

Chapter 1

Introduction

This book is long. It could have been somewhat shorter. But it couldn't have been just one sentence long.

This is a recurring pattern: for almost any familiar object, one can find respects in which it could have been somewhat different, but could not have been radically different. For example:

- (i) The Great Pyramid could have been a little bit smaller. But it couldn't have been as small as a thimble.
- (ii) The *Mona Lisa* could have been slightly different as regards its spatial distribution of colours. But it could not have had the spatial distribution of colours that is actually exemplified by Edvard Munch's *The Scream*.
- (iii) The Vienna Circle could have had somewhat different members. But it couldn't have been such that the only people who ever belonged to it were Sigmund Freud, Arnold Schoenberg, and Franz Kafka.
- (iv) The game of chess could have been generally played according to somewhat different rules: for example, it could have been generally played without the rule whereby a thrice repeated position leads to a draw, or without the *en passant* rule. But it could not have been generally played according to the rules of Twister.
- (v) The table before us could have been originally made using somewhat different parts: for example, any one of its legs could have been different. But this very table could not have been originally made of a completely different collection of parts.

These are what tradition calls '*de re*' modal claims: they have to do with what *specific* things could have been like, not just with what general *sorts* of thing there could have been. For example, we are claiming that the Great Pyramid

itself could have been somewhat smaller, not just that there could have been *some pyramid or other* (in the same place, made by the same people...) that is somewhat smaller than the Great Pyramid actually is. Unlike those other claims, ours implies that *there is something* that could have been somewhat smaller than it in fact is. Likewise, when we say that the Great Pyramid couldn't have been thimble-sized we mean that *that very pyramid* couldn't have been thimble-sized. The ancient Egyptians could of course have placed *a* thimble-sized pyramid in the Valley of the Kings right around where the Great Pyramid actually stands; but it would not have been the Great Pyramid itself.

Our primary concern in this book will be with some arguments that threaten to undermine these obvious-looking judgements. These arguments purport to show that if the objects in question are *tolerant*—capable of being somewhat different in the relevant respects from the way they in fact are—they are also *hypertolerant*—capable of being vastly different in these respects from the way they in fact are.

The motivating thought in these arguments is that the fact that the relevant objects are tolerant—assuming it *is* a fact—doesn't seem to be a mere accident. It is hard to believe that it's even *possible* for there to be a pyramid for which a certain slightly smaller size was impossible, or a table that couldn't have been made of certain slightly different parts. There is thus pressure to think that it is *necessary* that the objects in question are tolerant, in the relevant respect. But since modest differences can add up vast differences, being necessarily tolerant seems to entail being hypertolerant. For example, suppose the Great Pyramid is necessarily tolerant in the following sense: necessarily, however tall it is, it could have been 10% shorter. Since it is in fact approximately 134 m tall, it could have been approximately 121 m tall. But if it were 121 m tall, it would then be possible for it to be ten per cent shorter than *that*, i.e. approximately 109 m tall. So it is possible for it to be possible for it to be 109 m tall; but from this, it seems to follow that it is just *possible* for it to be 109 m tall. Multiple applications of this mode of reasoning will force us to say that the Great Pyramid could, after all, have been thimble-sized.

These arguments can be fitted to a general schema. Each begins with a *Tolerance* premise that says that a certain thing or certain things could have been slightly different in a certain way or range of ways. Each then adds a *Non-contingency* premise according to which the Tolerance premise is necessarily true if true at all. With the necessitation of the Tolerance premise in place, one can use certain principles of modal logic—most interestingly, an *Iteration* premise that says that what is possibly possible is possible—to de-

rive a *Hypertolerance* conclusion, that says that the thing or things could have been *arbitrarily* different in the given respect. Many arguments of this form are puzzling in the way that many well-known philosophical arguments are: the premises are rather plausible, but the conclusion seems unacceptable. While the label ‘paradox’ may be a bit strong, these arguments at least present us with an intriguing class of *puzzles*, since it is not obvious what to say about them. Should we accept the conclusion that the objects in question are hypertolerant? And if not, which premise should we deny?

While these “Tolerance Puzzles” may initially seem like mere curiosities, they offer an entry point into an array of deep and challenging metaphysical questions having to do with modality. This book will be structured as an exploration of various options for addressing Tolerance Puzzles. But along the way we will be investigating most of the questions about modality that have taken centre stage in metaphysics since around 1970.

