APQ LIBRARY OF PHILOSOPHY

Edited by NICHOLAS RESCHER

# THE LOGIC OF INCONSISTENCY

## A Study in Non-Standard Possible-World Semantics and Ontology

NICHOLAS RESCHER
AND
ROBERT BRANDOM

Oxford, 1980

PUBLISHED BY BASIL BLACKWELL

ISBN 0 631 11581 1

British Library Cataloguing in Publication Data

Rescher, Nicholas
The logic of inconsistency.—('American philosophical quarterly' library of philosophy; Vol. 5).
1. Belief and doubt 2. Rationalism
I. Title II. Brandom, Robert III. Series
121'.6 BD215

For
GEORG HENRIK VON WRIGHT

PRINTED IN GREAT BRITAIN
Billing & Sons Limited
Guildford, London and Worcester

# CONTENTS

# SECTION 18

## The Lattice of Possible Worlds

The manifold of non-standard possible worlds contains an element of structure which its basis, the realm of standard worlds, does not. For each element in this domain is constructed in a definite way by superposition and schematization, out of other elements. In this section we will give a precise mathematical characterization of this added element of structure, and in the sequel we will exploit it to show how the whole apparatus of non-standard entities can be turned to the philosophical purpose of an analysis of belief and cognitive inquiry. Finally, we will use some of the details of this application of the structure of inconsistent and incomplete entities to argue for the ontological *parity* of standard and non-standard individuals and worlds.

Non-standard possible worlds as we have defined them can be made to form a mathematical *lattice* under the operations of superposition and schematization. To see this, consider the set of all standard algebraic structures which are models of theories of some first-order language *L*, according to the standard, Tarskian account of satisfaction. Call these non-standard worlds of *constructive degree zero*. The set of non-standard worlds of constructive degree one is the set of all non-standard worlds which are the result of schematizing or superposing together any set of non-standard structures of constructive degree zero. In general, the set of non-standard worlds of constructive degree $n+1$ is the set of superpositions and schematizations of subsets of the set of non-standard worlds of constructive degree less than or equal to $n$. The set of non-standard structures is the set of all non-standard structures of any constructive degree.

We will construct a lattice on this domain of non-standard possible worlds (which we have just seen includes the standard worlds), using the algebraic operations of superposition and schematization which we have defined on that domain. Thus given any two elements $a$, $b$ of the set of non-standard possible worlds, that set contains also $a \uplus b$ and $a \Cap b$. We cannot simply identify our operations of superposition $\uplus$ with the lattice-theoretic "join" operation, and our operation of schematization $\Cap$ with the lattice-theoretic "meet" operation, however. To form a lattice, the operations must be shown to meet the following four conditions[87] (of which the first is a consequence of the last three, which are mutually independent).



L1: $a \Cap a = a, \quad a \uplus a = a$
L2: $a \Cap b = b \Cap a, \quad a \uplus b = b \uplus a$
L3: $a \Cap (b \Cap c) = (a \Cap b) \Cap c, \quad a \uplus (b \uplus c) = (a \uplus b) \uplus c$
L4: $a \Cap (a \uplus b) = a \uplus (a \Cap b) = a$

Non-standard structures as we have defined them do not quite satisfy these conditions, but we can easily transform the set of worlds which doesn't meet these conditions into one which does. We do this by defining an equivalence relation on the original set which is *stipulated* to satisfy L1–L4, and then work on the lattice of equivalence classes under that relation. Such a stipulation requires some care, however, so let's look a little more closely at the problem.

Our major concern is to keep the semantic consequences of our algebraic construction (as summed up in (T1) and (T2)) intact. In Appendix IV we saw how the semantic properties of our construction could be retained through various sorts of generalizations of the construction for which we gave detailed demonstrations of the central theorems. We must be sure that the same properties will be retained in the lattice generated by shrinking the domain of non-standard possible worlds to equivalence classes according to L1–L4, and using the meet and join operations induced on those equivalence classes by schematization and superposition in the original domain. What are those semantic properties? If we represent the set of sentences satisfied (non-standardly) by a structure $w$ as $Tw$, the condition is that

(1) $T(w_1 \uplus w_2 \uplus w_3 \uplus \ldots \uplus w_n \ldots) = Tw_1 \cup Tw_2 \cup \ldots Tw_n \cup \ldots$
and
(2) $T(w_1 \Cap w_2 \Cap w_3 \ldots \Cap w_n \ldots) = Tw_1 \cap Tw_2 \cap \ldots Tw_n \cap \ldots$

Consider for the moment just the entities which we operate on the right side of these equations. These are sets of sentences, which we union and intersect at will, and all are subsets of the set of all sentences of *L* (we can describe them better than that, as we will see below). Such subsets clearly form a lattice under set-theoretic union and intersection. What we want to do is to construct a homomorphic lattice of non-standard possible worlds such that when we *shrink* the superposition–schematization lattice of non-standard algebraic structures according to the equivalence relation "(non-standardly) satisfies the same set of sentences as", we get just the semantic lattice defined by union and intersection of sets of sentences. Our constructions so far do not enable us to do this, even with the generalizations we sketched, since algebraic superposition and schematization don't meet the lattice axioms L1–L4. However, what we propose to do is to take an equivalence relation stipulated to satisfy those axioms, shrink the set of non-standard possible worlds according to *that* relation, and then establish the desired homomorphism between the resulting lattice of equivalence classes and the semantic lattice. This procedure will accordingly be justified just in case the equivalence relations needed to convert our set of non-standard worlds with algebraic operations of superposition and schematization into a genuine lattice are *restrictions* of the equivalence relation "(non-standardly) satisfies the same set of sentences as". In short, we can simply *stipulate* that, e.g., commutativity (L2), holds for non-standard structures, just in case $T(w_1 \uplus w_2) = T(w_2 \uplus w_1)$ (and dually).

This condition turns out to be satisfied easily. Although our notation as presented so far distinguishes between $w$ and $w \Cap w$ and $w \uplus w$, the generalized forms of our major theorems (T1) and (T2) show us that $T(w \Cap w) = Tw = T(w \uplus w)$. Thus we are entitled to consider $w$, $w \Cap w$, and $w \uplus w$ as all one non-standard possible world, specified in various (and variously perspicuous) ways. Indeed, it is clear that each of the identities asserted by L1–L4 above are asserted between non-standard worlds which satisfy the same sets of sentences of $L$, and are thus acceptable.

More specifically, we define a relation R which is symmetric, transitive, and reflexive, and such that the following eight relations hold for all $w_1$, $w_2$, and $w_3$, non-standard worlds of any constructive degree:

$$\mathrm{R}(w_1 \Cap w_1, w_1),\ \mathrm{R}(w_1 \uplus w_1, w_1),\ \mathrm{R}(w_1 \Cap w_2, w_2 \Cap w_1),$$
$$\mathrm{R}(w_1 \uplus w_2, w_2 \uplus w_1),\ \mathrm{R}(w_1 \Cap (w_2 \Cap w_3), (w_1 \Cap w_2) \Cap w_3),$$
$$\mathrm{R}(w_1 \uplus (w_2 \uplus w_3), (w_1 \uplus w_2) \uplus w_3),$$
$$\mathrm{R}(w_1 \Cap (w_2 \uplus w_1), w_1 \uplus (w_2 \Cap w_1)),\ \mathrm{R}(w_1 \Cap (w_2 \uplus w_1), w_1).$$

Henceforth the expression "(non-standard) possible world" is to be understood as referring not to such worlds as they have been defined up to now, but to the equivalence classes of such worlds defined by the relation R. The stipulations specifying R above are tailored to make the domain of possible worlds (equivalence classes under R) thus defined into a lattice, according to conditions (L1)–(L4).

