

Models and reality*

Robert Stalnaker

Linguistics and Philosophy, Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

Kripke models, interpreted realistically, have difficulty making sense of the thesis that there might have existed things that do not in fact exist, since a Kripke model in which this thesis is true requires a model structure in which there are possible worlds with domains that contain things that do not exist. This paper argues that we can use Kripke models as representational devices that allow us to give a realistic interpretation of a modal language. The method of doing this is sketched, with the help of an analogy with a Galilean relativist theory of spatial properties and relations.

ARTICLE HISTORY Received 17 February 2016

KEYWORDS Actualism; contingentism; necessitism; possible worlds; Kripke models

The philosopher of 'possible worlds' must take care that his technical apparatus not push him to ask questions whose meaningfulness is not supported by our original intuitions of possibility that gave the apparatus its point.

Saul Kripke¹

1. Introduction

'Actualism' is an old label for a thesis in modal metaphysics. The general thesis of actualism is that to be real and to be actual are the same thing. All that exists is what actually exists. This was not intended as a substantive thesis, setting limits on one's ontological commitments, since it says nothing about what actually exists, and the thesis would seem trivial if it were not for David Lewis's modal realism, which rejected it. On Lewis's picture,² reality consists of a plurality of universes that are spatio-temporally disconnected from each other. 'Actual' is

CONTACT Robert Stalnaker  stal@mit.edu

*Apologies to Hilary Putnam for appropriating the title of one of his famous papers (Putnam 1980). My topic is not unrelated to his, but I am not going to talk about his paper. I chose the title because it says what my paper is about.

© 2016 Canadian Journal of Philosophy

an indexical term applying only to a proper part of reality, the part that is spatio-temporally connected with us (or when used by speakers of our language in other parts of reality, to refer to the parts spatiotemporally connected with them).

'Contingentism' is a more recent label for a thesis in modal metaphysics, a term coined by Timothy Williamson to contrast with 'necessitism', which labels a doctrine he defends. The necessitist's thesis is simple: everything that exists exists necessarily, and nothing could exist except what actually does exist. Contingentism is the contrary thesis, the denials of the two parts of necessitism: There are things that exist, but might not have, and there might have been things other than those there are. Necessitism is not modal realism, a doctrine that Williamson rejects, and it has no need for a distinction between what really exists and what actually exists. But the thesis does allow for the existence of things that might be people or physical objects, but in fact are not: things whose only properties are modal properties, such as the property of being possibly (but not actually) a table, a cabbage, an aardvark, or a person. And while the thesis holds that actual persons, tables, cabbages, and aardvarks are necessary beings, it allows that these things might have existed only in this form of pure potentiality.

Lewis acknowledged that his modal realism conflicts with unreflective common sense, and he took unreflective common sense seriously. (He thought the 'incredulous stare' that his doctrine elicited to be the strongest argument against it.) His defense of the doctrine was highly theoretical: he judged that despite the counterintuitive character of the doctrine, it offered the best theory for unifying a rich system of concepts involving modality, and for explaining the roles of those concepts in our cognitive and practical lives. The conflict with unreflective common sense is perhaps not quite so strong in the case of Williamson's necessitism as it is with Lewis's modal realism, but I think it should be acknowledged that it seems *prima facie* reasonable to think that the ordinary things we find in the world, including ourselves, are things that might not have existed at all, and that there might have existed things that in fact have no existence at all. It is, according to an intuitively natural view, contingent not only how existing things are arranged; it is also contingent what things there are to be arranged. So I think one needs a strong theoretical reason to reject these contingency intuitions, but it also must be acknowledged that the tensions between the contingency intuitions and a natural theoretical account of modality lie close to the surface. Orthodox Kripke semantic models for modal logic are hard to reconcile with the thesis that there might have been things that do not actually exist, but that semantic framework nevertheless seems to provide both intuitive insight into the relations between a diverse range of modal concepts, and powerful formal tools for getting clear about compositional semantics for the languages we use to describe modal phenomena. If the contingency intuitions are to be successfully defended, we need to find a way to interpret the orthodox

models in a way that reconciles them with the contingency intuitions, and that explains their conceptual and technical success. I think the contingentist theses can be defended, and reconciled with orthodox Kripke semantics (extended to higher order modal languages, as Williamson's theory does). My book, *Mere Possibilities*, is an attempt at such a defense, but that defense was not as explicit as it should have been about the status that I take Kripke models to have, and about the relation between these models and the reality that one is using them to model. This paper is an attempt to spell out in a little more detail what I take this relation to be.

Here is my plan: In section 2 I will describe the orthodox semantics for modal languages, and say what the tension is between the contingency intuitions and this way of modeling modal statements. In section 3 I will make some general remarks about models and model structures, the idea of an *intended* model, and the usual way of understanding the contrast between realistic and instrumental uses of models. I will then suggest the possibility of a more complex relation between models and reality, an interpretation of models that treats them neither as merely instrumental, nor as intended models in the simple sense. The claim will be that we can use models to give a fully realistic interpretation of a modal language without giving a realistic interpretation of the models we are using. In section 4 I will try to clarify the general strategy with an analogy with the use of geometric models in a theory of physical space. I will suggest that the way the Galilean relationist about space uses Newtonian geometric models is analogous to the strategy I will propose for reconciling modal semantics with the contingency intuitions. In section 5, I will sketch the strategy as applied to Kripke models, and argue that it helps, not only to reconcile the contingency intuitions with the model theory, but also to motivate an attractive general metaphysical view of the nature of modal phenomena. In section 6, I will respond to some objections, and then close, in section 7, with some general methodological remarks.

