What Are Possible Worlds?

JOHN E. NOLT

My question concerns possible worlds. By 'possible worlds' I mean, quite literally, worlds. Worlds are universes. The most interesting of them are spatiotemporal manifolds in which people live, time passes, and events unfold. Of these one, the actual world, is especially important to us.

What I do not mean by 'possible worlds' are maximal consistent sets of propositions, maximal states of affairs construed as actual but abstract objects, set-theoretic models, or similar mathematical or Platonic structures.[1] Frankly, I don't believe in such things. But whether they exist or not, they do not concern me here.

What do concern me here, however, are possible *worlds*. I don't believe in possible worlds, either—except of course for the actual one. That is, I don't believe they exist. But I do contend that they are possible and that understanding their possibility is a way of understanding what they are. Such is the thesis of this paper.

I hold the same to be true *mutatis mutandis* if for 'world' we substitute such terms as 'situation', 'scenario' or 'state of affairs'. (All of these suggest entities "smaller than" worlds.) Merely possible situations, scenarios, and states of affairs do not exist, but they are possible, and understanding their possibility is a way of understanding what they are. These "smaller" entities, however, will be mentioned only briefly and incidentally in the course of this paper.

Crucial to the thesis just stated is a distinction between existence and possibility. Consider the statement

(1) It is possible for spacetime to be Newtonian

and its equivalent

(2) A world in which spacetime is Newtonian is possible.

Both are true, provided that the term 'possible' is read in a logical or conceptual sense. (This is the only sense in which 'possible' will be used in this paper.) But many of us have serious reservations about the truth of

(3) A world in which spacetime is Newtonian exists.

[1] For an example of a theory that takes possible worlds to be actual but abstract states of affairs, see Alvin Plantinga, 'Actualism and Possible Worlds', *Theoria*, 1976, 139–60; reprinted in Michael J. Loux, ed., *The Possible and the Actual: Readings in the Metaphysics of Modality*, Ithaca and London, Cornell University Press, 1979. According to Plantinga, all possible states of affairs exist, but most of them do not obtain. I criticized this view in 'Sets and Possible Worlds', *Philosophical Studies*, 1983, 21–35.

(3) says more than either (1) or (2). A sufficient condition for the truth of (1) and (2) is that we can assert without self-contradiction that spacetime is Newtonian. But this is hardly sufficient for (3). (3) says that a Newtonian world exists, and we are reading 'world' literally. The mere consistency of an assertion does not guarantee the existence of a universe in which that assertion is true. Consequently, the mere possibility of a world is something quite distinct from its existence.

It might be objected that since consistency is simply existence of a model, the consistency of a statement *does* entail the existence of a "world" in which it is true.[2] But this objection confuses models with worlds. Newtonian spacetime can be modelled in set theory and perhaps even in some physical systems, but such models are not literally Newtonian universes. Nor does their existence entail the existence of any Newtonian universe. Yet the consistency of a statement does entail the logical possibility of a world in which the statement is true. Hence the conclusion stands: the mere possibility of a world is not the same thing as its existence.

But what does it mean to say that something exists? In plain English 'exists' and 'is actual' are synonymous. Hobbits, for example, do not exist, because they are not actual. Only confusion can result from the rejoinder that they exist as possible hobbits, as if 'possible' were a predicate satisfied simply by some subset of the domain of existing things. On the contrary, the totality of what is possible vastly exceeds the totality of what exists in the ordinary sense of the term. Thus, if we treat 'exists' as a quantifier in the usual way and regard 'is possible' as a predicate, the domain of the quantifier is too narrow to accommodate the full extension of the predicate. If we now expand the domain of quantification to include merely possible objects, we can no longer consistently interpret the quantifier as 'exists', for not all the objects in that domain exist. Rather, the existential quantifier now means 'there is possible' or 'there could be'. Similarly, the universal quantifier does not mean 'for all', where 'all' means 'all that exists'; it means 'for all possible'. Thus, 'is possible' functions not as a predicate, to be applied within the domain of existing things, but as a quantifier whose domain subsumes the domain of existing things.[3]

We are right, then, to have reservations about (3); for either it is straightforwardly false (if 'exists' means 'exists') or it is just a misleading formulation of (2), in which case it would be better to stick with (2) itself.