But can we simply identify a world with a *class* of worlds? What is the significance of such a stipulation, aside from the nice formal structure which results from it? Consider lattice axiom L2, commutativity. It is clear that our introduction and motivation of the



operations of superposition and schematization did not involve assigning any significance to the *order* in which two worlds are fused. It is merely an unfortunate side-effect of the notation we use to specify the product of such a world-fusion that two different such specifications ($w_1 \Cap w_2$ and $w_2 \Cap w_1$) are possible. Clearly stipulating that these two world specifications refer to the *same* world (which is the effect of turning our attention to the R-shrink of our original domain) merely undoes a spurious distinction otherwise enforced by an unperspicuous notation. Similarly, lattice axiom L3, associativity, also undoes a notational infelicity. Nothing in our motivating remarks suggested that the schematization (or superposition) of three worlds had to proceed in any particular order, or even had to proceed by the fusion of *pairs* of worlds at all. Our previous notation provided two different ways of specifying what is intuitively *one* world, $w_1 \Cap w_2 \Cap w_3$. Our current stipulation just removes this misleading suggestion.

The remaining two lattice-axioms, the absorption axioms L1 and L4, resolve another such difficulty. Our original motivating remarks concerning the ontological operations of schematization and superposition envisaged the results of these operations as worlds where, respectively, all and only those situations obtained which obtained in *both* of the relevant base-worlds, and where all and only those situations obtained which obtained in *either* of the base-worlds. The superposition or schematization of a world with *itself* ought then just to be that world again. Further, $w_1 \uplus (w_1 \Cap w_2)$ ought to consist of those states of affairs which obtain in $w_1$ *or* obtain in $w_1$ *and* in $w_2$, namely just those which obtain in $w_1$.

In fact these stipulations make less difference than at first might appear to be the case when we think of them as replacing single worlds with equivalence classes of worlds. In fact the equivalence relations holds between world-*designations* just in case all refer to a single world. In the case of the equivalence of specification induced by stipulation L2 and L3, we simply agree not to distinguish worlds with mutually commutative or associative specifications. In the case of L1 and L4, we can simply take the world $w_1$ to be the *canonical representative* of the equivalence class of world-designations which includes $w_1 \Cap w_1$, $w_1 \uplus w_1$, $w_1 \Cap (w_1 \uplus w_2)$, $w_1 \uplus (w_1 \Cap w_2)$, and so on.

It is important to notice that we are thus semantically entitled to stipulate the identities L1–L4 only at the level of *worlds*. The counterpart principles, applied to the *individuals* in those worlds will not apply. Thus we may not stipulate that every element specification of the form $x_1 \mathrm{O}_1 x_2 \mathrm{O}_2 \ldots x_n \mathrm{O}_n x_{n+1} \mathrm{O}_{n+1} \ldots$ (each $\mathrm{O}_i = \uplus$ or $\Cap$)

where $x_n = x_{n+1}$ can be replaced by one of the form $x_1O_1x_2O_2 \ldots x_nO_{n+1}x_{n+2} \ldots$ without violating our semantic principles. In this case, some element-specifications would have fewer components than others, and our definition of satisfaction presupposes that this was not the case. Nor can we specify even that when *every* element-specification in a given domain has such a redundancy at the $n^{th}$ and $n+1^{st}$ places, that *all* should be replaced as above. For this would reduce $w_1 \cap w_2$ to $w_1$ in every case in which the base worlds $w_1$ and $w_2$ had a common *domain*. Since such structures can obviously satisfy different sets of sentences, this would be an illegitimate stipulation.

We can thus stipulate the identities L1–L4, and the resulting set of non-standard structures with the operations of superposition and schematization form a lattice. This lattice is *not* isomorphic to the semantic lattice of sets of sentences with the operations of union and intersection, since *many* non-standard structures will satisfy the same set of sentences. Thus, if we take any two non-standard worlds that satisfy exactly the same set of sentences, their schematization and superposition will be further non-standard worlds in which just the same sentences are satisfied. We can be sure that there are such pairs of non-standard worlds because there are such pairs of standard worlds, and they just *are* non-standard worlds of a special kind. Nor is this all. It is clear that any consistent and complete set of sentences will not only have standard models, but will also have *non-standard* models. Thus it is clear that the fusion (whether by schematization or superposition) of any set of standard models of some maximal consistent set **T** of sentences of a language $L$ will also be a model of **T**. Again, if we take for each sentence $t \in \mathbf{T}$ the schematization of the set of all standard models which satisfy $t$, and superpose all of these sets, we will get a non-standard model of **T**. And still other constructions of non-standard models of maximal consistent theories are possible. These non-standard models have domains consisting of quasi-objects with their strange properties of identity and individuation, but in these cases $L$ is not sufficiently expressive to distinguish such models from standard ones. Thus it will turn out that even if in the end some theory emerges from an inquiry as consistent and complete, this will not settle the question of the standardness of the domain it discusses. For the language in which the theory is couched may not be powerful enough to distinguish the two cases. (This situation should be compared with that in arithmetic and analysis upon Robinson's introduction of non-standard models.)[85]



In standard model theory, one seeks to connect a syntactically defined notion of *consistency* with an algebraically defined notion of *having a model*. The semantics which results is *sound* in case every set of sentences that has a model is consistent, and it is *complete* in case every set of sentences that is consistent has a model. In keeping with the terminology of section 7 above, we may call a set of sentences *minimally consistent* iff it has no logical falsehoods as members. Similarly, we may call a set of sentences *minimally complete* iff it contains all the logical tautologies of the appropriate language. Then the model theory we have developed is sound and complete with respect to the set of all sets of sentences which are both minimally consistent and minimally complete. It is sound, because every set of sentences that is satisfied by a non-standard model is minimally consistent. It is complete, because for every minimally consistent set of sentences which is also minimally complete, there is a non-standard model which satisfies *exactly* the elements of that set. This may be seen as follows. By the way in which we defined non-standard structures, each structure of constructive degree n is the result of applying certain well-specified operations to a particular set of structures of constructive degree $n-1$ (recall that each degree includes the union of all structures of smaller degrees). Since structures of constructive degree zero are standard structures, each structure of whatever degree has a definite set of standard structures out of which it is composed. Our stipulation of L1–L4 made it possible for different paths from a base set of standard structures to lead to the same non-standard structure of constructive degree $n$, but not for one such structure to be constructible starting from two different sets of standard base worlds. Algebraically, then, the standard worlds are a *basis* for the non-standard ones, with respect to the operations of superposition and schematization. Further, the generalized versions of the superposition and schematization theorems tell us how to determine what set of sentences a complex non-standard world satisfies, based on its algebraic construction out of standard worlds. So (T1) and (T2) serve to assure that our non-standard model theory will be sound and complete with respect to weakly consistent sets of sentences just in case all such sets are constructible by union and intersection of the sets of sentences satisfied by standard structures. But this last is clear. For we can construct the non-standard world in which, besides the tautologies, only the sentence s holds, by taking the intersection of all the maximal consistent sets of sentences containing s. Any desired weakly consistent set can then be constructed by taking the union of all such intersections, indexed over the sentences in the weakly consistent set. *Semantically*, the set of maximal consistent sets of sentences, and the set of one-element sets of sentences, both form bases for the minimally consistent and complete sets of sentences, with respect to the operations of union and intersection. Thus the mapping which assigns to every non-standard structure the set of sentences of L which it satisfies is a homomorphism from the lattice of non-standard possible worlds with algebraic superposition and schematization as operations into the lattice of minimally consistent and complete sets of sentences with union and intersection as operations. Calling this mapping $h$,

$$h(w_1 \cup w_2) = hw_1 \cup hw_2 \text{ and } h(w_1 \cap w_2) = hw_1 \cap hw_2$$

where the joins are superposition of structures on the left side of the equation, and unions of minimally consistent and complete theories on the right, and the meets are schematizations of structures on the left side of the equation, and intersections of minimally consistent and complete theories on the right.

