Possible worlds

Michael De

Forthcoming in the Routledge Handbook of Modality (Otávio Bueno and Scott Shalkowski, eds.)

1 Introduction

Modality concerns what might or must be the case. We use modal expressions such as 'might', 'must' and cognates in a variety of ways, however. One can truthfully say, for example, that the neither yet proved nor disproved Goldbach's conjecture (that every natural number greater than two is the sum of two primes) *might* be true, even if the conjecture is false and hence, in a *metaphysical* sense, necessarily so. What one means here is that the conjecture is compatible with our evidence, since we have neither proven nor disproven it.

Modal expressions are typically *intensional* (with an 's'). An expression α is intensional just in case the substitution of extensionally equivalent expressions under the scope of α need not preserve truth. For instance, even though '8' and 'the number of planets' have the same extension, 'Necessarily, 8 is greater than 1' and and 'Necessarily, the number of planets is greater than 1' do not share their extension (i.e. truth-value). However, while many modal notions are intensional, not all are. Consider e.g. the *de re* modal relation expressed by 'it is necessary of x that it be such that A', where A is to be replaced by a declarative statement. Thus, even though the expressions 'modal' and 'intensional' are sometimes used interchangeably, modality and intensionality are not equivalent.

How best are we to analyze modality and intensionality? The most popular strategy since the mid twentieth century is to employ *possible worlds*. Reference to possible worlds dates back at least to Leibniz (see Mates (1968)), but they do not assume their familiar role until the 1940s when Rudolf Carnap (Carnap, 1946) gave an analysis of modal operators resembling a modern treatment in terms of quantification over what we would now call possible worlds. Such an analysis, often called *possible worlds semantics*, was later generalized throughout the 1940s and 1950s independently by a number of logicians including Saul Kripke and Jaakko Hintikka, and is now the standard treatment for a wide variety of intensional notions including modality. An impressive number of intensional notions have been given possible worlds analyses, only some of which include: conditionality, causation, knowledge, *de se* belief, intrinsicality, dispositionality, aboutness or subject matter, supervenience and dependence, truthmaking, the laws of nature, essence, property, propositional and intentional content, fictional worlds, and truth in fiction. The use of possible worlds in linguistics, logic, and computer science has also seen enormous success. The fact that possible worlds talk has become common parlance in many areas of contemporary analytic philosophy and other fields raises important questions concerning their ontological status and the explanatory value they afford.

Rather than answering the question *What are possible worlds?*, I wish to discuss what I think is a more tractable question, namely,

> *What theoretical roles are possible worlds supposed to play, and are they cut out to play those roles?*[1]

The question is tractable because we can simply look and see to what purposes possible worlds have been put and whether possible worlds analyses have survived the test of time or whether they have been succeeded by superior analyses which either do away with worlds altogether or else demote them to a lesser role.

In giving a partial answer to the question we will begin by looking at traditional possible worlds analyses of intensional and modal concepts (§2). We will then look at three possible worlds analyses that have played an important role in their perceived success, viz. the analyses of (i) modality and possibilities in counterpart theory (§3), (ii) belief contents (§4), and (iii) conditionals (§5).

2 Possible worlds semantics

Let us have a more careful look at possible world semantics and the reasons for its success. Let us denote 'Necessarily' by '$\Box$', 'Possibly' by '$\Diamond$', and 'It is not the case that' by '$\neg$'. Then, according to the simplest quantificational analysis of the broadest sort of necessity, often called *metaphysical*,

($\Box*$) $\Box A$ is true at a world iff A is true at every world.

Given the equivalence of $\Diamond A$ with $\neg\Box\neg A$, we have:

($\Diamond*$) $\Diamond A$ is true at a world iff A is true at some world.