In the last couple of decades, however, many metaphysicians seem to have been losing interest in these questions. Some herald an era of ‘post-modal’ metaphysics, in which modal concepts will no longer play a central role in setting the metaphysical agenda.¹ Sometimes, the thought is that modal notions are too *superficial* for debates involving them to be worth spending much time on.² Sometimes, the complaint is that the modal questions are too *coarse*, so that the answers to them fail to settle the answers to the really important “hyperintensional” questions in the vicinity. In some respects, then, we are swimming against the tide. But we make no apologies. The focus on questions involving modality brought a lot of clarity and discipline to metaphysics, a field notorious for its tendency to degenerate into obscurity and unruliness. This flowering was made possible by the wide uptake of modal logic, especially in the wake of Kripke (1959, 1963), which created a field rich enough to be of interest not just within metaphysics but across a wide range of fields within and far beyond philosophy. Whether or not one finds the modal questions interesting enough in themselves to deserve the central place in metaphysics they enjoyed in the late twentieth century, the new-found rigour that they made available should make them an obligatory part of any reasonable metaphysical education. And for what it’s worth, we think that the modal questions that will come up in the course of our exploration of Tolerance Puzzles are among the most interesting questions in metaphysics, neither too superficial nor too coarse to be set aside as somehow second rate.

¹See Sider 2020 (ch. 1).

²See Sider 2011 (ch. 11).

Although we are thus continuing a tradition, there is one important way that our work departs from its twentieth-century precursors: the formal language we use in regimenting claims and arguments is that of *higher-order* modal logic, which goes beyond the *first-order* modal logic that is central to Kripke's work. Higher-order languages are characterized by the availability of quantification into a wide range of different grammatical categories, not just into the category of 'singular terms' characteristic of first-order quantifiers. These languages provide a simple and rigorous way of regimenting informal talk of properties, relations, conditions, operations, and propositions or states of affairs, and thus serve as a helpful vehicle for expressing claims at the level of generality to which metaphysics tends to aspire. We don't mean to suggest that higher-order resources are indispensable for articulating Tolerance Puzzles; but as we will see, they are useful for capturing the general structure which the puzzles have in common, and in our discussions of particular strategies for solving Tolerance Puzzles higher-order regimentations will play an increasingly important role.

Higher-order languages are far from new: indeed, higher-order quantification is present already in Frege's *Begriffsschrift* (1879), perhaps the founding document of analytic philosophy. But thanks especially to Quine (1953a, 1970), higher-order formalization fell out of favour within metaphysics in the latter half of the last century, only quite recently making something of a comeback. Like first-order logic (modal and nonmodal), higher-order logic (modal and nonmodal) is a rich field, whose interest extends far beyond metaphysics, and indeed far beyond philosophy (e.g. into mathematics and computer science). We are thus hopeful that its broader uptake within metaphysics will bring further clarity and discipline to that field.

The Case for Tolerance

The chapters that follow offer a wide-ranging exploration of the main strategies for solving Tolerance Puzzles. After mapping out the argument more carefully and relating it to some other puzzles, we will spend several chapters each on the options of accepting Hypertolerance, denying Iteration, and denying Non-contingency. Of course, Tolerance Arguments can also be blocked by rejecting the Tolerance premises: for example, one might claim that the Great Pyramid couldn't have been even a little bit shorter than it in fact is, that the table couldn't have been made of even slightly different originating matter, etc. Logically speaking, it would have made sense to have included a chapter devoted to Tolerance-denial. But, while there is room for resistance to certain specific versions of Tolerance as they occur in

certain puzzles, we in general found the Tolerance premises so overwhelmingly plausible that we didn't have a chapter's worth of things to say about them. So in place of such a chapter, we will now briefly present our reasons for being sceptical of Tolerance-denial as a general strategy for handling Tolerance Puzzles.

The plausibility of Tolerance claims does not depend on any particularly philosophical sensibility. In many cases, there is a direct route to a Tolerance premise from certain very mundane beliefs that we could easily find ourselves expressing in everyday life. A carpenter, finding to her chagrin that a table she made doesn't quite fit through a doorway, might say: 'If only I had made the planks a little narrower, this table would have fit through here.' On the face of it, she is making a claim that entails that there is a table such that if she had made the planks a little narrower, that very table would have fit through the doorway. Similarly, consider 'If Neurath hadn't joined the Vienna Circle, it would have been much less influential.' Of course, these are counterfactual conditionals, not possibility claims. But in combination with other bits of mundane background knowledge they entail corresponding possibility claims: for example, if the table *would* have been small enough to fit through the doorway if the carpenter had made the planks a little narrower, then it certainly *could* have been small enough to fit through the doorway. And there is no need to appeal to counterfactuals, since possibility claims about particular objects are also common in our ordinary discourse. For example, a carpenter might say, 'I could have made this small enough to fit through the door, if only you had sent me the details', or a historian might say, 'Many versions of the castling rule in chess were proposed, any of which could easily have become standard for the game.'