This result, which was really implicit already in our discussion after the detailed demonstration of the initial forms of the superposition and schematization theorems, is not analogous to any result of standard model theory. For the lattice structure which is preserved by the mapping $h$ (which is both a join homomorphism and a meet-homomorphism) has no standard analogue. What operations on maximal consistent sets of sentences are we to liken to superposition and schematization in the non-standard case? Union and intersection of maximal consistent sets of sentences yield such sets only in degenerate cases. Nor are there parallel, suitably closed, operations on standard algebraic structures. In standard model theory one shows an isomorphism between the set of material-equivalence classes of sentences of $L$ with the partial ordering induced by the entailment relation, and the set of semantic-equivalence classes of standard models with the partial ordering of set-theoretic inclusion (The Deduction Theorem).[86] A version of this theorem can easily be proven with the materials now at hand for non-standard models as well, using collective rather than distributive readings of entailments. Our concern in the rest of this essay, however, will be with those aspects of the non-standard structures, codified in the lattice-homomorphism specified above, which *have* no standard analogue.

# SECTION 19

## Belief and Non-Standard Possible Worlds

The preceding discussion has shown how a consistent logic may be constructed for inconsistent and incomplete theories and how a sound and complete model theory for such weakly inconsistent and incomplete quantified theories can be produced. That model theory enables one to describe in detail the quasi-objects which inconsistent and incomplete theories are theories *of*. In the philosophical tradition, such Meinongian objects are generally thought of as *ideal* entities, objects of some propositional attitude, paradigmatically of *belief*. The next few sections will investigate the suitability of non-standard worlds for the representation of such attitudes.

Our concern throughout this essay has been with the ontological respectability of non-standard possible worlds, with their claim to citizenship in the republic of possible worlds with the same civil status as standard, consistent and complete, worlds. The consideration of the aptness of non-standard worlds for the representation of belief is not a retreat from this position. On the contrary, it will be shown that although non-standard worlds can usefully model what is *believed* to be the case, there is no reason to conclude that non-standard worlds are therefore ideal in some sense in which standard possible worlds are not. In particular, it will be argued that there is no more reason to believe that non-standard possible worlds represent *epistemic* states of affairs (e.g., merely subjective beliefs) while standard worlds alone can represent *ontic* states of affairs (e.g., how things objectively are or might be) than the other way around. Thus we will argue for *parity* of ontological status between standard and non-standard possible worlds and their inhabitants. We will show in some detail how standard worlds can be considered as the outcome of an *ideal* process of inquiry, while everything that *actually* occurs in the inquiry is represented by non-standard possible worlds. This demonstration will thus parallel our previous exhibition of non-standard worlds as the results of operations performed on standard ones. The conclusion towards which the discussion moves is that whatever ontological status is assigned to standard, consistent and complete possible worlds, ought also to be assigned to non-standard possible worlds.

Non-standard worlds are attractive as an apparatus for characterizing belief primarily because they provide a precise, formalizable semantics of minimally consistent and complete theories, and the set of beliefs held at a time by an individual is expressible by such a theory. The set of sentences believed by a speaker at a time, approximated by the set of sentences that speaker is disposed sincerely to assent to under some hypothetical standard conversational conditions, will in general meet our specifications for minimally consistent and incomplete theories, to which our non-standard meta-theory applies, and will not in general meet the specifications of a *strong* consistency within which context alone standard logic finds application. The minimal consistency and completeness condition demands that (i) all logical truths be believed, and that (ii) no logically self-contradictory thesis be believed, and (iii) that all sentences deducible from believed sentences by one-premise inferences of first-order logic is believed. Anyone whose beliefs putatively did not satisfy these three conditions, someone who is willing to assent sincerely to logical contradictions and is unwilling to so assent to tautologies or who does not believe the *immediate* consequences of his beliefs, simply offers the best evidence imaginable that his language is not the one we though it was. Far from forcing the admission of sets of beliefs which are not minimally consistent and complete, such a situation would show that the beliefs of the individual in question had not been understood.

On the other hand, intelligibility of beliefs imposes no stronger requirement than minimal consistency and completeness, since we have shown how a standardly consistent and complete meta-theory for minimal consistency can be developed. One therefore need not attribute consistent and complete sets of beliefs to an individual in order to express his attitudes coherently. It is here that the signal virtue of the analysis in terms of non-standard possible worlds for the expression of belief appears. For the paradigmatic expression of a belief is a disposition sincerely to assent to a statement when queried under standard conversational conditions, and such dispositions are typically neither consistent nor complete with respect to the statements of any language. As for completeness, even in the context of pure number theory there are conjectures whose truth or falsity has not been demonstrated and about which one may reasonably remain, not merely agnostic, but adoxastic—entirely without opinion. How much less complete are our beliefs outside of mathematics! Most of us have no views about the relative abundances of silicon and aluminium in distant galaxies, the presence or absence of red nettles in Tibet, whether the number of single-family dwellings in the Pittsburgh city-limits is odd or even, and so on. To believe rationally by our canons is often to refuse credence both to the claim, e.g., that it will rain tomorrow *and* to the claim that it will not (though of course, in accord with the requirements of minimal consistency, one will continue to believe that it either will rain tomorrow or not, and one will not believe that it will *both* rain tomorrow and not), when there is no evidence at all available, or if the only evidence available is the conflicting testimony of two equally unreliable informants.



Nor are our dispositions to assent to statements limited to consistent sets of statements. Even when we are maximally careful about the consistency of our theories, we may be done in by their complexity. Thus even in mathematics and logic, there are such familiar examples as Frege's inconsistent second-order logic, and Quine's difficulties with his version of set theory. At a more mundane level, it is common for an individual to be simultaneously disposed to assent to a statement if queried and to be disposed to assent to the denial of that statement if queried in some variant context (once again, this does not mean that one is ever disposed to assent to the conjunction of a statement and its denial).

With our apparatus, we can express such contextual relativity of belief, representing the various contexts by different subsets of beliefs of the speaker about the circumstances of the discourse, and the different dispositions to assent as the results of conjunctive multi-premise *inferences* from those beliefs.

As long as the dispositions to assent to the statement and its denial are not made the objects of conscious scrutiny (for instance by being actualized), they can coexist in peace. Such situations are common, and are most aptly described as the holding of (minimally) inconsistent beliefs. Any inquiry in which one comes to hold new beliefs, expressed by novel dispositions to assent in conversation, leaves the inquirer with the substantive task of canvassing old beliefs and their associated conversational dispositions for statements which entail the denial of some conclusion reached in the course of the inquiry. This separate process, phenomenologically familiar as "coming to realize the significance" of some view, or "allowing a conclusion to sink in", is the attempt to rectify the grosser inconsistencies in one's beliefs which result from reaching novel conclusions. Such a process is rarely completed, and its completion is certainly not a precondition of believing the novel conclusion which was the result of one's inquiry. So although there is indeed a force which acts on our beliefs at various times, seeking to drive them in the direction of greater consistency (about which more later), it is not the case that simply adopting a statement not previously believed automatically purges one's beliefs of all those statements which in the context of some conjunction of premisses entail the denial of the novel claim. (Indeed, as Socrates taught us, one's entire intellectual energies could usefully be expended in the effort to render more consistent even our commonest beliefs.)

The phenomenon of inconsistent belief would not be worth belaboring in this fashion were it not the case that it is the source of considerable difficulty in the representation of beliefs by standard possible worlds (which model only strongly consistent sets of sentences) and of corresponding ingenuity on the part of theorists seeking to employ the standard apparatus for that purpose. The difficulty arises not so much with expressing inconsistent beliefs in the first place, as in dealing with inference. For one can represent beliefs with a *set* of standard possible worlds, in some of which a statement p holds, and in others of which its denial $\bar{p}$ holds. But one of the simplest logical inferences one can make is from a statement and its denial to any arbitrary thesis whatsoever. The standard theorist is thus faced with a dilemma: Either deny minimal logical competence to believers, or face the collapse of all inconsistent belief-sets into the identical absurd set of all the sentences of the language (self-contradictory or not). The second alternative is clearly unacceptable, as it trivializes the notion of belief beyond its fitness for any explanatory task. The first alternative is only slightly more palatable. For what is the point of representing beliefs if one cannot also represent the inferences by which they are transformed in the course of inquiry? On the other hand, beliefs which were inferentially isolated from their fellows by the complete logical insulation entailed by a denial of even minimal logical competence can play little explanatory role in an account of behaviour or inquiry. Some sort of restriction on inference must be formulated which prevents the absurdity of identifying the sort of practically uninfectious inconsistency referred to above with believing all the sentences of one's language, while retaining enough inferential capacity for beliefs to be cognitively efficacious.