The quantificational analysis is simple and has the virtue of providing a way of determining whether complex modal sentences (or sentence forms), such as those containing a large number of iterated modalities (e.g. 'It is possibly necessarily possible that A only if it is necessary that A'), are true (valid) or not. Before the quantificational analysis, determining which modal inferences were valid rested mainly on potentially shaky intuitions concerning the plausibility of individual axioms or rules.[2] What possible worlds semantics provides is a translation of an obscure intensional language into the pristine clarity of an extensional (meta)language.[3]

More restricted versions of necessity can be given a similar analysis by making the notion of possibility a relative matter:

($\Box$) $\Box A$ is true at a world w iff A is true at all worlds possible relative to w.[4]

For doxastic modality, for instance, worlds represent possible states of belief of an agent, and one state of belief w' is possible relative to another w just in case w' cannot be ruled out by what the agent believes, as determined by w. ($\Box*$) is equivalent to the special case of ($\Box$) when every world is possible relative to every other, so ($\Box$) provides a semantics for a broader range of modalities. I will call ($\Box$) the *simple quantificational analysis* of modality.

What makes the simple quantificational analysis so attractive is that it provides, at a schematic level, simple and uniform semantics for a very broad class of modalities. In addition, there is a

[1] There is a vast literature on the ontology of possible worlds. Some excellent sources include Lewis (1986a), Armstrong (1989), and Divers (2002).

[2] The modal logic of the simplest quantificational analysis is called **S5**.

[3] See Routley and Meyer (1977) for an argument against extensional reduction.

[4] Relative possibility is sometimes referred to as *accessibility*.

natural correspondence between properties of relative possibility and modal validities. For example, if relative possibility is reflexive (i.e. if every world is possible relative to itself), then $\Box A$ semantically entails A, and *vice versa*. Thus, once we settle the properties relative possibility has, we settle the modal validities along with them.

Despite its advantages, the simple quantificational analysis has serious drawbacks. Perhaps the most widely discussed is that both possibility and necessity are closed under strict implication, which means that the following holds:

(K) $\Box(A \to B) \to (\Box A \to \Box B)$.

This property is sometimes referred to as *logical omniscience*.[5] The problem is that one may know or believe a proposition without knowing everything that necessarily follows from it. Just because you know the axioms of Peano Arithmetic, it does not mean that you know everything that follows from them. Moreover, since necessary truths are strictly implied by everything, the simple quantificational analysis yields that every necessary truth is known. While this may hold for the most idealized form of knowledge, it fails for any interesting notion that is the object of philosophical analysis. Relatedly, what would it mean to be morally obligated to see to it that Justin Trudeau be human? And can't one have inconsistent beliefs? If so, the simple quantificational analyses falls short as an analysis of knowledge, morality, agency, belief, and other notions for which closure under strict implication is highly implausible.

However, the fact that it falls short for certain intensional notions does not mean that an analysis in terms of possible worlds is unworkable. One way to achieve such an analysis is to associate to each world a set of propositions that are necessary relative to that world. Clearly assigning to each world a set of worlds that are possible relative to that world suffices to assign a family of propositions necessary there (namely, the set of propositions true at each relatively possible world), but as we have seen, this closes the family under potentially undesirable properties. Let N be a function that takes a world and yields the set of propositions necessary there; thus $N(w)$ is the set of worlds necessary relative to w. Let $\|A\|$ denote the set of worlds at which A is true, i.e. the proposition expressed by A. Then

(NS) '$\Box A$' is true at a world w iff $\|A\|$ is a member of $N(w)$.

We can see that necessity is no longer closed under strict implication, since both $\|A \to B\|$ and $\|A\|$ can be in $N(w)$ without $\|B\|$ being in $N(w)$. Indeed, necessity is not closed under much at all until some conditions are placed on N.

The main drawback of this sort of "analysis" is that it provides a poor explanation of when a sentence is necessary—it says it is necessary just in case the proposition it expresses is, which is true but trivial. We could, of course, provide a philosophical interpretation of N and corresponding analysis, but that will not make (NS) any less trivial. Moreover, determining which properties N has will depend on determining in advance which sentences are to be valid. So, unlike ($\Box$), the theorems are not derived from intuitive semantical properties, but instead the axioms determine what the semantical properties are.[6]

It is standard to assume that possible worlds satisfy the following two properties:

[5] $\Box(A \to B)$ expresses that A strictly implies B.

[6] The semantics given by (NS) is called *neighborhood* or *Scott-Montague* semantics. It can be seen as a generalization of the simple quantificational analysis, though itself not being a quantificational analysis. There are other ways of generalizing ($\Box$) while retaining the use of worlds that have not been mentioned. One is to distinguish two kinds of worlds, viz. so-called normal and non-normal or impossible ones, or to let relative possibility be a relation of more than two places, a strategy employed in relevance logic. See e.g. Priest (2008) for details.