Nevertheless, some philosophers hold statements like (3) to be true.[4]

[2] Actually, on my view, consistency is not the existence of a model but the possibility of a model. The set-theoretic structures we invoke to prove consistency are in general merely possible, not actual, objects. This view is elaborated in 'Sets and Possible Worlds', Nolt, op. cit.; 'Mathematical Intuition', *Philosophy and Phenomenological Research*, 1983, 189–211; and 'Abstraction and Modality', *Philosophical Studies*, 1980, 111–27.

[3] Of course, once such quantifiers are introduced, there can be no objection to having a predicate meaning 'is possible'. But it is of little use, since it is true of every object in the domain of quantification.

[4] Most notably, David Lewis in *Counterfactuals*, Cambridge, Mass., Harvard University Press, 1973, chapter 4, and *On the Plurality of Worlds*, Oxford, Basil Blackwell, 1986.

The reason is simple. There are occasions on which it is both useful and natural to quantify over possible worlds. This is so, for example, when we give truth conditions for the operator '◇':

(4) ◇p is true iff $\exists x$(x is a (possible) world & p is true in x).

Now these philosophers, for various reasons, do not want to abandon such quantification. And they assume that quantification entails commitment to the existence of the values of the quantified variables. Hence they conclude that they are committed to the existence of possible worlds.

I agree that quantification of this form should not be abandoned. But I do not think it commits us to the values of the variables. Just writing the backwards 'E', of course, does not commit one to the existence of anything, since it is simply a syntactic form. Only if it is interpreted to mean 'there exists' does it bear an ontological burden. But I have argued that when used to quantify over possibilia, it should be read not as 'there exists', but as 'there is possible' or 'there could be'. Thus the appropriate English rendering of (4) is

(5) ◇p is true iff there could be x such that x is a world and p is true in x,

not

(6) ◇p is true iff there exists x such that x is a possible world and p is true in x.

Likewise, the appropriate English rendering of

(7) $\exists x$(x is a (possible) world & spacetime is Newtonian in x)

is (2), not (3).

This is not to say that when '∃' is used to quantify over possible worlds it should be read as '◇∃'. The '∃' in '◇∃' ranges not over worlds, but over the domains associated with them, and '◇' is in effect an existential quantifier over worlds. But on the interpretation of '∃' that interests me here, '∃' ranges straightforwardly and objectually over worlds themselves; yet it asserts only possibility, not existence, and hence carries no existential commitment. Such an interpretation can be seen to be perfectly intelligible, provided that we can make sense of a simple game of make-believe.

To describe this game, we must make some distinctions. We shall need to consider both those quantifiers that assert existence and those quantifiers that do not. Let us call the existential quantifier that asserts existence, together with its universal dual, *standard quantifiers*. Standard quantifiers are interpreted objectually over domains of existing objects. And let us call the "existential" quantifier over worlds, together with its universal dual, *possibilistic quantifiers*. Possibilistic quantifiers are to be interpreted objectually over a domain of worlds, not all of which exist.

Let us further distinguish two languages, which for our purposes need be

characterized only roughly. The first is the language of possibilistic quantification, or *possibilistic language*, whose predicates are predicates of worlds. The second is a standardly quantified language, which we may take to be the set of declarative statements of English, except for those that contain modal operators, possibilistic quantifiers, and other modal locutions. To avoid irrelevancies involving semantic paradox, we shall also omit statements containing such expressions as 'true' and 'satisfies'. For the sake of clarity, we shall sometimes write the sentences of this language semiformally, using regimented notation for quantifiers and connectives. The predicates of this standardly quantified language are typically predicates of individuals, not of worlds. We shall call this language the *descriptive language*.

Now, the only thing that matters in determining whether a world of a certain description is (logically) possible is whether the description itself is consistent. More precisely, what counts as a logically possible world is governed in part by the following rules:

(i) The actual world counts as a possible world.
(ii) If A is a consistent statement (or consistent set of statements) of the descriptive language, then there counts as possible at least one world in which (each member of) A is true.
(iii) If A is an inconsistent statement (or set of statements) of the descriptive language, then no world in which A is true counts as possible.