Our account of non-standard deductive theory and model theory, of course, is designed precisely to resolve this dilemma. Each non-standard world satisfies the logical consequences of every thesis which it satisfies. But it will satisfy the standard consequences of a collection of statements (e.g., $p$, $\sim p$) *only* if their conjunction is also satisfied (as will *never* be the case in a non-standard world for $p$, $\sim p$). Deductively, this is a consequence of the demand that rules of inference be read collectively rather than distributively, and semantically it is a consequence of the way in which non-standard model structures were constructed. So both the inferential paralysis by deductive insulation of the first horn of the dilemma and the inferential collapse into universal assertion of the second are avoided. The two respects of non-standardness which generate non-standard worlds are thus both natural expressions of principles one would want antecedently to apply to beliefs.



Non-standard worlds thus give us the means to represent a whole spectrum of varieties of belief for each particular claim $p$ which might in some sense be believed. We can, for instance, represent *bare* belief that $p$ by a world in which $p$ holds but none of the conjunctions of $p$ with other beliefs holds in that world. Although $p$ would be assented to, it has not been absorbed and integrated into the beliefs of the individual. At the other end of the spectrum, *total* belief that $p$ can be represented by a world in which $p$ holds, and if $q$ is any other thesis consistent with $p$ which holds in the world, $p \,\&\, q$ holds. Because of the way we developed the semantics of non-standard worlds, this entails that the consequents of any conditionals with $p$ as their antecedents will hold in the world as well. In such a case, the belief that $p$ has "sunk in," been considered in the contexts of all the other beliefs, and its consequences in those contexts rigorously adhered to. Along another dimension, we may consider the minimal belief in $p$ as one where its negation is represented as also holding in the world, or, less drastically, where the negations of some conjunctions containing $p$ hold in that world. Thus possible worlds induce a natural array of degrees of belief for each proposition, but these varieties are indexed by the non-standard truth tables involving $p$ (rather than for instance, the real interval between 0 and 1). We believe this greater flexibility in the representation of the varieties of belief-attitudes one may have toward a single claim to be one of the prime virtues of the apparatus of non-standard possible worlds.

To be sure, using the apparatus of non-standard possible worlds to describe belief is not simply tantamount to the claim that the beliefs of an individual speaker at a time constitute a set of sentences in some language which is only weakly consistent. For the model of beliefs ought to be, not the set of sentences of some language, but a belief *world* or set of worlds, which *satisfy* such a set of sentences, that is, *the world which is just exactly as some individual believes it to be*. This is non-standard world, inconsistent and incomplete. The paradigmatic expression of a belief is a disposition to sincerely assent to a statement in some language, but this does not entail that all beliefs are adequately captured in a set of sentences in any particular language. We have argued above that the set of sentences of a language which are the expressions in that language of a set of beliefs will be only minimally consistent. This is not to say that there is not more to belief than this set of sentences. That component of one's beliefs which is not expressible in the particular language being considered is to be represented by the particularities of a non-standard world which, while it satisfies the minimally consistent and complete set of sentences of *L* which are the linguistic expression of those beliefs, also differs from other models of those sentences in various ways not specifiable in *L*. This stipulation amounts only to the recognition that there need be no unique language in which all of one's beliefs are expressible.

Let us represent the beliefs of an individual at a time with a non-standard possible world, his *belief*-world at that time. It was argued just above that this association need not be taken as implying that the individual whose beliefs are so represented could tell us everything about the world we use to represent his beliefs (even though this "everything" makes a weaker claim than might at first appear, since the world to be specified is itself incomplete). It is also worth pointing out that while the believer on this account believes in the existence of things we would describe as quasi-objects, with relations obtaining among them which conform to the principles expressed in our model theory (by means of a complicated structure of sub-relations into which the quasi-relation is partitioned), this does not imply there is anything in those beliefs which corresponds to the particular notational devices we have used in the metalanguage in which the model theory of non-standard theories is couched. In our treatment of non-standard worlds, for instance, we represented element-specifications in a notation which exhibits their compositional properties in a particularly perspicuous manner—each element-specification had the explicit form of a product by superposition and schematization of standard individuals. This notation is not expressible in the language L, whose weakly consistent theories are being modelled. One should therefore not require that a believer evince an ability explicitly to deal with element-specifications. We seek to represent beliefs by non-standard possible worlds, composed of quasi-objects and quasi-relations among them. The minimally consistent set of statements of *L* believed are true in this world. But the machinery of element-specifications and partitioned relations is *our* notation, a particular metalanguage we have adopted for discussing such worlds of quasi-objects. It is the worlds themselves,



not our metalinguistic representations of them in standard algebraic language, which are here to be taken as models of belief. Belief is a relation to such a world. That relation need not express perspicuously the potential decomposition into standard worlds which the notation in which the model theory was expressed was designed to make clear. In particular, the representation of belief by an associated non-standard belief-world is not to be taken as a claim that introspection on the part of the believer can recover a representation of that world such as we have employed, from which an account of its decomposition into standard worlds is immediate. It is not a trivial task for a believer even to be able to represent the (minimally consistent set of) sentences of *L* which he believes as the product in some (not unique) fashion of the union and intersection of complete and consistent theories of *L*. How much more difficult, then, it is to represent the facets of one's beliefs which are *not* expressible in *L* (indicated by a selection among the non-standard worlds satisfying the right sentences of *L*) by the schematization and superposition of *objects*. This last sort of project is internal to the theory we are presenting of non-standard worlds as representatives of beliefs, and is a task of the theoretician, not a capacity attributed to believers simply in virtue of being such.

So belief is to be taken as a relation a believer has at a time to a non-standard world. That world must satisfy all and only the statements of a language *L* which the individual believes. But one can believe more than can be said in any particular language, and this unverbalizable surplus is represented by the selection of a particular non-standard world from all of those satisfying the right sentences. The believer is not presumed to have the conceptual means for expressing this selection explicitly (everything so expressible gets packed into the sentences of a suitably powerful language *L*).

In the next section we will consider some aspects of *inquiry* (the methodical transformation of belief) which become apparent when one considers belief according to this model. We will not here address the sticky and interesting question of how linguistic, behavioral, and other kinds of evidence can justify the attribution of particular beliefs (the assertion of a relation to a particular world). We have advanced only general reasons to conclude that this issue will be more tractable once non-standard worlds and incomplete and only weakly consistent sets of sentences are admitted as legitimate tools.

# SECTION 20

## Methodological Realism and the Convergence of Inquiry

The project which will occupy the rest of our essay is that of giving an account of the ontological status of non-standard worlds as compared to standard, consistent and complete possible worlds. The apparatus of non-standard possible worlds as we have developed it so far is invaluable in the perspicuous presentation of an account of the various senses of ideality and reality whose claims must be adjudicated in order to settle the issue of ontological status. Accordingly, we will use that apparatus to define the notion of an *inquiry* which seeks an accurate and complete representation of things as they really are, and to define various constraints on the methodologies governing such inquiries as they relate to the reality which is presumed to be their subject-matter. Thus equipped, we will appeal to the example of the metaphysical views of Peirce as a concrete historical reference-point for a canvassing of alternative metaphysical views couched in the language of non-standard possible worlds.