Consistency: a world *w* is *consistent* iff there is no collection Γ of sentences true at *w* such that Γ (in some given sense) entails a contradiction;

Maximality: a world *w* is *maximal* iff for every sentence *A*, either *A* or its negation is true at *w*.

It is important to keep in mind that the notion of entailment that figures in consistency need not be logical. For instance, a world that makes true both '*x* has mass *m* kg' and '*x* has mass *n* kg' for distinct *m* and *n* will, given the relevant sense of entailment, entail a contradiction. At a certain level of abstraction, consistency and maximality are the only properties we need care about. In providing a semantics for certain bits of language, for example, it will not matter whether worlds are concrete or abstract entities, or whether they are built up from properties or propositions or something else entirely. For instance, we could let a world be a class of atomic facts, and say that an atomic fact is true at a world just in case that fact is a member of the world. The truth of complex facts is given as usual; a conjunctive fact is true (at a world) just in case each conjunct is, a negative fact is true just in case the negand isn't, and a universal fact is true just in case each instance is. Then, on the essential assumption that any collection of atomic facts is consistent (which one may reasonably deny), it is easy to show that a world as a set of atomic facts is maximal consistent.[7] A variety of other constructions will similarly yield the same result.

In what follows I illustrate, using three central examples, why worlds taken as maximal consistent objects cannot play the role they were initially assigned to play. That role includes giving an analysis, even a mere truth conditional semantics, for (i) modality, (ii) belief, and (iii) conditionals. We could list other examples that give compelling reasons for supplanting possible worlds with something better suited to the task, but I think these three examples, which figure prominently in the literature, do well to illustrate the limitation of possible worlds traditionally construed and their role in philosophical theorizing.

3 Counterpart theory and possibilities

Modal realism is, roughly, the view that modal propositions are grounded in the existence of concrete, non-actual individuals and worlds. It is given its fullest defense by David Lewis (Lewis, 1986a). According to modal realism,

Plenitude: for any way the world (or a part of it) could possibly be, some world (or a part of it) is.

A world, as Lewis defines it, is a mereological sum of spatiotemporally- (ST-) related individuals, satisfying the condition that if *w* is a world, *x* is part of *w* and *x* and *y* are ST-related, then *y* is part of *w*. If ST-relations are non-modal, as is very plausible, then modal realism provides a reductive analysis of modality, a feature that is touted as one of its main virtues.

According to a standard possible worlds analysis, a *de re* modal statement such as 'Hilary Clinton could have won the election' is true just in case there is a world where Clinton—she herself and not some simulacrum—wins the election. According to Lewisian modal realism, things exist in precisely one world, so if they are not to have all their properties necessarily, *de re* modal claims cannot be given the standard analysis.[8] Lewis proposes instead what he calls

[7] The use of 'maximal consistent' instead of 'maximally consistent' is to avoid the confusion that 'maximal' is modifying 'consistent'.

[8] Since Lewis accepts unrestricted mereological summation, there are individuals that do not wholly exist in one world, but have parts from different worlds. We can ignore such individuals since they play no role in Lewis's analysis of modality.

counterpart theory, according to which the statement 'Clinton could have won the election' is true just in case there is a world in which a counterpart of our actual Clinton wins the election. What is a counterpart of Clinton? It is someone who sufficiently resembles her in the relevant (i.e. contextually determined) respects, e.g. someone with a similar history to Clinton's, someone who physically looks like Clinton, and so on. The counterpart relation therefore serves as a more flexible substitute for the usual relation of (transworld) identity, more flexible because it need not be transitive, symmetric, or one-one. In particular, an individual may have multiple counterparts within a single world.[9]