It is evident that talk about logically possible worlds presupposes these rules. I have used the locution 'counts as possible' to emphasize that rules (i)–(iii) are prescriptive, not descriptive. They are to be read, not as a definition that picks out from among already given worlds those that are possible, but as instructions for playing a game of "make-believe". It is this game that we must understand in order to see what possible worlds are.

An immediate consequence of (iii) is:

(iii′) If A and B are consistent statements (or consistent sets of statements) of the descriptive language whose conjunction (union) is inconsistent, then any world in which (each member of) A is true is distinct from any world in which (each member of) B is true.

The effect of rules (i)–(iii) is to "set up" the domain of quantification for the possibilistic language. To some extent we can tell what is included in this domain and what is excluded from it by determining which statements of the descriptive language are consistent.

Of course, rules (i)–(iii) are not the only principles governing the (logical) possibility of worlds. Certainly, if we are talking about worlds (as opposed to "smaller" entities like situations, scenarios, or states of affairs) we will need a maximality condition—something like this: for every possible world w and

every statement A of the descriptive language, either A or its negation is true in w.

Yet even with the addition of this condition, the totality of possible worlds is not fully defined, for there is nothing to prevent many worlds from corresponding to the same maximal consistent set of statements of the descriptive language.[5] That, I will argue later, is as it should be.

There is, however, another respect in which rules (i)–(iii) fail to provide a completely rigorous characterization of the domain of possible worlds, and that is their use of the term 'consistent'. The descriptive language, we have said, is a fragment of English. But the notion of consistency is well-defined only for certain formalized languages. Therefore, the notion of consistency used in rules (i)–(iii) must be to some degree vague and informal.

Where we draw its boundaries is largely a matter of convenience, not of logic or of fact. Minimally, it will exclude formal contradictions (i.e. those statements which cannot be true in virtue of the semantic rules governing the standard logical operators). But we may also want it to exclude more.

We may want it, for example, to exclude arithmetical falsehoods. Consistency, then, will also be a function of the semantics of such terms as '$+$' and '$\times$'. Such terms will have for us a fixed interpretation (their standard interpretation over the natural numbers), just as the usual logical operators do, and any statement false on this interpretation will thereby be regarded as inconsistent.[6] Our concept of consistency will now be *essentially* informal, since the set of arithmetical falsehoods cannot be formally (i.e. recursively) specified.

We may also want our concept of consistency to exclude statements like

(8) $\exists x(x$ is round & x is square),

or

(9) $\exists x(x$ is taller than $x)$,

which, though not formally contradictory, are nevertheless false on any interpretation in which their predicates have their usual senses. Their inconsistency, in other words, is a function of the semantics of both logical operators and predicates.

We may even wish to exclude statements whose falsehood is not a function of semantics alone. Consider, for example, the statement made by the sentence

(10) This substance does not have atomic number 79,

[5] A maximal consistent set of statements for a language L is a consistent set of statements of L which contains for each statement A of L either A or its negation.

[6] Numbers are generally regarded as abstract entities, and since I announced at the beginning of this paper that I do not believe in such things, I may seem to have contradicted myself. For the resolution of this apparent conflict, see the works cited in n. 2. In short, the resolution is this: talk of numbers is really possibilistic talk about familiar, concrete things.

uttered with reference to a lump of gold (whose atomic number is 79). Since having atomic number 79 is arguably essential to the substance ostended, we may wish to regard this statement as inconsistent in an informal sense.[7] Not, however, that the same sentence uttered with reference to a hunk of lead is not only consistent, but true. The semantic contribution of the sentence is the same in each case; what differs is the context. Thus, if we regard (10) as inconsistent, we will conclude that context, too, may play a role in consistency.

The context dependence of the truth of (10) has its source in the demonstrative 'this substance'. The semantic principle governing this phrase fixes its reference to whatever substance is indicated by the context. But it does not stipulate which substance this is, nor what properties are definitive of (i.e. essential to) this substance. These things can only be known empirically. If the substance ostended is gold, one of the essential properties (given current scientific doctrine) is the property of having atomic number 79. It is only by "completing" the interpretation of the phrase with these context-dependent considerations that we can recognize the inconsistency of the statement in question. Thus, if we admit inconsistency in this sense, then inconsistency is not always detectable a priori. Where semantic interpretation is context-dependent, it may depend on aspects of context which can only be known empirically.