Inquiry is the rationally controlled transformation of belief over time. The classical philosophical tradition has focussed its attempts at epistemological explication of this process on two particularly prominent motive forces which help drive inquiries in general. First, an *empirical* component of inquiry may be discerned: one constantly acquires new beliefs and discards old ones as the result of observations made in accord with widely shared non-inferential reporting practices. Second, a *rational* component of inquiry can be distinguished. Within a specified observational situation, one's beliefs are transformed according to complicated canons of inferential coherence, which dictate the practical incompatibility of some set of previously held beliefs, or require that such a set be amplified by the addition of a belief it entails (which entailment had been hitherto unnoticed). The observational and inferential practices which control the transformation of belief in this way may be referred to collectively as the *cognitive methodology* of the believer.

The stratagem of representing the beliefs of a speaker at a time by a non-standard possible world which satisfies all and only those (in general only weakly consistent) beliefs allows the exhibition of any particular inquiry as a *path* from one non-standard world to



another, on the complete lattice of all such worlds. When a measurement provides a firm figure for the charge of a new particle under investigation, for example, the inquirer moves from a belief-world in which the quasi-objects which are the particles as they are believed to be are relatively more indeterminate to one in which they are relatively less indeterminate. Prior to the experiment, the inquirer's belief-world may conceivably have satisfied mutually inconsistent claims about the charge of the particle in question, or it may have satisfied no claims at all regarding that charge (i.e. neither was it true in the initial belief-world that the particle had a determinate charge, nor that it did not). The experimental measurement, however, moves the inquirer to a new belief-world. Of course, the belief-world arrived at in this fashion may still satisfy mutually inconsistent claims about the particle, and may remain incomplete in other regards.

In the sorts of cases which come most readily to mind, change of belief may be represented by a move from a world which satisfies a set $S$ of sentences to a world which satisfies a different set $S'$ of sentences, relative to some fixed language. As we have pointed out, however, some changes of belief may be better represented by movement from one non-standard world to another *within* the equivalence class of worlds which satisfy the same set of sentences in that language. Thus, someone may come to believe that a clarinet sounds like *this* and not like *that*, and not have sufficient musical vocabulary to express that alteration of belief linguistically (though the change of belief may have behavioral consequences for discriminations the believer would be disposed to make). The acquisition of such a new belief ought then to be represented as a movement from a belief-world in which clarinets have exceedingly indeterminate sounds (clarinets are quasi-objects in such worlds, their phonic properties incomplete and perhaps inconsistent) to a world in which they have a more determinate sound-range. The sentences of an official language $L$ which are satisfied by these two worlds may be identical. The belief-worlds differ only in the way clarinets sound in them, and that difference need not be expressible in $L$.

A cognitive methodology with empirical and rational components must determine for each belief-world a set of admissible transformations by inquiry. We may represent such a methodology by an accessibility relation between possible belief-worlds, which will in turn determine a set of admissible paths through the lattice of non-standard possible worlds, corresponding to methodologically sanctioned inquiries. Our interest will not be in the details of the constraints which methodologies so codified impose on inquiries which recognize their authority, but concerns rather conditions which it is reasonable to insist that such methodologies must meet if we are to consider them candidate ways of seeking to find out how things really are in the first place. For it is clear that not all disciplined transformation of belief need have the representation of reality as its goal. A monastic community might take as the goal of its discipline the transformation of belief so as systematically to encourage or enforce the development of habits of action taken to be desirable from a religious or moral point of view, caring nothing for the *correctness* of the views thus arrived at, in the *cognitive* sense of that term. The question we wish to address is how to distinguish methodologies whose aim is the cognitive one of discovering how things really are. The task of offering conditions sufficient for such discrimination lies outside the scope of our enterprise. Instead, we will seek to elucidate one popular and plausible contender for the status of *necessary* condition for such discrimination, a view which can be discussed more precisely in terms of non-standard possible worlds.

The view in question may be called "methodological realism", enunciated well by Peirce:

> There may be some questions concerning which the pendulum of opinion would never cease to oscillate, however favorable circumstances may be. But if so, these questions are *ipso facto* not *real* questions, that is to say, are questions to which there is no true answer to be given. (CP 5.461)

Methodological realism requires two things of cognitive inquiries. First, that every question which is admitted as a proper subject of inquiry have some answer, and, second, that the aim of cognitive inquiries be to discover such answers. Thus formulated, the doctrine can admit (as Peirce would not) that there may be questions with true answers which cannot be established by inquiry. Methodological realism describes the *goal* of cognitive inquiry as the settlement of opinion by the determination of the truth. There are two issues here, what sort of methodological constraint expresses the *aim* of inquiry to bring our beliefs to some conclusion, and how "truth" or "how things really are" is to be interpreted as related to the ideal endpoints aimed at by cognitive inquiries. The second issue (including the question of the status of truths undiscoverable by inquiry) will be dealt with below when we discuss the thesis of *ontological* realism, contrasting it with Peirce's ontological idealism. The first issue Peirce approached through the notion of the methodological importance of the *convergence* of inquiry.



We will follow Peirce's suggestion, and express methodological realism by means of a distinction between inquiries, (represented by paths through the lattice of non-standard belief-worlds), which converge and those which diverge. We will not consider Peirce's account in its details since, as Quine has pointed out, it suffers from

> . . . a faulty use of numerical analogy in speaking of the limit of theories, since the notion of a limit depends on that of 'nearer than', which is defined for numbers and not for theories.[87]

Using the lattice-structure which relates non-standard worlds, on the contrary, will enable us to define a precise sense of "limit" and "nearer than" which applies to the belief-worlds which make up inquiries. Methodological realism will then be the claim that lattice-convergence is a regulative ideal of cognitive inquiries. Our apparatus will enable us to distinguish a number of different senses in which this ideal could be taken to constrain particular methodologies.

We have used non-standard possible worlds to represent beliefs, and sequences of them to represent inquiries. We have shown above how the set of non-standard possible worlds can be treated as a complete lattice generated by the operations of superposition and schematization. The notion of convergence (and hence of limit) is defined primarily for topologies, numerical convergence being a special case, deriving from a particular order topology on the natural numbers. Mathematicians have shown how to associate topologies with suitably well-behaved lattices. In Appendix V on Convergence, we show how the intrinsic order topology of the complete lattice of non-standard possible worlds can be used to define convergence of nets (indexed by upper-directed paths), of which sequences or paths are a special case. Intuitively, the sense of "nearness" which is employed in this definition of convergence is nearness in respect of non-standard composition out of standard worlds. It is with respect to the way in which non-standard worlds are constructed by the superposition and schematization of standard and non-standard worlds that a sequence of such worlds may be said to converge. We have *not* imposed any kind of similarity measure on standard, consistent and complete worlds. Our approach does not permit us to compare such standard worlds with one another for "nearness", nor to talk about a convergent sequence of them (except eventually constant ones). For the structure we use to define convergence is the *lattice* structure, which expresses only how a world may be resolved algebraically as the product by superposition and schematization of a set of standard worlds.

To see what is at stake here, it is best to look at a *resolution convergent* inquiry. Such an inquiry is a set of worlds, indexed over time, such that each successive member of the sequence is a component of the preceding world (i.e., is a world which was superposed or schematized with some other to yield the preceding term of the sequence). Such a sequence will always be of finite length, and will terminate in a standard world to which it converges (somewhat trivially, because of its finite length) by our criterion. This process of resolution proceeds by sorting out and untangling the various components of an initial belief-state. The result is a set of beliefs linguistically expressible by a consistent and complete theory, a set of standard beliefs which can be conceived of as having been observed and confused with alternative standard belief worlds, resulting in an overall, non-standard belief world. The process of resolution is thus a certain kind of clarification of one's views' One understands one's own views better when those views can be exhibited in a precise way as a certain kind of product, by superposition and schematization, of standard belief worlds.

