allow some with complex conclusions, the premisses continuing to be required to be atomic. Should we respond to our difficulty by admitting *non*-deductive boundary rules of this extended kind, with complex conclusions? That would put a great strain on the fundamental assumption as applied to those conclusions, since that assumption would still be required to hold good whenever the atomic hypotheses of the subargument were verified. Doubtless suitable restrictions on the non-deductive boundary rules could be framed; but it seems better to rely on the commonplace that an experientially based conditional will be asserted only as the tacit consequence of some generalised version of it, even though, in complicated circumstances, the proponent might be hard put to it to frame the relevant generalisation. We must therefore turn our attention to the universal quantifier—which we have in any case to consider—in the hope that it will bring with it a solution to our problem.

Universal Quantification

The application of the fundamental assumption to the existential quantifier obviously resembles its application to disjunction; but the universal quantifier, as ordinarily understood, appears not to fit that assumption at all, which amounts to saying that we are entitled to say that something holds of *everything* only when we can show that it must hold of *anything*. It seems highly doubtful that we can hit on a genuine sense in which anyone entitled to assert a universally quantified statement could have arrived at it from the corresponding free-variable statement. Intuitionists would agree. For them, a universally quantified mathematical statement has been proved when we have demonstrated an effective way of obtaining a proof of any given instance. How this can be done will depend on the domain being quantified over. In number theory, for example, the fundamental method is that of mathematical induction; but in all cases it must be allowed that the form of the proof, for a particular instance, may depend upon the instance, and need not take the simple form of replacing a free variable by a term for the element in question. The most natural view, for general contexts, is that our primitive understanding of 'all' is as extending over a finite, surveyable totality, as when a mother says to a child that all his fingernails are dirty, and that its extension over finite but unsurveyable totalities, and further over infinite ones, is arrived at by analogy with this primitive case.

If there is to be a defence of the appeal to the fundamental assumption, as applied to the universal quantifier, it therefore cannot rest

upon its unqualified truth. It is intuitively natural to regard the meanings both of the disjunction operator and of the existential quantifier as determined by their introduction rules: the fundamental role of each is to give an 'incomplete communication' of some more specific truth. For the universal quantifier, however, as for the conditional, it is equally natural to take its meaning as encapsulated in the elimination rule. It is in connection with the consequences we draw from a universally quantified statement, not with our means of arriving at it, that it is correct to say that we can assert about *every* object in the relevant domain just those things we are prepared to assert about *any* such object. That is to say, something will serve us as a ground for asserting a universally quantified statement just in case we take it as entitling us to make that assertion about an arbitrary member of the domain. An enthymematic derivation of the free-variable statement—that is, a logical deduction of the assertion, as applied to an unspecified element, from premisses established empirically or otherwise—is only one such ground. Inductive procedures form the most obvious alternative type.

A universal generalisation is sometimes based on purely deductive inference. From an open atomic sentence ' $F(a)$ ' we may deduce ' $G(a)$ ' by a boundary rule: by if-introduction and universal quantifier-introduction we then arrive at the quantified statement ' $\forall x (F(x) \rightarrow G(x))$ '. It is the evident fact that such methods will not yield all the universal statements we are willing to assert that has led to our doubts about the fundamental assumption. Does this fact show the fundamental assumption, as applied to the universal quantifier, to be false? That depends on a different point in its interpretation. If we take its content to be that we assert a universally quantified statement ' $\forall x A(x)$ ' only when we have a deductively valid derivation of the free-variable form ' $A(a)$ ' from established truths, then indeed it is manifestly false. For our purpose, however, we need not construe it in so strong a sense; it suffices for us that we can always regard the universally quantified statement as derived from the free-variable one, however the latter was arrived at. We can then allow some licence to the derivation of free-variable statements, because the validity of the subargument for a free-variable statement does not require us to be able to convert a supplementation of that subargument itself into a valid canonical argument; we need only be able so to convert a supplementation of an *instance* of the subargument, which is a much weaker condition. The question is, therefore, whether we can, within our framework, accommodate inductive inferences (in the sense of empirical induction) without so far disrupting that framework as to invalidate our proof-theoretic justification procedure.