2. The contingentist intuitions and Kripke models

The contingentist intuitions consist of the following two claims

- (1) Some things exist only contingently.
- (2) There might have been things that do not in fact exist.

It is (2) that is the more problematic, but there are good reasons to think that the two claims must go together. In any case, (2) seems as intuitively compelling as (1). Just as it seems plausible that Hillary and Bill Clinton might have had no children, in which case Chelsea Clinton would not have existed, so it seems equally plausible that they might have had a son rather than, or in addition to, a daughter. But it is at least *prima facie* reasonable to think that there is no actual thing that could have been Bill and Hillary Clinton's son.

The contingency intuitions proliferate: If Hillary Clinton is a contingently existing object, then it seems reasonable to think that the properties of being identical to Hillary, or being the daughter of Hillary are also contingent. And if there are singular propositions about Hillary (for example, that she is running for President in 2016) then one might think that these propositions are *object-dependent*, and so also things that exist only contingently. This extension to higher order contingency is not compulsory: Alvin Plantinga accepts the contingency intuitions about individuals, but rejects the extension to properties and propositions. But the generalization of contingentism is natural, once the first step is taken, and the tension between the orthodox semantics and the contingency intuitions is there even if one's contingentism stops at the individual level.

So what is the tension? The idea motivating possible-worlds semantics is that a proposition will be possibly true only if it could be realized in a specific way, which is to say only if there could be a complete possible situation in which that proposition is true. So the formal semantics for a modal language is given in terms of a structure that includes a set of possible worlds (usually with one of them representing the actual world of the model), thought of as maximal ways things might be. A statement will be possibly true if and only if it is true in at least one of these maximal possibilities.

A semantic framework of this kind provides a natural generalization of the semantics for an extensional language. Whatever structure the extensional semantics provides for the interpretation of the extensional language, the modal version of that logic will provide a structure of that kind for each possible world. So, since a model for first-order extensional quantification theory postulates a domain of individuals, the semantic strategy suggests that a model for first-order modal quantification theory should postulate a domain of individuals for each of the possible worlds.

The contingency intuitions imply (in the context of models of this kind) that the domains will vary from world to world. In particular, (2) implies that there will be counterfactual worlds with domains that contain things that are not in the domain of the actual world. But on the intended interpretation, the domain of the actual world is the domain of everything there is. For those who reject Lewis's modal realism, there is no room for anything outside of this domain. Our theorizing about modality is taking place in the actual world, and the materials available for theorizing – the things we appeal to in interpreting the modal language – must all be actual things. So what are these mere possibilities that inhabit the domains of the possible worlds that verify examples of our second contingency intuition?

It should be noted that the contingency intuitions, stated in ordinary modal language, do not, on the face of it, imply that there are merely possible things – things that might, but don't, exist. What must exist, according to claim (2) is only the possibility of things that do not in fact exist, and we have no reason to deny that this is an actual possibility. The move from the possibility of things of

a certain kind to the existence of things that might be of that kind comes from the model-theoretic representation. But that representation has both a compelling intuitive motivation, and considerable success in clarifying the structure of modal discourse, so we should be reluctant to give it up, or even to treat it as a merely instrumental device. I think we can reconcile the orthodox model theory with a realistic account of the language that the models are used to interpret, but doing so requires a more complex understanding of the representational role of models. In the next section I will sketch the standard simple account of the relation between models and reality, as I understand it, and the complications that I think need to be added to this simple account.

3. Models, intended models, and instrumental uses of models

A model (in the model-theoretic sense) for a language has two parts: first, a structure of some kind that is specified independently of any language; second, a valuation function from expressions of the language to objects that are defined in terms of elements of the structure. In first-order *extensional* quantification theory, the structure is just a domain of individuals. The language contains n -ary predicates for each n , and sometimes names as well. The function assigns individuals from the domain to the names, and subsets of the set of all n -tuples of members of the domain to the n -ary predicates. In the case of a model for first-order modal quantification theory, the model structure consists of a non-empty set (the possible worlds), with one of them designated as the actual world of the model, sometimes a binary relation on the set,³ and a domain of individuals for each possible world. The valuation function assigns to each primitive descriptive expression an *intension*, which is a function from possible worlds to extensions, where extensions are as they are in extensional quantification theory. So, for example, the intension of a binary relational predicate is a function from possible worlds to sets of ordered pairs of members of the domain of that world. In general, the definition of a model structure is abstract: the set of 'possible worlds', for example, might be any arbitrary non-empty set – say the numbers 1, 2 and 3. Models of this kind might be used for various purposes, for example to show the logical consistency of some set of sentences of the language, but for realistic interpretations of a language, one talks of *intended* models. On the usual way of understanding intended models, the component of the model that is the structure is supposed to be the actual subject matter of the language – what the language is being used to talk about. It is assumed that a use of a model is either intended, or merely instrumental, which in the latter case means that the model is used to show something about the logical capacities of a language to represent something, but not anything about what the language is actually used to represent.