It does not follow, of course, that statements that are actually consistent might not have been consistent if the empirical facts had been different. We have seen that when (10) is uttered in one context (with reference to a hunk of lead) we get consistency and when it is uttered in another (with reference to a lump of gold) we get inconsistency. But this does not mean that the consistency of a statement is subject to the whim of empirical fact. Rather, it shows that statement identity may depend on empirical fact. What we have here is not a single statement that is consistent in one context and inconsistent in another, but rather two statements that are distinguished by context. Changing the context replaces one statement by another. It does not make a formerly consistent statement inconsistent.[8]

My aim in discussing the various notions of consistency has been merely to sketch the range of options available to us in interpreting rules (i)–(iii). For the purposes of this paper we need not choose among these options. In

[7] For an account of the essentialist intuitions on which this claim of inconsistency might be based, see Saul Kripke, *Naming and Necessity*, Cambridge, Mass., Harvard University Press, 1972, esp. pp. 123–5. Kripke himself, however, does not say that statements like (10) are inconsistent.

[8] We might suspect, however, that contingent fact affects consistency in a more subtle way. Perhaps one day empirical discoveries will force us to revise our conception of a physical substance so that atomic number is no longer a criterion of substance identity. Then we might no longer see any inconsistency in asserting with reference to a lump of gold that *this substance* does not have atomic number 79. Some big questions lurk in the shadows of this suspicion—including a very big one about reidentification of statements, and ultimately of possible worlds, across differing conceptual schemes. But this is not the place to draw them out.

practice, the choice may be made differently for different purposes.[9] No matter which choice is made, however, consistency is determined at least in part by semantic rules, and perhaps also in part by those aspects of context needed to fix an interpretation. If these factors do not rule out the truth of a statement, then that statement is consistent. If they do, it is inconsistent.

Since we lack a general account of semantic rules, this account of consistency may seem unconscionably vague. But in specific instances in which the relevant rules are thoroughly understood, it is self-evidently correct. Consider, for example, the inconsistency of statements of the form 'p & $\neg p$'. This inconsistency is a product of the semantic rules governing the operators '$\neg$' and '&'. The rules are: '$\neg$' reverses truth value; '&' forms compounds that are true only if both components are true. Thus, it is perfectly evident that these rules prohibit the truth of any statement of the form 'p & $\neg p$'. Because the semantic rules governing '$\neg$' and '&' are thoroughly understood, the nature of the inconsistency is utterly transparent.

It may be objected, however, that with respect to the more informal notions of consistency this sort of explanation is ultimately circular. I set out to explain what possible worlds are in terms of consistency. Consistency is a function of semantic rules and (perhaps) context. But how are these rules to be spelled out, if not with reference to possible worlds? The inconsistency of (8), for example, is a result of the senses of 'is round' and 'is square'. But the most obvious way to represent the semantic rules that constitute these senses is as functions from possible worlds to extensions. Hence the explanation runs in a circle.

There *is* a circle here, but it is not vicious. Of all the elements of this circle, the semantic rules which constitute senses are the most familiar. The rules governing expressions like 'is round' and 'is square', however, are familiar, not in the sense that we can formulate them fully and precisely, but in the sense that we know how to use them. That is how we detect consistency and inconsistency in practice.