A universal generalisation based on inspection of a surveyable totality rests, of course, on the knowledge that the elements t_1, \dots, t_n of the totality comprise all those objects satisfying the predicate $\mathbf{A}(\)$ defining the totality. Such knowledge certainly requires that, for each i ($1 \leq i \leq n$), we can give a proof of $\lceil \mathbf{A}(t_i) \rceil$; we are here adopting the convention that t is a constant term denoting the object t . The further component can be expressed by the statement

$$\forall x (\mathbf{A}(x) \rightarrow x = t_1 \vee \dots \vee x = t_n)$$

To avoid having to involve ourselves in a formal treatment of identity, let us instead conceive of it as a quasi-empirical rule of inference, allowing a passage from premisses

$$\mathbf{B}(t_1) \dots \mathbf{B}(t_n)$$

to the free-variable conclusion

$$\mathbf{A}(a) \rightarrow \mathbf{B}(a)$$

here neither $\mathbf{A}(\)$ nor $\mathbf{B}(\)$ is required to be atomic.

Something similar may be envisaged for inductive generalisations. We make a finite number of observations, the results of which may be broken down into a finite number of closed atomic statements. These, combined with premisses $\lceil \mathbf{A}(t_i) \rceil$, may lead to a finite number of conclusions $\lceil \mathbf{B}(t_1) \rceil, \dots, \lceil \mathbf{B}(t_n) \rceil$ of the same form, where, again, for each i ($1 \leq i \leq n$), we can give a valid argument for $\lceil \mathbf{A}(t_i) \rceil$. We shall be prepared to generalise that $\mathbf{B}(\)$ holds good of all the members of a totality if we regard t_1, \dots, t_n as constituting an adequately representative sample of it. Where $\mathbf{A}(\)$ defines the totality, we may then assume that we have a rule of inductive inference allowing us to infer $\lceil \mathbf{A}(a) \rightarrow \mathbf{B}(a) \rceil$ from the premisses $\lceil \mathbf{B}(t_1) \rceil, \dots, \lceil \mathbf{B}(t_n) \rceil$. This is, naturally, a highly cavalier way of describing inductive inference, which, in its sophisticated forms, involves assurance of total relevant available evidence, estimation of a priori probabilities, sampling methods, tests of statistical significance, and much else. That does not matter here: we are concerned not to analyse inductive inference but to defend our justification procedure from doubts that assail the fundamental assumption that underlies it. The tests of statistical significance, and so on, may therefore be viewed simply as conditions of application for the schema of inductive inference—that is, as criteria for the representativeness of the sample.

What is a representative sample? This must depend on the conclusions $\lceil \mathbf{B}(t_1) \rceil, \dots, \lceil \mathbf{B}(t_n) \rceil$ we are aiming to generalise. We shall regard a sample as representative, relative to those conclusions, if we are confident that, if we were or had been suitably placed to observe any

given object, we should be or have been able to make observations establishing either that it is not a member of the totality, or that, if it is, it satisfies $\mathbf{B}(\)$. This is not something we can hope to *prove*: our criteria for sound inductive inference are designed to ensure that, if the universe is sufficiently orderly, our confidence will be misplaced as seldom as possible. The orderliness of the universe of course requires that what appears random should usually *be* random.

Given an argument whose last step is an application of universal quantifier-introduction, and whose initial premisses are all atomic, our criterion for its validity is the possibility of transforming any instance of the subargument leading to the penultimate line $\lceil \mathbf{P}(a) \rceil$ into a valid canonical argument. There has to be, in other words, an effective means of finding how a canonical verification of any instance of that free-variable statement can be or could have been obtained. Suppose, now, that $\lceil \mathbf{P}(a) \rceil$ is of the form $\lceil \mathbf{A}(a) \rightarrow \mathbf{B}(a) \rceil$, and has been obtained by the inductive inference rule from premisses $\lceil \mathbf{B}(t_1) \rceil, \dots, \lceil \mathbf{B}(t_n) \rceil$. We have, for any instance $\lceil \mathbf{P}(s) \rceil$, to find a valid canonical argument for it: and that involves finding a valid deduction of $\lceil \mathbf{B}(s) \rceil$ from the hypothesis $\lceil \mathbf{A}(s) \rceil$. If s is one of the terms t_1, \dots, t_n , we already have an outright proof of $\lceil \mathbf{B}(s) \rceil$. For the rest, we must allow for the fact that t_1, \dots, t_n are not logically guaranteed to have formed a genuinely representative sample of the totality, even if the formal conditions for a correct inductive inference were satisfied. If it is, the conditions stated above for a sample to be representative ensure that the required argument can be given, invoking only new observational premisses which we can obtain, or could have obtained, but no non-logical rule of inference. If we do not have a genuinely representative sample, we shall of course be unable to prove the argument valid; but that is as well, since there will be no sense in which it is.