Some may talk of *the* intended model of a language, or *the* intended structure. One may, for example, take the intended domain of extensional quantification

theory to be the domain of absolutely everything that exists. Jon Barwise and John Perry take this view of what is required by a realistic interpretation of possible worlds semantics: 'If the model theoretic structures of possible worlds semantics . . . are supposed to be a model of something, say superreality . . . then there ought to be *one* that is an intended or standard model, *the one* that really corresponds to superreality.'⁴ (They might better have said, 'the one that *is* superreality. The correspondence is between the expressions of the language and its subject matter – a structure that constitutes modal reality itself.) But realism about model structures does not require completeness or comprehensiveness, and does not even require that there be a coherent notion of metaphysical completeness or comprehensiveness. So one who rejects the intelligibility of the idea of a domain of absolutely everything can still give a realistic interpretation of extensional quantification theory – an interpretation that specifies or presupposes a domain that is what the quantified sentences of the language are taken to be about. A model, to be an intended model in the standard sense, could be a model with a domain that is potentially extendable, but it must be a domain that includes *only* things that exist. (Does that imply that if gods do not in fact exist, the theist who intends a realistic interpretation of her language cannot include gods in her domain of discourse? Yes it does imply that, though of course the theist can coherently claim that there are gods. It is just that if there in fact are no gods then these claims will be false, on the theist's intended interpretation.)

But the situation of the theorist who accepts both the contingency intuitions and the Kripkean semantics is worse than that of a theist in a godless world. This theorist acknowledges the non-existence of some of the things that his model theory seems to be quantifying over (the merely possible things that inhabit the domains of some counterfactual possible worlds), and this seems to foreclose the possibility of a realistic interpretation of the model theory. I think this is right: the contingentist cannot take the model structures that are appropriate for modeling metaphysical modality to *be* the reality, or the part of reality, that our modal languages are talking about. But the relationship between models and the reality that is modeled may in some cases be more complex.

The central claim of this paper is that one may intend a realistic interpretation of a *language*, but use nonrealistically interpreted models as aids in stating the compositional semantics for the language. That is, one may use a model as a representational device, interposed between the language being interpreted and the reality that is its intended subject matter. So there will be a correspondence between the language and the model structure, and a further correspondence between certain features of the model structure and the reality being modeled. The interpretation of the modal object language will be fully realistic in the following sense: Each closed expression of the language (predicates, names, and sentences) will denote something that is real (at least according to the contingentist's theory of what there actually is.) This will be true, not only of

the first-order modal language, but also of the higher order extensions of the language of the kind that Timothy Williamson discussed in his book, *Modal Logic as Metaphysics*. But, with the more complex use of models, there will be two steps to be taken on the road from an expression of the language to the entity that it denotes: from expression to entity in the model structure to which it corresponds, and from entity in the model structure to object, property, relation, proposition that it denotes. I will try to explain how this works in the next two sections, starting in section 4 with an analogy that provides a precedent for the strategy I want to promote, and then in section 5 by sketching the application of this strategy to the use of Kripke models to interpret modal languages in a way that is compatible with the contingency intuitions.

4. Geometric models and relativism about space

Consider the contrast between a Newtonian conception of absolute space and a Galilean conception of relative space. The relativist claims that there is no fact of the matter about absolute motion, no difference between an inertial frame moving at constant velocity through space and the frame at rest. The Newtonian acknowledges that according to the physical theory they both accept, there is no difference in the way bodies behave in an inertial frame at rest and the way they behave in one moving at constant velocity. The issue between the absolutist and the relativist is whether there is a *metaphysical* difference. According to the Galilean, physical geometry is a system of spatial relations between bodies, but the relations are not grounded in intrinsic spatial properties. There are no *locations* in space: the spatial points in a Euclidean geometric model of space are just devices for representing the relations. There is, for example, a real relation that holds between two bodies that are 10 m apart, but the fact that this relation holds is not grounded in the fact that one of the bodies is located at a certain spatial point that is 10 m from the point where the other is located. The Galilean can model his relativism by using the same models that the Newtonian absolutist uses, but he supplements the models with a family of permutation functions each of which maps space-time points onto space-time points. The set of mappings will be defined so that the distance relations between spatial points are preserved. Suppose we use quadruples of real numbers to represent space-time points, and that there is a permutation function in the family that maps each point $\langle t, x, y, z \rangle$ to a point $\langle t, x^*, y^*, z^* \rangle$. (Since time is absolute in the physics shared by the Newtonian and the Galilean, the time coordinate will always map to itself.) This will be an admissible permutation function only if the distance from the point $\langle x, y, z \rangle$ to point $\langle u, v, w \rangle$ in Euclidean space is the same as the distance from $\langle x^*, y^*, z^* \rangle$ to $\langle u^*, v^*, w^* \rangle$. But the points themselves may be different. If, in the model, the center of mass of the sun is located, at a certain time t_0 at point $\langle x, y, z \rangle$, then in the permutation of the model, it will be located (at time t_0) at point $\langle x^*, y^*, z^* \rangle$. The sun will be at rest from t_0 to t_1 in the basic

model if it is located at the same spatial point throughout this interval (that is, if it is located at $\langle x, y, z \rangle$ for all t from t_0 to t_1). But it might be in motion (at constant velocity) in an admissible permutation of the original model, since it might be that as the time shifts from t_0 to t_1 , the permuted spatial location at that time, $\langle x^*, y^*, z^* \rangle$, moves uniformly in some direction in a straight line. The Galilean's claim is that the original model and the permuted model represent the same factual situation. The facts about the motion of bodies – the changing spatial relations between them – are the facts that are invariant under the admissible permutation relations.⁵