Since explanation should proceed from the more familiar to the less

[9] Since consistency is linked by rules (i)–(iii) to logical possibility (and thereby to related notions such as validity, necessary truth and logical equivalence), choosing any but the minimal notion of consistency will, in effect, broaden our conception of logic. If we treat (8) as inconsistent, for example, then we shall treat the inference

$$x \text{ is round}$$
$$\therefore x \text{ is not square}$$

as deductively valid even without the additional premiss

$$\forall x(x \text{ is round} \rightarrow x \text{ is not square})$$

that would be needed to make it valid when transcribed into the predicate calculus. The result, if this idea is taken seriously, is an *informal* logic, which (I have argued elsewhere) is more useful for some purposes than standard formal logics. See my 'Possible Worlds and Imagination in Informal Logic', *Informal Logic*, 1984, 14–17 and *Informal Logic: Possible Worlds and Imagination*, New York, McGraw-Hill, 1984.

familiar, these practically and intuitively understood rules are the appropriate starting point. The question 'Why can't there be worlds in which statements of the form 'p & $\neg p$' are true?' is correctly and enlighteningly answered by explaining the semantics of '$\neg$' and '&'. Similarly, the question 'Why can't there be worlds containing round squares?' is correctly and enlighteningly answered by appealing to our pre-theoretical understanding of the way 'is round' and 'is square' work. The latter explanation may utilize imprecise intuitions, but it does not involve possible worlds. If we reverse the direction of explanation, however, the explanation falls flat. The question 'Why do the semantic rules governing '$\neg$' and '&' work the way they do?' is not enlighteningly answered by appeal to the fact that there can be no worlds in which 'p & $\neg p$' is true. Nor do we enhance our understanding of how the senses of 'is round' and 'is square' make (8) inconsistent by pointing out that there can be no worlds containing round squares. The semantic rules constituting senses are thus conceptually prior to possible worlds. (It is not wholly irrelevant in this regard that consistency and inconsistency were recognized and understood to a considerable degree at least two millennia before Leibniz popularized the notion of a possible world.)

It is by appeal to our practical and intuitive understanding of semantic rules that I hope to shed light on the question of what possible worlds are. Perhaps on other occasions we will find it useful to represent these rules as functions from worlds to extensions. Still, that will not alter the fact that they were initially understood independently of the worlds concept. The movement, then, will have been from a working understanding of semantic rules, to a working understanding of possible worlds, back to a theoretical representation of the rules with which we started. If this theoretical representation were the only way of understanding semantic rules, then the circle would be vicious. But since our initial understanding is independent of this theoretical understanding, the circle is not vicious, but hermeneutic.

Having laid the groundwork, let us now return to the question 'What are possible worlds?' Two answers suggest themselves immediately: (1) possible worlds are possible states of affairs that are maximal in some sense, and (2) possible worlds are maximal consistent sets of statements of the descriptive language.

The first answer I take to be true but unilluminating. It needs to be augmented by an account of possible states of affairs. This account will either regard possible states of affairs as genuine possibilia, components of worlds (universes) that are possible but (except in one case) do not exist, or it will not. If not, then it will take them to be actual abstract entities of some sort. In that case it is not an answer to my question. For no worlds other than our own are actual; hence these actual abstract entities are not worlds. But my question was about worlds.

On the other hand, if possible states of affairs are regarded as genuine

possibilia, then I accept (1). But I do not think it is particularly useful to analyse possible worlds in terms of possible states of affairs. Rather, I want to elucidate both concepts from a totally different perspective.

Answer (2) proposes an ontological reduction of worlds to sets of statements. But, as I have already indicated, this sort of answer will not do. For sets of statements, like other sorts of actual abstract entities, are not worlds.

At this point, those who favour (2) or some other sort of ontological reduction are likely to reply that I have missed the point. Reduction establishes identity; possible worlds are just these sets, and hence they are the very things that interest me.

That, however, is not true. The things that interest me have properties (being a universe in which people live and breathe, for example) that sets of statements and similar abstract entities lack. There is no identity here.

Other proponents of reduction might reply that reduction means elimination. The (alleged) fact that we can replace possible worlds talk with talk of sets of statements (or similar abstract entities) shows, they may hold, that the non-actual worlds that interest me do not exist.

But I have already granted that; we are in agreement.

This odd twist of the dialectic points to a deeper *dis*agreement. The reductivist sees the non-existence of merely possible worlds as an *objection* to possibilistic discourse. This is because the reductivist's paradigm of language use is theory or description of what actually exists. Thus, when a form of discourse concerns itself with non-existent things, it strikes him as pointless and illegitimate.