Whether a world represents of an individual that it have such and such properties therefore depends not just on the counterpart relation determined by the context, but also on which counterpart of the individual we choose to do the representing. For instance, suppose a world w contains two identical twins, each of whom is a counterpart of Jane. Then if one is born at twelve o'clock and the other at a quarter past twelve, it cannot be possible that Jane both is and is not born at twelve o'clock. If the counterpart relation is purely qualitative, as Lewis argues, then one can even have multiple counterparts within their own world.[10] This is important because it shows that counterpart theory (as the qualitative counterpart theorist conceives it) cannot identify possibilities, not even maximal ones, with possible worlds. Rather, a possibility must be identified with something finer-grained, such as a pair consisting of a world and a counterpart function f mapping individuals to at most one counterpart per world. For example, in the case of Jane, if we call the younger twin Molly and the elder Holly, and f and g are counterpart functions mapping Jane to Molly and to Holly respectively, then $\langle w, f\rangle$ and $\langle w, g\rangle$ represent two distinct possibilities for the same individual in the same world.

Why does this matter? First, and as Lewis himself notes, this marks a break from "established theory" according to which (maximal) possibilities just are possible worlds. Lewis later makes the break by dropping the assumption that no two worldmates can be counterparts, but that break was made at the inception of counterpart theory by allowing things to have multiple counterparts in the same world.[11] Second, it shows that worlds by themselves cannot do the work the theory needs them to do. Worlds alone can neither serve as possibilities nor can sets of them serve as the contents of propositions, which are two of their primary roles.

The fact that possibilities cannot be worlds has reaching implications. By way of example, consider the doctrine known as *haecceitism*, which as Lewis puts it is the claim that two *possibilities* may differ in what they represent *de re* concerning an individual without thereby differing qualitatively. Lewis rejects haecceitism for a number of reasons, all of which rely on the crucial assumption that what does the representational work on a standard possible worlds account are worlds *and worlds alone*.[12] But given the fact that possibilities need to be something such as world-counterpart-function pairs, Lewis turns out committed to the doctrine. Very briefly, the reasoning goes roughly as follows. First, we need to say when two possibilities differ qualitatively. Most plausibly, we should say that two world-counterpart-functions pairs

[9] Allowing for multiple counterparts within a world invalidates the necessity of identities: that if two things are identical, they are necessarily identical. See the translation scheme of Lewis (1968) for details, and for Lewis's original presentation of the theory.

[10] Consider a world consisting of only two qualitative duplicates.

[11] See (Lewis, 1986a, §4.4) for Lewis's motivation for dropping the assumption. The break from established theory is made as early as Lewis (1968) and not in the much later (Lewis, 1986a, §4.4), as Lewis suggests. It is also made by allowing multiple counterpart relations relative to a single context, a strategy already employed in Lewis (1971) for dealing with puzzles of coincidence. Consider a statue and the coinciding lump of clay that constitutes it and suppose, as Lewis does, that they are identical. If the expressions 'statue' and 'lump of clay' evoke different counterpart relations even relative to the same context, then one and the same world can represent differently concerning one and the same individual via multiple counterpart relations.

[12] See (Lewis, 1986a, §4.4) for Lewis's attack on haecceitism.

$\langle w, c\rangle$ and $\langle w', c'\rangle$ differ qualitatively if their world components do *or* their second components map the same thing to qualitatively distinguishable individuals. Next, suppose our world, call it '@', is one of one-way eternal recurrence. Then we have intrinsic duplicates and hence counterparts in each epoch, even though we ourselves inhabit exactly one of them.[13] Let $\langle @, c\rangle$ and $\langle @, c'\rangle$ be such that c is the identity function and c' is just like c except that it maps me to a qualitatively indiscernible other-epoch worldmate. Then these possibilities are qualitatively identical even though they differ in what they represent concerning me: one represents that I inhabit one epoch, and the other represents of me that I inhabit another. In other words, a qualitative counterpart theorist such as Lewis is committed to geniune haecceitism (and not just his "cheap substitute") precisely because possibilities need to be played by entities having a richer structure than just worlds.

4 Belief and centered worlds

Belief has posed a problem for the simple quantificational analysis for a variety of reasons, including the following main ones:

1. one can have inconsistent beliefs;
2. belief is not closed under strict implication;
3. first-personal or *de se* belief poses a unique challenge.