Of course, there are many other exigencies of clarity and cogency to which inquiry is subject. Resolution of one's views in this sense is not an over-riding imperative. It may be methodologically more appropriate at a certain stage in an inquiry to seek more empirical information, for instance, than to try to untangle one's present views to see just how one's current belief world is constructed. We shall have more to say on this issue below when we discuss methodologies from the standpoint of our lattice structure. The introduction of this notion of a resolution sequence is meant to give some indication of the kind of convergence according to non-standard construction by superposition and schematization we have captured in our technical formulation. It is insofar as the belief-worlds are inconsistent and incomplete that sequences of them converge on our lattice. It is nearness in point of non-standard construction out of standard worlds which is utilized in the definition. This is not the sort of convergence of sequences of worlds which other methodologists have considered, for the simple reason that they did not consider non-standard worlds at all. Our discussion below will accordingly deal with those aspects of the methodology of inquiries which concern the non-standardness of belief-worlds. The structures we define do *not* (and are not meant to) provide an account of what "convergence" might mean for a sequence of consistent and complete scientific theories, for instance.

The Appendix on convergence defines a precise sense in which inquiries as we have represented them can be said to converge on a settled opinion or way things have been found to be, and we have interpreted methodological realism as the claim that such convergence is a regulative ideal of cognitive inquiry. We have not yet said how this ideal might be expressed in constraints on methodologies which govern inquiries. Before we address such questions as whether it is reasonable to require that every methodologically admissible inquiry converge to some world, it will be useful to consider some more specific conditions which can be formulated in terms of our notion of convergence of inquiry. One generally laudable characteristic of an inquiry is that over the long run it converges to a *consistent* world. That is, we have a bias towards wanting our inquiries to be progressive in the sense of weeding out inconsistent beliefs. Of course it is the import of this whole study that such a requirement ought *not* to be taken as an absolute methodological requirement, but as a nice property, desirable when other things are equal. Our specification of the semantics of inconsistency is meant to show that in various circumstances the general desirability of consistency may be over-ridden by other methodological desiderata. A strategic bias toward consistency need constrain our methodological tactics less than has been previously thought.

The dual notion of *completeness* convergence concerns not the deductively motivated preference for consistency, but the inductively motivated search for completeness in our beliefs. This motive is, it would seem, less urgent than that of consistency. The push for a unified and universal science is not methodologically negligible, but not as urgent as the push for consistency where it is obtainable. That requirement is even more obviously merely a desideratum, to be carefully weighed against competing methodologic considerations. Other things being equal, it is a greater criticism of a theory that it is not consistent (though as we have shown, this is not an insurmountable objection) than it is that it is not complete. Completeness convergence is accordingly a less important methodological recommendation than consistency convergence.

The conjunction of these two conditions is the requirement that an inquiry converge to a standard, consistent *and* complete world. *Standardness convergence* in this sense is a very strong condition to put on one's methodology. To see what sort of a condition it is, consider methodologies and requirements on them somewhat more generally. A methodology is a specification of a set of *admissible* inquiries. One way of doing this is to express the constraint of admissibility by transitive and reflexive accessibility relations on non-standard belief worlds. A sequence of worlds with an initial element, that is, an inquiry, will then be admissible according to the methodology captured in the accessibility relations just in case each element of the chain is accessible from every earlier element. One may, of course, simply specify a set of admissible chains directly, simply as a set of such chains. Any characteristic shared by all inquiries which are admissible according to a particular methodology is *methodologically enforced* or guaranteed.

One such interesting condition is *universal convergence*. This is the requirement that every admissible inquiry converge to some belief world. The methodologies which meet this condition *ensure* that no divergent inquiry will be methodologically acceptable. It is not a trivial task to put conditions on an accessibility relation between belief worlds in such a way as to guarantee convergence in a non-trivial fashion. This condition can be strengthened by combining it with the other conditions we have considered. Thus *universal consistency convergence* will stipulate that each admissible inquiry converges to a consistent world. *Universal completeness convergence* and *universal standard convergence* conditions may be constructed similarly. Strengthening along an independent dimension, we can formulate *unique universal convergence*. This is the requirement that all admissible methodologies converge to the *same* world, although different paths to that goal are allowed. This condition may obviously also be combined with those of consistency, completeness and therefore standardness.

Each of these possible methodological requirements stipulates *a priori* an important characteristic of all inquiries admissible according to a methodology imposing the conditions. The conditions described in the last paragraph all entail the Universal Convergence condition. This means that they all remove from the realm of empirical possibility otherwise methodologically acceptable inquiries which do not converge. We may ask whether this is a question which is appropriately addressed at the level of *a priori* methodology.

SECTION 21

# Peirce and Empirical Convergence

Peirce was inclined to insist that every admissible inquiry be such that it converges. In a passage which combines what we will call below Peirce's ontological idealism with this insistence he says:

> . . . to assert that there are external things which can be known only as exerting a power on our sense, is nothing different from asserting that there is a general *drift* in the history of human thought which will lead it to one general agreement, one catholic consent. (CP 8.12)

We are currently interested in reading the asserted identity in the other direction, as claiming that to be finding out about the external reality Peirce invokes one must drift toward that general agreement. Peirce's views will be considered in more detail below, but we see here that he is willing to insist on convergence to a single belief world as a criterion of an inquiry which claims to address itself to an empirical reality.

A number of questions should be asked about this view. First, even if it is accepted, does it rule out as empirically significant a methodology which does not *guarantee* convergence *a priori*, but merely aspires to it? If a methodology comprises some convergence a *sine qua non* of admissibility, then whether a particular inquiry converges or diverges is an *empirical* matter. One just has to wait and see, engaging in actual inquiries, transforming beliefs according to the methodology, moving from one belief-world to another. Luck may be required to make the transformations of belief which will lead to convergence. Many admissible or obligatory transformations of beliefs may not contribute to convergence at all except in fortunate, empirically determined circumstances. Empirical inquiries must aspire to convergence. But actual inquiries we class as empirical do not obviously *guarantee* convergence, and there are horrible, unempirical methodologies which do. Thus, the method of Tenacity[88] which consists in maintaining one's current beliefs come what may, or various more socially totalitarian versions of this superhuman dogmatism would satisfy the Universal Convergence condition without being empirically praiseworthy thereby. Is it not imaginable, for instance, that a perfectly acceptable methodology would have the characteristic that any admissible inquiries which began with a belief world in a certain region of the lattice (or one with some other property) would converge, while those initiated elsewhere would not? This would be a methodology with an empirical presupposition to its "guarantee" of convergence. Perhaps even our own scientific methodology has this characteristic. Is it clear that current ideas of scientific methodology, if applied to, say, Parmenides' belief-world, would result in the sort of empirical-technical convergence we fondly believe ourselves to have assured beginning with Newton's world? Methodologically *guaranteed* convergence may exhibit only the advantages of theft over toil usually associated with a priorism.

Another question we must ask is whether empirical methodology must seek to ensure *unique* convergence. Peirce says:

> Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside of themselves to one and the same conclusion. This activity of thought by which we are carried, not where we wish, but to a fore-ordained goal, is like the operation of destiny. No modification of the point of view taken, no selection of other facts for study, no natural bent of mind even, can enable a man to escape the pre-destinate opinion. This great hope [in first draft: "law"] is embodied in the conception of truth and reality. The opinion which is fated to be ultimately agreed to by all who investigate is what we mean by the truth, and the object represented in this opinion is the real.(5.407)

Peirce is unambiguously ruling out a methodology according to which, for instance, some inquiries converged to one belief world and the others converged to another, or diverged. The difficult status of this claim is indicated by Peirce's shift over the years in the precise official formulation of it which he endorsed. In the early days of "The Fixation of Belief", he thought of this unique convergence as a *law*, constitutive of empirical inquiry and reality alike. In later years, however, he treats unique convergence as a merely regulative ideal. Unique universal convergence is a state empirical inquirers *aspire* to, an outcome they *hope* and even *believe* will apply to their efforts, but an issue which is essentially empirical and undemonstrable *a priori*. In short, he came to think of unique universal convergence as something which could be *guaranteed* in advance only by such ferociously anti-empirical methodologies as that of Authority ("Everyone believe what *S* now believes!") Such a position is entirely compatible with concurrent belief in a real world that is independent of our actual beliefs and to which those beliefs should aspire to be faithful, that is, is equivalent to belief in the empirical convergence of appropriate inquiries.