I, on the other hand, am suggesting that the paradigm by which possibilistic discourse is understood should be games of make-believe. From this perspective, the non-existence of merely possible worlds is not an objection, but precisely what one would expect. Possibilistic quantification is quantification over a mostly make-believe domain. But that does not make it pointless. Since the domain is structured by rules (i)–(iii), possibilistic discourse conveys quite definite information. This information has two aspects; it concerns both statements and worlds. Consider, for example, the assertion

(11) There could be a world in which human civilization is destroyed by nuclear war.

This assertion conveys the information that the statement

(12) Human civilization is destroyed by nuclear war

is consistent. But that is not all it conveys—nor, perhaps, what is most important. For the possibilistic language game is played by taking (11) to be referring, not to (12), but to a world in which (12) is true. Of course that world may be make-believe. But then again it may not. What endows possibilistic discourse with utility (and sometimes poignancy) is that one

world in its domain is not make-believe—and (since we lack a complete description of the actual world) in a sense we do not know which one. It is, therefore, of more than just semantic concern that there could be a world in which human civilization is destroyed by nuclear war.

Yet those for whom theory or description are the paradigms of language use are likely to remain sceptical of the legitimacy of possibilistic discourse. Someone might argue as follows. Let the peculiarities of this sort of discourse be granted. Still, *something* must make possibilistically quantified statements true or false, and the theory of this something is their semantic metatheory. Now this metatheory *is* a theory, and as such it must live up to its ontological commitments; it must therefore provide truth conditions for possibilistically quantified statements solely in terms of existing things. Without such a metatheory, possibilistic quantification is unintelligible.

This appears to be a forceful objection. Indeed, I am unable to provide the semantic metatheory that it demands. (One could, of course, give truth conditions in terms of maximal consistent sets of statements. But not only would this be the wrong semantics for talk about *worlds*; it would also be inadequate, since the number of worlds may exceed the number of maximal consistent sets of statements—a consideration to which I will return shortly.) This does not mean that possibilistic quantifiers lack truth conditions. In fact, their truth conditions are trivial. In the simplest case, where ϕ is an open sentence with only x free, they are just this:

$\exists x\phi$ is true iff there could be a world satisfying ϕ,

$\forall x\phi$ is true iff every possible world satisfies ϕ.

The problem is that possibilistic quantifiers reappear in the metalanguage. The metalanguage requires us to engage in the same sort of make-believe that the object language employs.

The objection claims that this is not enough. It demands that truth conditions be formulated solely in terms of actually existing things. Now, a purely actualistic metatheory may well be desirable. But if we cannot have it, must we really concede, as the objection demands, that possibilistic quantification is unintelligible? What is the argument?

One might suspect a vicious circle here. To understand the possibilistic language, we must first understand its truth conditions. Hence, if its truth conditions can only be given in possibilistic language, we cannot come to understand it without first having understood it. Therefore we cannot come to understand it at all; it is unintelligible.

But this reasoning rests on a false assumption. It is not true that to understand a language we must first understand its truth conditions. Long before anyone had explained the truth conditions of standard quantification to us, we understood standard quantification and used it correctly. Indeed, we could not have understood its truth conditions until we had first learned

to use it, since those truth conditions employ standard quantification in precisely the way in which the truth conditions for possibilistic quantifiers employ possibilistic quantification. We first understood standard quantification, not by being given its truth conditions, but by observing it in use.

In the preceding pages I have explained and exemplified the use of possibilistic quantifiers. In particular, I have tried to show how, in accordance with rules (i)–(iii), patterns of consistency in the descriptive language are taken to guarantee the possibility of corresponding worlds and hence to "set up" the domain of possibilistic quantification. To understand these functions of language simply *is* to understand what possible worlds are, for the possibility of worlds (note that I do not say 'their existence') is constituted solely by the absence of inconsistency from the corresponding descriptions. (The actual world is no exception, since its inclusion in the domain via rule (i) makes sense only on the presupposition that it is consistently describable.) Non-actual worlds do not "exist out there" as structured wholes independent of our linguistic inventions. That way of thinking about them confuses them with actual things and misconstrues talk about them as theory or description. Rather, their identity and structure as a domain of objects is the product of a game of "make-believe" played with statements of the descriptive language.