We have discussed the first two problems and have seen one way of at least formally dealing with them, i.e., in terms of (NS), but even given the immense flexibility such a framework provides, many are convinced that it is still unable to capture the distinctive feature of *de se* belief.[14]

Indexical expressions are those that are context-sensitive. This includes expressions such as 'I', 'you', 'here', and 'now'. A sentence containing an indexical expression cannot be assigned a truth-value until a context supplies a value for the indexical. Once those values are supplied, the sentence is assigned a content, that is, the proposition expressed by the sentence relative to the context. Thus, 'Jane is here' expresses different propositions relative to contexts where 'here' refers to different places. If 'here' refers to her house then the proposition expressed is <Jane is at Jane's house>, and if it refers to the skating rink the proposition expressed is <Jane is at the skating rink>.

Now suppose Jane and John have gone for a hike in the woods and are confronted by a bear who begins to chase Jane, and suppose she yells 'I am being chased by a bear!' Then according to the traditional account, the proposition expressed by her utterance is <Jane is being chased by a bear>. If the contents of beliefs are propositions, as is commonly assumed, then John and Jane share the same belief concerning the situation in question; for they both believe that Jane is being chased by a bear. However, if what explains their behavioral differences (as only Jane is climbing a tree) is a difference in belief, then John and Jane do not share the relevant belief after all. Many take the difference in behavior to be explained by Jane's distinctively *de se* belief that

[13] Lewis uses the example of duplicate worlds of one-way eternal recurrence to show that qualitatively identical, overlapping worlds can exhibit haecceitistic differences when the counterpart relation is identity. See (Lewis, 1986a, p. 228).

[14] Concerning inconsistent beliefs, (NS) clearly allows for them, but every outright contradiction is still represented by the empty set. The centered worlds analysis of *de se* content discussed below has the advantage that one can have inconsistent beliefs in numerously distinct ways: one can, e.g., believe they are someone else, or that they lack a property they necessarily they have.

she herself is being chased by a bear, versus John's merely *de re* belief that Jane is being chased by a bear.[15]

What Lewis proposes is to treat belief as the self-ascription of a *property*, rather than the belief in a content understood as a set of worlds. Call a *centered world* a pair consisting of a world and an individual on which the worlds is centered. To simplify matters, let us assume with Lewis that individuals are worldbound (i.e. exist at precisely one world), so that a centered world can be represented by a single individual, i.e. its center. (Centered worlds can be thought of as world-indexed individuals.) Then call a doxastic alternative of an agent an individual the agent cannot rule out as being herself, and say that an agent believes that she has property ϕ iff all her alternatives have ϕ. Instead of treating *de se* differently from *de re* or *de dicto* belief, we can say, for any A whatever, that an agent believes that A just in case each of her alternatives inhabits an A-world. Finally, to go back to our example involving Jane and John, it is clear that they no longer share a belief, for only Jane's alternatives are being chased by a bear, and this difference in belief can be used to explain their difference in behavior.

A proposition as a set of worlds can be equally well represented as a set of individuals, so the centered worlds analysis is not committed to two types of contents. For every set of worlds there corresponds the set of individuals each of which inhabits one of those worlds: say that such a set of individuals is true at a world just in case it contains an inhabitant of that world. On the other hand, not every *de se* content can be equally well represented as a set of worlds. So what to say about the truth of a *de se* content that corresponds to no set of worlds? Is singleton {Jane}, for instance, true or false? The question needs answering because we need to know when someone has a true belief about oneself, and not just when it is true that one believes something about oneself. The question does have an answer, but it can only be given once we supply an individual relative to which the content can be evaluated. {Jane} is true only relative to Jane, since only she can have the true, first-personal belief <I am Jane>. *De se* contents, then, do not stand alone in the same way sets of worlds do, which is a sign of their irreducibly indexical nature.

The role that worlds traditionally played as maximal consistent doxastic states has been assumed by individuals in order to capture the distinctively *de se*. There simply is no way to account for the difference in Jane and John's behavior in terms of a proposition understood as a set of worlds. There are other ways of accounting for the difference that allows us to hang onto the traditional view about propositions, but we will have to leave matters here.

5 Conditionals and impossible worlds

Consider the following pair of conditionals:

1. If Oswald didn't shoot Kennedy, then someone else did (true);
2. If Oswald hadn't shot Kennedy, then someone else would have (false).