These considerations suggest a weakening of Peirce's condition.



If we deny Universal Convergence, *conditional unique universal convergence* may still hold. This condition requires of every admissible inquiry that *if* it converges, then it converges to the same world that every other admissible convergent inquiry of that methodology has as its limit. In such a methodology, it is an empirical matter whether any admissible inquiry will converge or not, but there is only one possible admissible outcome for such convergence if it does occur. The other strengthenings of Universal Convergence which we considered can also be expressed in weaker, conditional forms. Thus *conditional consistency convergence* states that every convergent admissible inquiry in a methodology converges to a consistent world, *conditional completeness convergences* states that every convergent admissible inquiry converges to a complete world, and *conditional standardness convergence* is the conjunction of these two conditions. The sort of methodological thesis exemplified by these various conditions is characteristic of a certain kind of *realism*, which we will consider further below. The leading intuition is first that "good" cognitive inquiries ought to lead us to find out something (settle our beliefs, represented here by convergence, not just of the sets of sentences held true, but of everything believed, convergence of sequence of *worlds*). Second, good inquiries ought to lead us to believe in the *right* world, the one which is real, actual, or distinguished in some other inquiry-driving fashion. One's prejudices about the nature of that world will accordingly color the methodological proscriptions concerning admissible inquiries one is prepared to subscribe to. In particular, if one is in some fashion convinced *a priori* that the world our empirical inquiries strive to be adequate to just *must* be a consistent one, then conditional consistency convergence is a reasonable methodological constraint to impose on inquiries, and one would look next at the various detailed ways of formulating this requirement in terms of practically describably accessibility relations from one non-standard world to another (e.g. stipulating that no world $w'$ which is inconsistent in some respect in which world $w$ is consistent be methodologically accessible from $w$, and so on). Our current interests are more abstract than this project, however. We are considering various kinds of methodological convergence on the lattice of non-standard possible worlds in order to elucidate the ontological and epistemological status of non-standard possible worlds.

One worry which our use of the notion of convergence might raise is the following. It might seem that treating various kinds of convergence as methodologically meritorious is a concession to *conservatism*, which by ruling conceptual revolutions out of court at the outset seeks to guarantee that an inquiry will not arrive at conclusions radically different from those beliefs which occur early in the inquiry. If this objection were correct, it would be a serious criticism of our account. But in fact the global constraint of convergence does *not* rule out locally discontinuous, radical, or revolutionary changes of belief. Specifically, in our mathematical model (as in the special case of numerical sequences) convergence does not mean that if $m < n$ that the $n+1^{st}$ element must be "closer to" the $n^{th}$ element than the $m+1^{st}$ is to the $m^{th}$ (where the lattice ordering determines the sense of "nearer"). Nor is there any pair of worlds $w_1, w_2$ (or, again, numbers, in the special case) which are too "far apart" for $w_1$ to be the $n^{th}$ and $w_2$ the $n+1^{st}$ elements of a convergent sequence of worlds (numbers). This is clear from the consideration of the trivial case of a sequence which is constant after a certain point. So there is no change of belief which is too radical or revolutionary to be contained in a convergent inquiry, and convergence is not the counsel of conservatism. More precisely, a methodology may br called *conservative* insofar as it attributes a special positive weight, in its evaluation of possible transformations of belief, to *retaining* beliefs currently held, or which have remained unchanged over the recent course of inquiry. Conservatism thus described is a constraint on methodologies at the same level as our various convergence conditions, and the present point is just that it is a constraint which is completely independent of considerations of convergence.

Having pointed this out, however, it might seem that we have opened the door to a more serious criticism. If the requirement of global convergence excludes *nothing* at the local level (forbids no particular transition), then isn't it a *vacuous* constraint on methodology, which after all can only operate at the local level, governing individual changes of belief via an accessibility relation? The requirement is *not* empty, as we see when we distinguish the *two* levels of constraint which are relevant here. First, we have constraint of an inquiry (path) by a methodology (accessibility relation determining admissible paths). Second, we have the constraint of methodologies by meta-theoretic principles like unique universal convergence. That permission for any particular transition at the first level is consistent with any particular prohibition (at least of the well-behaved ones we have considered) at the second level does *not* mean that second-level constraints *on methodologies* aren't genuine. To be genuine is to exclude *some* methodologies, partition the infinite space of possible methodological accessibility re.ations into those which permit and



those which do not permit paths prohibited by a meta-level principle (such as the requirement that all pairs of paths with a common element converge to the same world). In fact a requirement such as that just stated clearly *is* a significant constraint on methodologies. The *apparent* vacuity of these meta-level principles arises only from the confusion of the two different sorts of constraint on inquiry which are at issue here.

Finally, we may consider briefly the complementary objection, that our talk of convergence of cognitive inquiries is utopian and hopelessly optimistic, since actual inquiries cannot be expected to conform to such an expectation. (Such a view might be motivated either by a low opinion of current cognitive practice or contemporary methodological sophistication, or by pessimism about the influence of the latter on the former.) To this we respond that methodology is primarily a matter of how we *appraise* inquiries, and such appraisals may be related to the actual course of previous inquiries only in historically and sociologically complex ways. Future inquiries can be affected by our methodologies, but only via the appraisals of possible course of inquiry which we make according to *them*. The objection would have a point were we making specific methodological recommendations, but is without force at the level of meta-theoretic conditions which might be imposed on methodologies, which, after all, is where our concerns lie.

# SECTION 24

## The Example of Peirce

It will be useful to consider a particular example of the general approaches we have presented abstractly in the terminology of idealism and realism. Accordingly, we shall, in this section, examine briefly the views of the early Peirce on methodology and ontology. This discussion serves a dual purpose. On the one hand, we will see several important aspects of Peirce's views which have seemed mysterious to those approaching his work with standardist prejudices, and which become clear and persuasive when viewed from the perspective of the apparatus of non-standard possible worlds. On the other hand, the consideration of Peirce's sophisticated views about convergence, inquiry, and reality provides a concrete setting in which to argue the thesis of the ontological parity of standard and non-standard possible worlds which is our ultimate aim.

Peirce is a methodological realist, but an ontological idealist. One of his more important legacies to later pragmatists (taken up particularly by Dewey) was his doctrine that the objects of knowledge inhabiting the limit-world approached by a convergent methodologically admissible inquiry are the intelligible *products* of such disciplined inquiry, not its ineffable *causes*.[95] Peirce's Kantianism consists precisely in this claim that the methodologically real is ontologically ideal.[96] This combination of views is often obscure, since Peirce is not careful to distinguish the two senses in which "real" is used. Because it is precisely in the interface between the methodological realism and the ontological idealism that Peirce's views about inconsistency and incompleteness lie, it is worth looking more closely at this region.

Peirce does not talk about inconsistency and incompleteness in those terms, but in terms of *vagueness* and *generality*. The following passage is characteristic:

> Perhaps a more scientific pair of definitions would be that anything is *general* in so far as the principle of excluded middle does not apply to it and is *vague* in so far as the principle of contradiction does not apply to it. (5.448)

Peirce called anything which is vague or general a *sign*, and claimed that ideas were signs in this sense.[97] His ontological idealism thus consisted of the claim that the signs in this sense which make up our



beliefs at a time are what actually exist. In fact, Peirce has a more radical view still. He identifies a human being with his beliefs and the transformation of them over time which is inquiry.