Notice that the Galilean is not aiming to give a general criterion for distinguishing features of the geometric model that are artifacts from features of the model that will be shared by *any* model representing the same reality. I am not sure what it would mean to do that. The aim is the more modest one of representing the distinction between certain spatial properties and relations (represented in the models by sets of locations, or sets of ordered n -tuples of locations) that are real from objects of this kind that are artifacts of the model. So consider the fact that in our basic model, we used quadruples of real numbers as our space-time points. The model presupposes an arbitrary coordinate system. We might connect our model with physical phenomena by stipulating that the point $\langle 0, 0, 0, 0 \rangle$ shall be the center of mass of the sun at the moment of Isaac Newton's birth. The fact that this particular spatial point is the origin of our coordinate system, of course, has no representational significance. But once we have made this stipulation in our basic geometric model, the fact that the center of mass of the sun at the time of Isaac Newton's birth is the origin, $\langle 0, 0, 0, 0 \rangle$, will be invariant under the admissible permutation functions. The general idea behind the permutation strategy is to use models (or model structures) as representational devices, and it is the model plus the family of permutation functions that is the representational device. *Any* representational device will have features that distinguish it from other representational devices for representing the same thing. If the model structure were the reality itself, as it is on the simple picture, then (trivially) nothing about it would be an artifact of the model. The Newtonian absolutist might say that in the intended model, the points are represented, not by quadruples of real numbers, but by the locations themselves. The Galilean can't say this, since he claims there are no such things as locations themselves, but he can still use mathematical models (with a particular conventional coordinate system) to say in a precise way which of the properties and relations of a particular kind (as represented in the model) are real, and which are only virtual properties and relations.

One can generalize about the kind of situation in which the permutation strategy used by the Galilean is appropriate. Leibniz stated and defended the principle that *all relations between substances are reducible to (supervenient on, grounded in) the intrinsic properties of the relata*, but this is a substantive metaphysical thesis. Galilean relationism⁶ about space is a paradigm example of

a theory that rejects this Leibizian thesis,⁷ since it holds that there are spatial relations that are not grounded in the intrinsic locations of the related bodies. There are also other structures that seem to have this feature. Preference relations, for example, seem to be essentially comparative. One might try to ground them in some measure of the absolute degree to which one likes or wants something, but this would not be very plausible. Measures of the extent of preference (as in utility theory) are usually taken to be irreducibly relational, and one uses the permutation strategy to make this way of understanding utility values explicit. All the positive linear transformations of a utility function are taken to be equivalent. A second example: I suggested once that a functional theory of phenomenal qualia might be an example of an ungrounded relational structure, since the capacity to discriminate between experiences is the central capacity on which such a theory is based.⁸ The inverted spectrum example is a *prima facie* problem for a functional theory of qualia, since it suggests the possibility of a perfectly symmetrical system of the relations between types of qualitative experience. The permutation strategy helps to explain how symmetrical types of qualitative experience that are in a sense functionally equivalent, can still be different. There is a permutation that maps red to green and green to red, preserving discriminatory capacities, but it is an invariant feature of the structure that red is discriminable from green. (Whether such an account is a plausible account of qualia is another question, but at least it illustrates the abstract idea.)

5. Modal reality, and model structures

As I said in the last paragraph of section 3, the contingentist's strategy is to use a model structure, not understood as the structure of an *intended* model in the usual sense, but instead as a representational device interposed between the language and the reality that is being modeled. So let me start with some general remarks about the subject matter that we will be intending our modal language to be about. The type theory developed by Daniel Gallin⁹ and discussed by Williamson in his book,¹⁰ is a useful framework for describing the subject matter of a modal language, for both necessitists and contingentists. The type theory is a typology of entities that can be specified independently of any language that the type theory might help to interpret. There are two basic types, first, *e* for 'entity' (the domain of individuals)¹¹ and second a type that is mysteriously labeled $\langle \rangle$ (a pair of empty brackets) for propositions.¹² The other types are defined by a single recursive rule: if t_1, \dots, t_n are types (basic or derived), then $\langle t_1, \dots, t_n \rangle$ is a type. The interpretation of the derived types is this: $\langle t_1, \dots, t_n \rangle$ is the domain of *n*-ary relations between things of types t_1, \dots, t_n , respectively. So for example, $\langle e, e \rangle$ is the domain of binary relations between individuals, $\langle e, \langle \rangle \rangle$ is the domain of binary relations between an individual and a proposition, $\langle \langle e, e \rangle \rangle$ is the domain of monadic properties of binary relations.