A Summing Up

Our examination of the fundamental assumption has left it very shaky. As applied to the disjunction operator, we have had to interpret it quite broadly; the need for this exemplified a general feature of reasoning about empirical matters, namely, the pervasive decay of information. Unlike mathematical information, empirical information decays at two stages: in the process of acquisition, and in the course of retention and transmission. An attendant directing theatre-goers to different entrances according to the colours of their tickets might even register that a ticket was yellow or green, without registering which it was, if holders of tickets of either colour were to use the same entrance; even our observations are incomplete in the sense that we do not and

cannot take in every detail of what is in our sensory fields. That information decays yet further in memory and in the process of being communicated is evident. In mathematics, any effective procedure remains eternally available to be executed; in the world of our experience, the opportunity for inspection and verification is fleeting.

Worse yet, as applied to the conditional and the universal quantifier, we have had to concede that the fundamental assumption is not literally true. The meaning neither of 'if' nor of the universal quantifier is completely determined by the introduction rule governing it: rather, that rule is, in each case, a specialisation to the realm of logic of a more general principle. In both cases, we recognise as legitimate grounds for assertion what does not guarantee the correctness of the assertion, being willing to believe, and to assert, much more than we have *conclusively* established. In both cases, the fundamental assumption can be maintained in the narrow sense that the last step in establishing a conditional or universally quantified statement as true can be taken to be an application of the introduction rule. In order to do this, however, appeal to non-deductive principles must be admitted into the subordinate deduction—the subargument to the consequent of the conditional from the antecedent as hypothesis in the one case, or to the free-variable statement in the other; and the meaning of 'if' or of the quantifier depends on what non-deductive principles are allowed. Accordingly, neither operator is a purely logical constant, judged from this standpoint. The full content of either, in empirical or even in mathematical contexts, cannot be expressed in purely logical terms. In a broader sense, therefore, the fundamental assumption fails for both operators.

Thus our problem has been to find a way to cordon off those operations, other than appeal to the introduction rule, leading to the assertion of a conditional or a universally quantified statement, so that the falsity of the fundamental assumption would not invalidate the proof-theoretic justification procedure that apparently depends on it. The admission of non-deductive principles of inference has entailed severe disadvantages. We have had to loosen our conception of a valid argument: by allowing an argument to invoke non-deductive rules, in order to arrive at universally quantified statements, we have had in effect to replace the notion of a valid argument by something like that of an admissible one, all this to ensure that all closed logically complex statements, if correctly arrived at, can be arrived at by an introduction rule. We have also had to permit the valid canonical argument for a supplementation of a given argument, which we want to show to be valid, to appeal to *new* atomic premisses, as long as we are in a position to feel assured that they will be, or would have been, available.

For all that, it is clear that, however urgent these matters are for one wishing to construct a verificationist meaning-theory reasonably faithful to our practice, they no more affect our estimation of the validity of logical rules of inference than the fact that we sometimes make faulty observations, and hence draw conclusions from false premisses. As already observed, our justification procedure will readily validate all the laws of first-order intuitionistic logic, at least of its negation-free fragment. Those laws are not going to be called into question by any uncertainties over the scope or status of the fundamental assumption, precisely because the classical logician will admit that assumption, interpreted in terms of an ideal observer. Nevertheless, we have seen that the fundamental assumption, even so interpreted, will not suffice to validate all the laws of classical logic by proof-theoretic means. That is not a condemnation of classical logic, since there is no a priori reason to assume that meanings of the logical constants can be wholly specified by any set of self-justifying laws. The proof-theoretic justification procedure itself is elegant; but, in vindicating its applicability to arguments within empirical discourse, we have had to exchange this elegance for an unattractive messiness. It remains that the laws that would hold good if our introduction rules really did completely determine the meanings of the logical constants, and if the fundamental assumption held literally and under its most straightforward interpretation, are just those that hold good when we allow both for the decay of information and for reliance on less than conclusive grounds for assertion. It is only if we begin with logical laws, like those of classical logic, which violate the fundamental assumption even before such allowance is made, that we shall be unable to justify those laws by our proof-theoretic procedure.