> . . . the fact that every thought is a sign, taken in conjunction with the fact that life is a train of thought, proves that man is a sign . . . (5.314)

Man is thus the inhabitant of an inconsistent and incomplete world, a world of *signs* in Peirce's special usage. Since signs are coordinate with ideas, this view was entitled, Peirce felt, to be called an idealism. But idealism in this sense is perfectly consistent with methodological realism, the view that the controlled settlement of opinion (to which we have given precise expression in terms of the resolution convergence of methodologically admissible inquiries) is the criterion of truth. Thus Peirce wrote:

> When we busy ourselves to find the answer to a question, we are going upon the hope that there is an answer, which can be called *the* answer, that is, the final answer. It may be that there is none. If any profound and learned member of the German Shakespearian Society were to start the inquiry how long since Polonius had had his hair cut at the time of his death, perhaps the only reply that could be made would be that Polonius was nothing but a creature of Shakespeare's brain, and that Shakespeare never thought of the point raised. Now it is certainly conceivable that this world which we call the real world is not perfectly real but that there are things similarly indeterminate. We cannot be sure that it is not so. In reference, however, to the particular question which we at any time have in hand, we hope there is an answer, or something pretty close to an answer, which sufficient inquiry will compel us to accept. (4.61)

Here Peirce is envisaging methodological convergence to a non-standard world. Not only is every stage of inquiry a world of signs, fraught with vagueness and generality (that is, in his somewhat confusing usage, *ideals* or signs rather than *reals*) but so may well be the limit world which they seek to be adequate to, to find out about. The doctrine that this limit-world, consisting of the objects of knowledge *produced* by methodologically admissible convergent inquiry, contains vagues and generals (quasi-objects in our technical sense) is what Peirce calls Scholastic Realism. It is the view that inconsistent and incomplete objects are methodologically real.[98] Since man is such an object (or quasi-object) according to Peirce, this is also the question of whether *man* is methodologically real, i.e. whether he exists in the limit-world as an object of knowledge to which appropriate inquiries will converge.[99]

It is clear, then, that an ontological idealist like Peirce can tolerate the methodological reality of non-standard worlds and their constituents. If reality in this sense is ultimately *defined* by convergence of methodologically admissible inquiries, there seems no particularly compelling motive for requiring *standardness* convergence (i.e., convergence to a *standard* world). The ontological realist, on the other hand, does seem to have the intuition that the (methodologically and ontologically) real world just *must* be standard, that standardness is a necessary condition of actual existence, as opposed to the merely virtual status of inconsistent and incomplete objects. Peirce acknowledges such an intuition in the passage quoted above when he describes the situation in which inquiry converges to a non-standard world as one in which the (methodologically) real world is not perfectly real (i.e., standard). This is an intuition of the greatest importance for our overall project.

## SECTION 25

# A Critique of Standardism

Our purpose throughout has been to set out a framework to justify giving inconsistency and incompleteness an appropriate place within the arena of rational inquiry. The ideas we have considered are intended to show that inconsistency and incompleteness—and even the ambiguous identity relations of quasi-objects which such situations include—have a perfectly intelligible and rigorously formalizable structure. We have urged that non-standard possible worlds should be accorded exactly the same ontological status which the standard, consistent and complete worlds enjoy. (After all, from one point of view, those standard worlds are merely a special, degenerate case of non-standard worlds.) Thus we want to maintain the ontological *parity* of standard and non-standard worlds—the claim that *whatever* ontological status one accords to standard, consistent and complete possible worlds must be accorded as well to their non-standard brethren.[100]

In particular, we are concerned to raise the possibility that standardness convergence ought *not* to be taken as an absolute methodological constraint on all serious inquiry—that, on the contrary, inconsistency and incompleteness are tolerable and manageable in the progressive transformation of belief and practice which is inquiry. One should not in other words, consider convergence to a non-standard world as the hallmark of cognitive defeat, but rather must prepare for the possibility that *our* world—the real world, in which we live and breathe and have our being—is non-standard. The ontological realist and the ontological idealist differ, as we have seen, in their interpretation of which world *our* world is—whether it is the belief-world which represents the current stage of inquiry, or the limit-world to which that inquiry will (hopefully) converge eventually.

Our concern is with the relation of this dispute to ontological *standardism*, the view that all that exists, actually *or possibly*, is consistent and complete objects and the worlds comprising them. Ontological standardism depends upon maintaining a disparity of status, roughly, between beliefs and the things those beliefs are *about*. The former are allowed to be inconsistent and incomplete, the latter precluded from being so. More precisely, standardism claims that inconsistency and incompleteness are characteristics always of

# APPENDIX V

## Convergent Nets and Paths on a Lattice

The definitions of this appendix are all standard mathematical fare. See for instance G. Birkhoff's *Lattice Theory* (Volume XXV American Mathematical Society Series 1967, pp. 244ff.)

1. A lattice is *complete* iff in addition to satisfying the lattice axioms (L1)–(L4) stated in the text, every subset of the lattice has a meet and a join. In our case, this means that superposition and schematization must apply to arbitrary sets of worlds, not just countable sets. Our motivations and detailed implementations allow this, so we will take advantage of it.

2. An *upper-directed* set is a set $I$ ordered by $=$ such that for any $i, j \in I$, there is a $k$ such that $i \leqslant k$ and $j \leqslant k$.

3. A *net* is a set indexed by an upper-directed set. If the upper directed set is totally ordered, we may call it a *path*. In our applications, paths will be employed rather than the more general nets, and the ordering principle of their upper-directed index sets is taken to be *temporal*, so that progress along a path in a lattice corresponds to ever *later* stages in some inquiry. Our present interest lies in exploiting the formal relations between such an ordering principle and the order codified in the lattice structure induced by superposition and schematization.

4. That lattice-order, symbolized by $\leqslant$, is defined by $w_1 \leqslant w_2$ if and only if either $w_1 \cup w_2 = w_2$ or $w_1 \cap w_2 = w_1$. This induces a partial ordering on the lattice of non-standard possible worlds.

5. For any subset $X$ of a partially ordered set, an *upper bound* of $X$ is an element $a$ such that for all $x \in X$, $x \leqslant a$. A *least* upper bound is an upper bound $a$ of $X$ such that for any other upper bound $b$ of $X$, $b \leqslant a$. The greatest lower bound of $X$ is defined dually. We write the least upper bound of $X$ as *sup* $X$, and its greatest lower bound as *inf* $X$.

6. Given a net $\{x_i/i \in I$ an upper-indexed set$\}$, we define $\mathrm{Lim}(\inf\{x_i\}) =_{\mathrm{df}}$. Sup $\leqslant \{y_j/y_j = \inf \leqslant \{x_i/i \geqslant j\}\}$, where sup $\leqslant$ indicates that it is the lattice ordering, not the index ordering, with respect to which we consider bounds. Similarly, we define $\mathrm{Lim}(\sup\{x_i\}) =_{\mathrm{df}}$. Inf. $\leqslant \{y_j/y_j = \sup. \leqslant \{x_i/i \geqslant j\}\}$.



7. With these definitions in hand, we can say that a net $\{x_i/i \in I\}$ *converges to a* on a complete lattice (sometimes this is called *order convergence*) iff:

$$\mathrm{Lim}(\inf\{x_i\}) = \mathrm{Lim}(\sup\{x_i\}) = a$$

This is motivated by the familiar fact that for the real numbers the convergence of a sequence $\{x_i\}$ to the limit $a$ is equivalent to the condition that $\lim(\sup\{x_i\}) = \lim(\inf\{x_i\}) = a$. Our definitions simply generalize convergence as we understand it on real numbers, with the index ordering being induced by the ordinality of various positions in the sequence of real numbers, and the analogue of our lattice ordering $\leqslant$ coming from the intrinsic ordering of the reals.

We may utilize these formal definitions so as to be able to define inquiry as a path from non-standard belief-world to non-standard belief-world, totally ordered by the relation "arrived at *after*", and thus be able to discuss the *convergence* of such indexed inquiries on the lattice of non-standard possible worlds with *its* ordering principle induced by the behavior of schematization and superposition.