This is not to say that possible worlds are nothing more than language. They are worlds, as I have argued all along. But in saying that they are worlds, I am already using the possibilistic language. That is, I am already engaged in the sort of "make-believe" that we are trying to understand.

An analogy may be helpful here. Anyone who has read the works of J. R. R. Tolkien and who understands their use (story-telling) knows what a hobbit is. To such a person it is evident that hobbits do not somehow "exist out there" as fully defined structured wholes. Their "existence", such as it is, is nothing more than a make-believe projection of Tolkien's language. Yet this is not to say that hobbits are nothing more than language. On the contrary, they are small rational animals with woolly toes. But in saying this, I am already using the fictive language, i.e. engaging in the "make-believe", that we are trying to understand. (I am certainly not describing or theorizing in a way that involves ontological commitment; to think that I am is to misunderstand me completely.)

Possibilistic language is the apparatus of a game in which we speak of possible worlds as objects, even though (apart from the actual world) no such objects exist. Rules (i)–(iii) provide specific directions for play. As we noted, however, these rules, even together with a maximality condition, do not completely define the domain of possible worlds. Though they do imply that to each maximal consistent set of statements of the descriptive language there corresponds *at least* one possible world, they do not imply that for each there is *only* one. They leave that question open. This, I contend, is not a deficiency of the rules, but a virtue. For the domain of possible worlds is

indefinite in just the way that these rules are. In other words, rules (i)–(iii), plus a maximality condition, define this domain as completely as it can be defined without imposing arbitrary constraints. To see this, note that the descriptive language may grow. Tomorrow it may contain more expressions than it does today. These expressions will enable us to formulate new pairs of consistent statements whose conjunction is inconsistent and hence (by rule (iii′)) to discriminate a multiplicity of worlds for each of those that we can discern today.

To illustrate this process of discrimination, let us consider a simplified example. Suppose our descriptive language contains only one name (say 'the Earth') and one predicate (say 'is spherical'), plus the logical operator for negation. In this language we can make only two non-equivalent statements: 'The Earth is spherical' and 'The Earth is not spherical'. These are mutually inconsistent, and so by rules (ii) and (iii′) *at least* two worlds are possible—one in which the Earth is spherical and one in which it is not. Now, suppose we add the predicate 'is cubical' to our language. The set containing the statements 'The Earth is spherical' and 'The Earth is cubical' is inconsistent; so by rule (iii) no world corresponding to this set is possible. But by rules (ii) and (iii′) at least three worlds are possible, one corresponding to each of the following three pairs of statements:

(A) 'The Earth is spherical', 'The Earth is not cubical'
(B) 'The Earth is not spherical', 'The Earth is cubical'
(C) 'The Earth is not spherical', 'The Earth is not cubical'.

The possible world(s) associated in our original language with the statement 'The Earth is spherical' is (are) clearly the same as the world(s) now associated with (A).But the world(s) originally associated with the statement 'The Earth is not spherical' has (have) been further articulated into at least two worlds, those associated with (B) and (C). Note, however, that these are not "new" possibilities. They are already implicit among the possibilities originally associated with the statement 'The Earth is not spherical'. Thus expansion of our language does not create new possibilities. It merely subdivides and articulates the old ones.

Now, there is no reason to think that language growth ever comes to an end. That is, there is no reason to believe in some ultimate descriptive language so rich that it could not be richer. (We could always, it seems, add one more predicate or one more name, just as we can always increase any ordinal number by one. Thus, the idea of an all-inclusive language seems no more credible than the idea of a greatest ordinal number.) If an all-inclusive descriptive language were possible, then each maximal consistent set of its statements would correspond to exactly one ultimate possible world. But if, as seems likely, there can be no such language, then possible worlds are in a sense infinitely divisible. At every given point in history we can distinguish as many of them as there are maximal consistent sets of statements in our

descriptive language. But later each of these may be resolved into many more. Any attempt to bound the domain of possible worlds by adding further conditions to rules (i)–(iii) and the maximality condition would arbitrarily limit this process of resolution. There is thus no generally applicable justification for adding such conditions.