The basic type e has no intrinsic structure – that is, no structure that is imposed by the type theory itself. But the abstract interpretation of the type theory does impose a structure on the other types, and on the relations between the entities of the different types. The domain of propositions, for example, brings with it some basic properties and relations that are essential to the idea of a proposition, for example the relation of entailment that holds between a class of propositions and a proposition. The general abstract theory of propositions will impose certain closure conditions on the class of propositions, for example that every proposition has a contradictory (where the relation of being a contradictory is definable in terms of entailment). The derived types will also have related essential structural properties, required by the most general assumptions of what properties and relations are, and there will be a rich structure of interrelations between the members of the different types. An individual (member of type e) and a monadic property of individuals (member of type $\langle e \rangle$) will determine a (presumably unique) proposition (member of $\langle \rangle$) – the proposition that is true if the individual instantiates the property.

The *categories* of the type theory are built recursively from the basic categories, and the theory will impose an elaborate structure on these categories, specifying certain closure conditions, and relations between elements of the different categories. But we need not assume that the membership of the domain that is specified by a derived type is determined by the membership of the domains of the types from which it is derived. That is, it could be that the structural conditions imposed on the types and the relationships between them do not ensure that once we have specified the domain of individuals, and the domain of propositions, we have thereby determined the domains of all properties and relations, even on the most abundant construal of properties and relations. If our metaphysics makes certain necessitist assumptions, then a reduction of the domains of properties and relations to the domains of individuals and propositions may be possible. Consider the class of all functions from maximal consistent propositions to subsets of the domain of individuals (things of type e). The class of maximal consistent propositions is definable in terms of the domain of propositions, given its essential structure, so the necessitist might identify the domain of monadic properties of individuals (things of type $\langle e \rangle$) with this set of functions, and so reduce the domain of monadic properties to a domain definable in terms of the basic types. If this works, one could reduce all the higher types in a similar way. But the contingentist cannot do this, since for the contingentist there may exist a property F that is possibly exemplified by something that does not actually exist (and so is not in the domain of the type e). For the contingentist, the relational structure exhibited by the type theory will be ungrounded in a way that it would be grounded if the necessitist's metaphysical assumptions were true. This kind of ungrounded relational structure is, in general, the kind of phenomenon that motivates the permutation strategy illustrated by the Galilean theory of space.

Recall that the tension between the contingentist intuitions and Kripke models arose because the possibility that there is a thing of a certain kind is modeled by a thing (a member of a domain of a nonactual world) that is possibly of that kind. The necessitist takes this commitment at face value, but the contingentist notes that (aside from what the models seem to require) the inference does not seem to be mandatory. That is, for example, it does not seem incoherent on its face to assume that there really is the possibility that Hillary Clinton have had a son, even though there is nothing that might have been Hillary Clinton's son. The aim is find a way to use orthodox models to clarify modal phenomena without justifying this kind of inference.

So how can we use Kripke models to interpret a language for talking about a modal reality that has this kind of ungrounded relational structure? Just as the Galilean uses exactly the same geometric models as the Newtonian, so the contingentist uses the same Kripke model structures as the necessitist – model structures that would be necessitist if they were interpreted realistically. The materials in a model structure (a set of **possible worlds**, with a domain of **individuals** for each **world**) suffice to determine **propositions** (represented by subsets of the given set of **possible worlds**) and **properties** and **relations** (represented by functions from **possible worlds** to subsets of the domain of the **world**, or more generally to subsets of the set of n-tuples of members of the domain of the **world**.) Since the plan is to use the model structure itself merely as a representational device, there is room for confusion in the use of terms for the elements of the model structure. So when the terms ('property', 'relation', 'possible world', 'proposition', etc.) are used to refer to elements of the model structure, I write them in bold italics, while when they are used as terms for components of the reality we mean to talk about (members of the categories in the type structure) I will use them unaccented. **Possible worlds** and **possible individuals** are just arbitrary objects – perhaps they are all real numbers. Some, but not all, of these things, and some but not all of the functions definable in terms of the primitives of the model structure – the **propositions**, **properties** and **relations**, will represent real entities – real propositions, properties, and relations. Which ones? That question is answered by classes of permutation functions mapping the **possible worlds** of the model structure onto the **possible worlds**, and mapping the members of the domains of those **worlds** onto the domains of corresponding **worlds**. The representational device we are interposing between the language and the alleged reality will consist of a Kripkean model structure plus a class of permutation functions of this kind associated with each **possible world** of the structure. A **proposition p** (a set of **possible worlds**) will represent a real proposition (relative to **possible world w**) if and only if it is invariant under all of **w**'s permutation functions, which is to say if and only if each of the permutation functions takes the set **p** to itself. The **proposition p** will be merely virtual if it is not invariant. Of course the set **p** itself, whether invariant in this sense or not, is a real set – perhaps, as I suggested, it is a set

of real numbers. By calling some of these sets virtual, I am saying that they do not represent genuine propositions, where genuine propositions are the members of the domain of propositions in the type theory that we are taking to be the reality that our model language represents. We interpret the model structure with a correspondence relation between the invariant elements of the domains of *individuals*, *properties*, *relations* and *propositions* and elements of the appropriate kind in the corresponding categories in our type theory.

One of the *possible worlds* in our model structure is identified as the *actual world* of the model, and this component of the model structure will correspond to the maximal true proposition – the way the world actually is.¹³ The classes of permutation functions are defined relative to each *possible world*, but it is the *actual world's* permutation functions that are relevant to determining the actual commitments of our theory. (The other classes of permutation functions will be relevant to the interpretation of iterated modal claims.)