Given our knowledge of the facts, viz. that Kennedy was shot, the first is true. But the second seems false, for we know of no other possible shooters besides Oswald. This difference in truth value between the two conditionals implies a difference in meaning of their respective conditionals, the first being referred to as the *indicative* and the second as the *counterfactual*.

[15] One might think that some other difference in belief is responsible for the difference in behavior, but we could assume that Jane and John share *all* of their beliefs, if none of them are distinctively *de se*. Moreover, we need not assume more generally that behavior is to be explained in terms of belief-desire psychology. See Ninan (2016) for a defense of the *de se* from skeptics.

Both kinds of conditionals are in need of analysis as it is clear that classical material implication fails to capture the meaning of either.[16] We will focus here on the counterfactual. The most natural intensional conditional we could define in a basic modal language is the strict conditional, but there is good reason for thinking that the counterfactual, which is also intensional, is not a strict conditional. In particular, the strict conditional satisfies antecedent strengthening:

(AS): $\Box(A \to B) \to \Box((A \wedge C) \to B)$,

since if B is true in all A-worlds, it must also be true in all A-worlds that are also C-worlds. The counterfactual, on the other hand, seems not to. For consider the following sequence:

1. if Otto had come, it would have been alively party, but
2. if both Otto and Anna had come it would have been a dreary party, but
3. if Waldo had come as well, it would have been lively...

Sequences sharing this pattern are called *Sobel sequences* and have attracted considerable attention.

What Stalnaker and Lewis proposed independently and at around the same time was roughly the following analysis:

($>$): $A > C$ is true at a world w just in case C is true at all the A-worlds *closest* to w.[17]

This analysis has a number of important features. First, it is easy to see that antecedent strengthening fails, as do a number of other properties of the strict conditional that are intuitively invalid for the counterfactual. Second, and what will be the most important to us, is that any counterfactual with an impossible antecedent—a *counterpossible*—is necessarily true, a property ($>$) shares with the strict analysis of the conditional. However, it is intuitively false, e.g., that if Anaxagoras had squared the circle, nobody would have squared the circle. For this reason, many have rejected the Stalnaker-Lewis analysis, at least without some further amendment.

One obvious way to circumvent this problem is by admitting impossible worlds, so that the antecedent of a counterpossible can be true at an impossible world without the consequent also being true there.[18] This amendment requires no change to ($>$), only a broadening of the class of worlds. The toughest challenge facing such an account concerns what to say about closeness now that impossible worlds are in the picture: e.g., are possible worlds always closer than impossible ones to possible worlds?[19]

What sort of entity would an impossible world be? Most construct them from fairly uncontroversial entities (e.g. sentences, propositions, or states of affairs).[20] Some believe them to be real or concrete worlds not unlike our own (ontologically speaking).[21] Others who accept the importance of impossible worlds reject the view that possible and impossible worlds ought to be the same sort of entity, ontologically speaking.[22]

[16] For arguments to the contrary, see e.g. Grice (1989), Jackson (1979), and Lewis (1986b).

[17] See Stalnaker (1968) and Lewis (1973). For discussion concerning the relation of closeness between worlds, see Lewis (1979).

[18] See e.g. Goodman (2004), Nolan (2013), Nolan (forthcoming), and Brogaard and Salerno (2013) on impossible worlds amendments to the Stalnaker-Lewis account.

[19] See Nolan (1997) for discussion concerning closeness between possible and impossible worlds.

[20] See e.g. Laan (1997).

[21] See e.g. Yagisawa (2010).

[22] See e.g. Berto (2009).

It would be strange to believe that the Stalnaker-Lewis account gets it right for counterfactuals with possible antecedents (and possibly false consequents, which pose a similar problem), but that it fails terribly otherwise. That is, it would be strange to believe that the conditional in counterpossibles is somehow different in meaning from the conditional in other counterfactuals. Thus any broadly Stalnker-Lewisian account of the counterfactual (i.e. any variably strict analysis) will need to employ impossible worlds.[23] Since the Stalnaker-Lewis account is the best on offer, there is good reason for thinking that impossible worlds are as central as possible ones in our understanding of counterfactual reasoning.[24]

6 Conclusion

We have looked at three important areas where possible worlds have been supplanted, or at least supplemented, by entities better suited to the task originally assigned to worlds. There are at least two other important areas we have not discussed which include the incomplete situations of situation semantics, and the propositions of two-dimensional semantics.[25]

The conclusion to draw from this is not that possible worlds have no important role to play in the analysis of modality, belief, and so on, but that the simple analyses of these notions that made possible worlds semantics initially attractive required, and continues to require, further refinement in light of the complexities exhibited by the phenomena under analysis.