The infinite divisibility of possible worlds may be abhorrent to classically tempered minds, but it does not matter much in practice. If we want the effect of a well-defined domain of worlds, we can always specify some precise descriptive language and then ignore differences among worlds that are not expressible in this language. As a result we will recognize only one world corresponding to each maximal consistent set of statements of this language. Technically, this result can be achieved by the device of identifying indiscernibles.[10]

But all this raises a final objection. Since non-actual worlds, on my view, derive whatever structure they have from patterns of coherence in the descriptive language, what would happen if there were no descriptive language at all? It would seem to follow that no non-actual world would be possible, which is surely absurd.

This objection, however, is mistaken. It invites us to consider a counterfactual situation in which there is no descriptive language and assumes that in doing so we are prevented from using the descriptive language we now possess. But that is not the way counterfactual discourse works. We always apply our current linguistic and conceptual resources to any possible situation we consider, and as a result counterfactual changes in languages have no effect on the structure of possible worlds *as determined by these linguistic and conceptual resources*. This practice implies that each world is accessible from all worlds, since the descriptive language and rules (i)–(iii) are held constant as we shift our viewpoint from world to world. Hence even in worlds where there is no language, all the worlds possible for us are still possible. Thus it is not true that if there were no descriptive language, no non-actual worlds would be possible.

Now, this answer may seem evasive. What the objection really wants us to do (and what I refused to do in the previous paragraph) is to imagine what happens to possible worlds from the point of view of a person who has no descriptive language. We are to consider this possible situation, not as outsiders looking in through the spectacles of our own language (which is the usual procedure), but empathetically, from the viewpoint of the languageless person himself. Now it ought to be clear that whatever we discover about this person's conception of possible worlds is irrelevant to the structure of possible worlds *from our point of view*. The possible worlds that matter to us are structured by our language games, not by his. Nevertheless, it might be illuminating to imagine how things seem to him.

[10] For an account of this device see W. V. O. Quine, 'Logic and the Reification of Universals', in *From a Logical Point of View*, New York, Harper & Row, 1961, 102–29.

To make this thought experiment easier, let us imagine that this person speaks our descriptive language initially, but then slowly loses it. His language, that is, can express fewer and fewer statements, until at last it can express none. The result with respect to his conception of possible worlds will be just the reverse of the progressive articulation that occurs with language growth. Worlds will seem to "merge", rather than to "fission". Previously distinguished worlds will become indistinguishable. Finally, no descriptive statements will be left. At this point, rules (i)-(iii), which "set up" the domain of possible worlds, imply neither the possibility nor the impossibility of any non-actual world. The only thing they still imply categorically is that the actual world is possible. The question of the possibility of other worlds is not decided in the negative; it is simply undecided. Thus it is in no sense true (even from this person's point of view) that nothing non-actual is possible. Rather, a precondition for distinguishing the possible from the actual is lacking. All previously distinguishable worlds (including the actual one) have from his viewpoint merged into undifferentiation.

Thus we can see that by enriching or impoverishing the descriptive language we can vary the degree of articulation of the domain of possible worlds. But in this process nothing becomes possible that was not possible before, and nothing ceases to be possible. Possible worlds or situations previously distinguished may cease to be distinguished. But possibility itself is neither created nor destroyed.

In this paper I have tried to explain what merely possible worlds (and also merely possible situations, scenarios, and states of affairs) are by offering an account of their possibility. I have not tried to do more than that. Thus my explanation leaves many questions unanswered. I have not, for example, addressed the question of how this game of make-believe is useful in decision-making and other practical activities. Obviously, part of the answer lies in our ignorance. Often we do not know which possibilities are actual and which are not; our decisions must therefore prepare us for non-actual as well as actual situations. But, of course, much more than this needs to be said. Nor have I discussed the problem of possibilistic quantification over individuals and the closely related problem of transworld identity. But these are issues for another context.[11]

The University of Tennessee,
Knoxville.

JOHN E. NOLT

[11] I would like to thank Steve Humphrey for the raucous conversations that stimulated my interest in this topic. This paper was extensively revised in the light of perceptive comments by Graeme Forbes, who undoubtedly still disagrees with most of it. Simon Blackburn also provided some very helpful criticisms.