We have said nothing yet about a modal language, or about any correspondence between a modal language and the model structure. Our focus has been on the characterization of the subject matter that our language will be talking about, on a model structure (which is specifiable independently of any interpretation of a modal language), and on a correspondence between the model structure and the elements of the different types of our type structure. The language will be interpreted in the standard way by a correspondence between the primitive expressions of the language and semantic values that are definable in terms of the basic components of the model structure. It will be a constraint on a realistic interpretation of the modal language that each name and each primitive sentence letter or predicate be interpreted (in the model structure) by an *individual*, *proposition*, *property*, or *relation* that is invariant, relative to the actual world's class of permutation functions. So, since the invariant elements of the model structure correspond to elements of the appropriate kind in the domains of the type structure, all of these primitive expressions will correspond indirectly to elements of things of the appropriate kind in the intended subject matter of the theory. And the structural conditions on the relational structure exhibited by the type-theory will ensure that each of the closed complex expressions of the language will correspond to real properties, relations, and propositions if the primitive ones do.¹⁴

My aim in this paper is to clarify in a general way the role of Kripke models in giving a realistic semantics for modal concepts that is compatible with a contingentist metaphysics. Discussion of the formal details are for another occasion, but the appropriate constraints on the permutation functions have been worked out, and result in interesting structures.¹⁵ The relation between models and reality is complex on the contingentist picture, but I think the tensions in our intuitive conception of modal phenomena are real: both the contingency intuitions and the Kripke structures have features that need to be reconciled. The overall comparative evaluation of necessitism and contingentism will have

many dimensions, but I believe there is a coherent contingentist picture that is attractive and worth developing. There may be some deep reasons for resisting the picture, but there are also some objections that are based on misunderstandings. In the next section I will respond to two of these.

6. Objections and replies

Nicholas Jones, in a recent paper,¹⁶ argues that the invariance strategy for reconciling the contingency intuitions with Kripke models 'is flawed: [Stalnaker's] characterization of the representational/non-representational distinction does not cohere with his metaphysical view.' His argument for this conclusion presupposes that the aim of this strategy is to give a general characterization of all the features of a model that are artifacts of the model – all the features of a model that distinguish it from any model that is representationally equivalent to it. I doubt that this is a coherent aim, but in any case, as I have tried to make clear in sections 4 and 5 above, it is not the aim of the invariance strategy. Jones argues that the strategy fails in the representation of the Galilean relationist theory of space as well as in the attempt to reconcile the use of Kripke models with contingentism. Let me look first at that simpler geometric case:

'Stalnaker's relationist denies that spatial locations exist: "there are really no such things as spatial locations – there are just spatial relations between things." But it is an invariant feature of *all* the spatial models in any permutation class that there are spatial locations. Because Stalnaker's relationist rejects the existence of spatial locations, she faces a choice: (a) reject every spatial model as inaccurate; or (b) reject the Invariance-based analysis of non-representational artefacts.' I choose (b), given Jones's interpretation of the invariance-based analysis of nonrepresentational artifacts, but we can still use the invariance strategy to distinguish the spatial properties and relations between bodies that are real (according to the theory) from those that are virtual. That is all the invariance strategy is trying to do.

Jones's argument for the failure of the invariance strategy in the modal case has the same structure, with possible individuals playing a role analogous to locations in the spatial case. Before looking at the argument, let me make a terminological clarification. The representational device that we are interposing between the language and the modal reality described by the type theory is a single Kripke model structure, supplemented with a family of permutation functions. Both the structure and the permutation functions are specified independently of any language that they might be used to interpret. So models, in the model-theoretic sense, that include, as one component, a valuation function mapping expressions to entities definable in terms of the elements of the model structure are not involved at this point. So we are not talking about invariance across a class of models, but invariance of permutation classes of **properties, relations, propositions**, etc. that are features of the one model structure. But

the claim that some sentence of a modal language is satisfiable will have implications for the existence of model structures meeting certain conditions.

Jones's argument concerns the interpretation of a sentence of the modal language that he calls ' $\exists\text{con}$ ': $\diamond\exists x@ \forall yx \neq y$, (where '@' is the actuality operator). The contingentist claims that $\exists\text{con}$ is consistent, and this will be true only if there is a model structure in which there is a **possible world** with a domain that contains an **individual** (call it ' n ': that is, let ' n ' be a metalinguistic name for that component of the model structure) that is not in the domain of the **actual world**. But the constraints on admissible permutation functions will ensure that every permutation takes n to a possibly different **possible individual** that is also not in the domain of the actual world. So (Jones argues) it is an invariant feature of the model structure, supplemented with the permutation structure, that, as he puts it, $\exists\text{con}$ has a 'witness': a merely possible individual in virtue of which it is true. The witness will be different for different permutations, but there will always be one. This fact about invariance is true, but irrelevant to the question whether the invariance strategy commits the theorist to the existence, in modal reality, of merely possible individuals. A **possible individual** of the model structure corresponds to a real individual (a member of the domain of individuals in modal reality) only if the relevant permutation functions are invariant (all take the **individual** to itself.) The **individual** n does not pass the test, nor do any of the different **possible individuals** that are the images of n for some admissible permutation function. So, Jones's argument does not show that the invariance strategy commits the contingentist to the existence of merely possible individuals.