References

Armstrong, D. M. (1989). *Universals: An Opinionated Introduction*. Focus Series. Westview Press.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. Bradford Books, Cambridge.

Berto, F. (2009). Impossible worlds and propositions: against the parity thesis. *The Philosophical Quarterly*, 60(240):471–486.

Brogaard, B. and Salerno, J. (2013). Remarks on counterpossibles. *Synthese*, 190:639–660.

Carnap, R. (1946). Modalities and quantification. *Journal of symbolic logic*, 11:33–64.

Divers, J. (2002). *Possible Worlds*. Routledge.

García-Carpintero, M. and Macia, J., editors (2006). *Two-dimensional semantics*. Oxford University Press.

Goodman, J. (2004). An extended Lewis/Stalnaker semantics and the new problem of counterpossibles. *Philosophical Papers*, 33(1):35–66.

Grice, P. (1989). *Studies in the way of words*. Harvard University Press.

[23] Of course, one could also provide a pragmatic defense of vacuously true counterpossibles similar in spirit to the Gricean defense of the material conditional.

[24] The use of impossible worlds is nothing new in the semantics of conditionals. They are employed, e.g., in the semantics of relevance logic, where a counterexample to the irrelevant $(A \wedge \neg A) \rightarrow B$ is needed. See (Priest, 2008, §9.7).

[25] See Barwise and Perry (1983) on situation semantics, and García-Carpintero and Macia (2006) on two-dimensional semantics.

Jackson, F. (1979). On assertion and indicative conditionals. *The Philosophical Review*, 88(4):565–589.

Laan, D. V. (1997). The ontology of impossible worlds. *Notre Dame Journal of Formal Logic*, 38(4):597–620.

Lewis, D. (1968). Counterpart theory and quantified modal logic. *The Journal of Philosophy*, 65(5):113–126.

Lewis, D. (1971). Counterparts of persons and their bodies. *The Journal of Philosophy*, 68(7):203–211.

Lewis, D. (1973). *Counterfactuals*. Harvard University Press, Cambridge, Massachusetts.

Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*, 13(4):455–476.

Lewis, D. (1986a). *On the plurality of worlds*. Blackwell Publishing.

Lewis, D. (1986b). Postscript to "Probabilities of conditionals and conditional probabilities". In *Philosophical papers*, volume II, pages 152–156. Oxford University Press.

Mates, B. (1968). Leibniz on possible worlds. In van Rootselaar, B. and Staal, J., editors, *Logic, Methodology and Philosophy of Science*, volume 3, pages 507–529. North-Holland Publishing Company, Amsterdam.

Ninan, D. (2016). What is the problem of *de se* attitudes? In García-Carpintero, M. and Torre, S., editors, *About Oneself: De Se Attitudes and Communication*, pages 86–120. Oxford University Press.

Nolan, D. (1997). Impossible worlds: a modest approach. *Notre Dame Journal of Formal Logic*, 38(4):535–572.

Nolan, D. (2013). Impossible worlds. *Philosophy Compass*, 8(4):360–372.

Nolan, D. (forthcoming). Causal counterfactuals and impossible worlds. In Beebee, H., Hitchcock, C., and Price, H., editors, *Making a Difference*. Oxford University Press.

Priest, G. (2008). *Introduction to non-classical logics: from ifs to is*. Cambridge University Press, second edition.

Routley, R. and Meyer, R. K. (1977). Extensional reduction i. *Monist*, 60(3):355–369.

Stalnaker, R. C. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in logical theory*, pages 98–112. Blackwell Publishers, Oxford.

Yagisawa, T. (2010). *Worlds and individuals, possible and otherwise*. Oxford University Press.