Of course, these ordinary possibility judgements involve notions of possibility much more demanding than the very permissive status of *metaphysical* possibility that features in many Tolerance Arguments.³ But the inference from the claim that something could *easily* have been the case to the claim that it is metaphysically possible for it to be the case seems unproblematic.⁴

One might try resist the case for Tolerance from ordinary practice by al-

³For the standard way of introducing these statuses see Kripke 1972. We will have much more to say about what metaphysical necessity and possibility are, and how they relate to other kinds of necessity and possibility, in Chapter 8.

⁴And even if one somehow rejected that inference, one can as we shall see generate gripping Tolerance Puzzles using certain other modalities, such as nomic possibility and *having nonzero objective chance*, for which the inference from the ordinary 'could's to the ones that generate the puzzle is if anything even more straightforward.

lowing, for example, that ‘This table could have been less than four feet wide but isn’t’ has a true reading, while denying that this reading is *de re* in the sense of entailing ‘There is something that could have been less than four feet wide but isn’t.’⁵ Analogously, the most salient reading of ‘The president could have been Hillary Clinton but isn’t’ is true but doesn’t entail ‘There is someone who could have been Hillary Clinton but isn’t.’ It is a familiar observation that some expressions—such as ‘someone’, ‘every boy’, and ‘the president of the US’—generate systematic ambiguity when they share clauses with modal words like ‘could’, ‘possible’, and ‘likely’. For example, ‘Someone was likely to come in’ can mean ‘It was likely that someone would come in’ or ‘There was someone who was likely to come in.’⁶ However, “narrow scope” readings—those which block existential generalization—seem absent, or at least hard to access, for names and demonstratives. ‘John is in the car, and John could have broken his leg’ and ‘That guy in the car could have broken his leg’ seem unambiguously to imply ‘There is someone in the car who could have broken his leg.’⁷ Perhaps there are a few special cases where demonstratives in modal sentences admit readings which don’t allow existential generalization: consider ‘That could have been me’ (pointing to the lottery-winner on television). But the target sentences seem manifestly not special in that way. On the intended reading of ‘This table could have been narrower than it is’, the inference to ‘There is a table that could have been narrower than it is’ looks completely straightforward. In any case, there

⁵Leibniz suggests that his Tolerance-unfriendly brand of essentialism can be reconciled with various ordinary claims which seem to conflict with it by treating certain uses of proper names as definite descriptions: ‘as, for instance, when we mean by Adam the first man, whom God puts in a pleasure garden, which he leaves through sin, and from whose side God makes a woman’ (Leibniz 1956: 515f.). Elsewhere he routinely uses proper names as predicates: ‘I will now show you some [worlds], wherein shall be found, not absolutely the same Sextus as you have seen... but several Sextuses resembling him.’ (Leibniz 1710: §414).

⁶The most familiar (though not the only) approach to such ambiguities explains them using the same notion of “scope” that also covers the ambiguities that arise when multiple quantificational expressions share clauses with one another. For example, ‘At least one person reads every paper published in *Mind*’ can mean either ‘There is at least one person who reads every paper published in *Mind*’ or ‘Every paper published in *Mind* is read by at least one person.’

⁷The idea that proper names are just as accessible to existential generalization when they occur in the scope of modals as they are anywhere else is one of the central themes in Kripke 1972. Kripke adopts a theoretical framework (originated by Carnap 1947) that uses the ideology of “designation relative to a world”. In this setting, the claim can be put by saying that names designate the same thing relative to every world: they are “rigid designators”. But for those who, like us, are not particularly fond of this semantic ideology, it is important to see that the central thought can be stated in a far more neutral way. (See Yli-Vakkuri unpublished for a critical discussion of Kripke’s semantic ideology.)

is no real need to conduct the discussion using names and demonstratives at all: we can just as well get the puzzles going using quantified judgements like ‘Every table you made today could easily have been made a bit smoother by sanding it more carefully.’

In making such appeals to “how we talk”, we are not putting forward arguments of the form ‘The folk say that *P*; the folk are to be trusted as regards whether *P*; therefore *P*.’ We don’t need to conceive of ourselves as philosophical anthropologists, documenting the practices of the folk and worrying about whether it would be in some sense chauvinistic or presumptuous to override them. We are in the relevant sense among the folk; the relevant ordinary claims are ones that we are highly confident in, and take ourselves to know using the same methods that non-philosophers use. Appeals in this book to “ordinary practice” should be taken in this spirit.