Williamson, in his book, gives a somewhat different argument for a closely related claim: that the contingentist cannot coherently reject the Barcan formula, (BF) $\diamond\exists x\phi \rightarrow \exists x\diamond\phi$. Williamson also seems to assume that the ambition of the invariance strategy is to give a general criterion of representational significance, but I think the main problem of his argument is that it conflates the **individuals** that are the components of a model structure with the individuals that, on the intended interpretation, are the range of the quantifiers of the language, and it conflates the classes of permutation functions on a model structure with different models.¹⁷

Here is the argument, as I understand it: Suppose we have a model M that invalidates BF. This requires a model structure in which there is a **possible world** w that has in its domain an **individual** o that is not in the **domain** of the **actual world**. But independently of the model, " o is an actual individual, just like anything else. . . . o is an actual individual that represents a non-actual individual (but not any particular non-actual individual)."¹⁸ (This last remark needs interpretation. It is right that o does not represent any particular non-actual individual, since there are no non-actual individuals. What it represents is the possibility that there be an individual that does not in fact exist.)

Now (the argument continues), since ***o*** is an actual individual, there must be another model, M^* in which ***o*** is in the domain of the actual world. This is true, but M^* will not be relevant to the class of permutation functions that map the ***worlds*** of the structure onto the ***worlds***, and the ***possible individuals*** onto the ***possible individuals***, since all of the permutations of ***o*** will take it to an ***individual*** in the domain of a non-actual world. But Williamson seems to be assuming that since ***o*** is an actual individual, it must represent an actual individual (itself) in some admissible permutation on the model structure, and he is right that this would not be compatible with the use of ***o*** to represent the possibility of an individual that does not in fact exist. It is also right that if ***o*** is in the domain of individuals in our type structure that represents the intended subject matter of our language, then there will have to be something in the ***domain*** of the ***actual world*** of our model structure that corresponds to ***o***. But there is no reason why it has to be ***o*** itself.

7. Concluding methodological remarks

The necessitist's theory is simpler than the contingentist's, and some may take this to be a reason to be a necessitist. Simplicity is a theoretical virtue, so long as the simpler theory accounts as well for the data, but it is less clear in metaphysics what the data are than it is in science. The contingency intuitions are not unassailable data. We do need to account for the distinctions that they reflect, but the necessitist has a story to tell about the special status of the things that the vulgar are inclined to say are merely contingent things, and about the facts that ground the possibilities of things that we may be inclined to say do not in fact exist. Why jump through the complicated hoops that the contingentist metaphysics forces us to jump through? One possible answer is that the questions that contingentism raises are good questions, and that being forced to answer them brings out features of reality, and of our conceptual resources for talking and thinking about reality. Whether this is right, or whether alternatively the questions are forced on us only by a false metaphysical picture is a delicate question to be answered only by developing the alternative conceptions of metaphysical modality.

The invariance strategy is an appropriate general strategy for representing relational structures that are in a sense ungrounded in the intrinsic properties of the things that are related. The Galilean relativist grants that the Newtonian's models are simpler, but insists that reality does not provide the distinctions that would ground those models. We can nevertheless use those same models, she claims, to represent physical reality, and to factor out the distinctions to which nothing in reality answers. I think the general strategy works, and that it helps us to understand any structure of relations that do not satisfy the Leibnizian metaphysical grounding principle. It is, however, a further metaphysical question

whether there are ungrounded relational structures of this kind, and if there are, where they are to be found.

The necessitist picture is an extensionalist picture, taking at face value the extensional, set-theoretic structures that modal model theory uses to clarify the compositional semantics of modal languages. Properties and relations are explained in terms of the individuals that instantiate them, or that might instantiate them, and propositions are explained in terms of sets of complete ways that the domain of all the things there are might be arranged. On the contingentist picture, possibilities, properties and relations are not reducible to the ways the things there are might be arranged. The hierarchies of propositions, properties and relations cannot be built with the materials found at the ground level: the individuals and possible worlds. But the contingentist can still use the same extensional set-theoretic tools that the necessitist uses, with the help of the permutation strategy. The necessitist picture may in the end be defensible on metaphysical grounds, but I think the contingentist picture provides a coherent and attractive vision of modal reality that is worth taking seriously.¹⁹

Notes

1. Kripke 1980.
2. See Lewis 1986.
3. I will ignore the binary accessibility relation that is a constituent of a frame in a general Kripke model structure for modal languages, restricting attention to the S5 case where the accessibility relation is universal; necessity (in a given world w) is truth in all possible worlds, and not just in all that stand in some accessibility relation to w .
4. Barwise and Perry 1985. 120.
5. The kind of permutation strategy I am sketching is familiar in mathematics, and in measurement theory, but it is important to distinguish the use of this strategy by the Galilean from its use to bring out the conventionality of units of measurement. Newtonians and Galileans will agree, for example, that the choice of a coordinate system for physical space and the unit of measurement for representing distance relations are conventional. Any permutation of the numbers used to represent the distance between two bodies that preserves the ratios between distances will be an equivalent representation. What is permuted, in this case, is the relation between a number and the relational property that holds between two bodies iff they are a certain fixed length apart. The relation, x being 2 meters from y is the same relation as x being 4 half-meters from y , but the Newtonian and the Galilean agree that this is a real relation. Consider a possible world like ours, except that all the distance relations between bodies were doubled. This is very different from considering a world in which we meant by 'one meter' a distance relation that was what we actually mean by 'one half meter' (cf. Grünbaum 1964). The Newtonian claims that there are location properties, even though it is conventional what quadruples of real numbers we use to represent them. The Galilean, in contrast, denies that there are such properties. This disagreement is reflected in the permutation functions that are permitted by the Galilean's theory, which cannot be interpreted as permuting just the numerical representation of spatio-temporal

locations. The disagreement is about whether there is a real difference between rest and uniform motion at constant velocity.

6. It is an irony that Leibniz is both the one who propounded this grounding thesis, and also a prominent defender of a relational theory of space. His views are reconcilable, at least on the surface, since on his metaphysical theory, substances do not stand in spatial relations. His relationism about space is more radical (at least on one way of reading him) than that of the Galilean: space itself is merely ideal. The view seems to be that spatial relations between phenomena are ultimately reducible to the intrinsic properties of substances.
7. I think one should think of physical space itself as a structure of properties and relations of the things in space. On this way of thinking, for the absolutist, spatial locations, and regions defined by sets of specific locations, are intrinsic properties of the things at the location, or in the region. Spatial relations are determined by these intrinsic properties: once you have identified the specific locations of the relata, you have determined how they are spatially related to each other. But for the Galilean, the Leibnizian principle fails.
8. Stalnaker 2000.
9. Gallin 1975.
10. See Williamson 2013, 221ff.
11. As I understand him, Williamson interprets this domain as the domain of absolutely everything, including all of the members of the domains that are the other types. But as I suggested in section 3, we need not take a realistic specification of the subject matter of some language – even a language for doing metaphysics – to be comprehensive. So we could understand the type *e* as just a class of basic individuals, and take the other types to be classes of entities that are disjoint from the things of type *e*.
12. The motivation for the label, and for the idea that this type is really a derived type, is the idea that just as *n*-ary relations can be thought of as properties of *n*-tuples, so propositions can be thought of as 0-ary relations, properties of 0-tuples.
13. So our model structure is what Williamson calls an ‘inhabited model structure’. There is a ‘fixed point’ constraint on the permutation functions, ensuring that in any of *world w*’s permutation functions will take *w* to itself, so the *actual world* of the model structure, with its domain, will be, by definition, invariant.
14. I say that if the primitive predicates sentence letters and names, all denote real entities of the appropriate kind, then so will the complex *closed* predicates and sentences. As Bruno Jacinto has shown (in as yet unpublished work), this is not true for *open* expressions, which have as their semantic values *propositional functions, propositional function-functions*, functions from *individuals* to *properties*, etc. Virtual propositional functions seem to play an ineliminable role in the compositional process by which complex predicates and quantificational constructions are interpreted, even though all of the complex predicates and closed quantificational constructions get an invariant interpretation. I should perhaps not have found this as surprising as I did, since the whole detour through virtual models (that in a sense, represent reality as more fine-grained than it is), together with permutation functions (that serve to bring us back to the right grain in our account of what is real) would be unnecessary if we could do the compositional semantics directly in terms of the real entities that our language is talking about.
15. I spell out some of the details in two appendices to my book, Stalnaker 2012, but the appendices are overly compressed, and contain some errors, which have been pointed out and corrected in work by Peter Fritz (Fritz, *forthcoming*).

16. Jones 2016.
17. I take some responsibility for the latter conflation, since as Williamson notes in a footnote (Williamson 2013, 189, note 48), I confusingly use the label 'model' for what is in fact a model structure. He assumes I must really have meant to talk about a model (the structure plus a valuation function), but what I meant was just a model structure.
18. Williamson 2013, 191.
19. Thanks to Peter Fritz, Bruno Jacinto, Nicholas Jones and Tim Williamson for very helpful discussion and correspondence about these issues.

Notes on contributor

Robert Stalnaker is the Laurance S. Rockefeller Professor of Philosophy at MIT. He is the author of *Inquiry* (MIT Press, 1984), *Our Knowledge of the Internal World* (Oxford U. Press, 2007), *Mere Possibilities* (Princeton U. Press, 2012), and *Context* (Oxford U. Press, 2015).

References

- Barwise, J., and J. Perry. 1985. "Shifting Situations and Shaken Attitudes." *Linguistics and Philosophy* 8: 105–161.
- Fritz, P. (forthcoming). "Propositional contingentism," *Review of Symbolic Logic*.
- Gallin, D. 1975. *Intensional and Higher-Order Modal Logic: With Applications to Montague Semantics*. Amsterdam: North-Holland.
- Grünbaum, A. 1964. "Is a Universal Nocturnal Expansion Falsifiable or Physically Vacuous?" *Philosophical Studies* 15: 71–79.
- Jones, N. 2016. "The Representational Limits of Possible Worlds Semantics." *Philosophical Studies* 173: 479–503.
- Kripke, S. 1980. *Naming and Necessity*. Oxford: Blackwell.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Putnam, H. 1980. "Models and Reality." *The Journal of Symbolic Logic* 45: 464–482.
- Stalnaker, R. 2000. "Comparing Qualia Across Persons." *Philosophical Topics* 26: 385–405.
- Stalnaker, R. 2012. *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton: Princeton University Press.
- Williamson, T. 2013. *Modal Logic as Metaphysics*. Oxford: Oxford University Press.