One might be tempted to contrast our “conservative” methodology that places a premium on these ordinary beliefs with some imagined “revolutionary” alternative, that sets ordinary beliefs aside and approaches the subject without prejudice. But it is hard to imagine a sensible version of what such a methodology could actually look like. If at the outset of philosophical theorizing we set aside our confidence in the things we ordinarily take ourselves to know about the topic at hand, where is our confidence in the premises of philosophical arguments going to come from? In practice, an attempt to follow the path of the revolutionary seems likely to amount to a credulous embrace of whatever collection of philosophical arguments happens to first catch your eye. We don’t want to go so far as to put forth some general doctrine that ordinary platitudes can never be overturned by philosophy. But the more carefully one looks at the bits of philosophy that would be required to turn our puzzles into an argument against some ordinary Tolerance belief, the harder it is to imagine how one could reasonably regard the conjunction of the arguments’ premises as more compelling than the tolerance claims they attempt to undermine.⁸

Tolerance Puzzles arise from certain tensions between ordinary Tolerance beliefs and certain other *prima facie* plausible claims, among which are the claims that the relevant objects are not *hypertolerant* in the relevant respects: it seems to many of us that the Great Pyramid could not have been thimble-sized, that this very table could not have been made of some completely non-overlapping piece of wood, etc. But there is a salient contrast in strength between the pressure to accept Tolerance and the pressure to reject Hypertolerance, at least when both are interpreted as claims about

⁸Our attitude here is similar to Moore’s (1959) attitude to arguments for skepticism.

metaphysical possibility. The tendencies of thought that makes Hypertolerance strike us as odd do *not* seem particularly firmly embedded in ordinary practice. Of course, with a little Socratic questioning, one can elicit judgments from non-philosophers about all sorts of philosophical issues including questions of Hypertolerance. But such dispositions are notoriously weak and variable, and it would be disastrous to lump them together with the more spontaneous and mundane dispositions that underlie our Tolerance beliefs.⁹

True, we have many mundane pre-philosophical views to the effect that particular objects *would* not have been created under various counterfactual circumstances ('This table wouldn't have been made if I hadn't come across that particularly beautiful slab of oak'), or to the effect that they *could* not have been created without certain circumstances obtaining, in some quite demanding sense of 'could' ('The Great Pyramid could never have been constructed without access to an abundant supply of slave labour'). But whereas there is a plausible route from ordinary 'would have' and 'could feasibly have' claims to claims of metaphysical possibility, there is no straightforward route from 'wouldn't have' or 'couldn't feasibly have' claims to claims of metaphysical impossibility. Giving up Tolerance would require positing pervasive error right at the heart of our everyday practices, whereas accepting Hypertolerance for metaphysical possibility would make barely any difference. So, if we were convinced that there was no way to avoid the choice between giving up Tolerance and embracing Hypertolerance, we are inclined to think that accepting Hypertolerance would quite obviously be the way to go.¹⁰

Despite this asymmetry in the initial state of play, when philosophical neophytes are confronted with the puzzles, their first impulse is often to reject the Tolerance premise in favour of some quite radical view on which the objects in question couldn't exist at all without being exactly the way they actually are in the relevant respect. Perhaps this impulse isn't so sur-

⁹Moreover, with artful Socratic questioning one can also warm people up to many initially odd-sounding Hypertolerance claims. For example, we might get someone to concede that an unfinished book is still a book, and that a book that has only just been begun is an unfinished book. At that point, the claim that this book could have been just one sentence long may seem much more tempting.

¹⁰As we will see, Tolerance Puzzles can be raised using many different interpretations of possibility, including closer-to-home statuses like easy or feasible possibility as well as the broad status of metaphysical possibility. Hypertolerance claims involving the narrower modalities would be much harder to reconcile with our ordinary practices. However the main motivations for thinking that Iteration holds for metaphysical possibility do not carry over to these narrower modalities.

prising: as everyone who has taught philosophy knows, many philosophical neophytes start out with an affinity for radically revisionary views about all sorts of questions. But the impulse to reject Tolerance doesn't just spring from a taste for denying the obvious. It is also due, in part, to a bad argument which purports to derive intolerance from the logic of identity. The argument goes something like this: 'The table that would have been made if the carpenter had made the boards a little narrower is not exactly the same size as this table. So, the table that would have been made if the carpenter had made the boards a little narrower is not identical to this table. So, if the carpenter had made the boards a little narrower, this table would not have been made.'¹¹

The problem here is with the starting premise, not with the subsequent inferences. Since everything is exactly the same size as itself, and since this table is the one that would have been made if the carpenter had made the boards a little narrower, the table that would have been made *is* exactly the same size as this table. The truth in the vicinity of the first premise is that the table that would have been made if the carpenter had made the boards a little narrower *would not have been* exactly the same size that this table *actually* is. But that truth does nothing to support the claim that that this table wouldn't have been made if the carpenter had made the boards a little narrower, just as the fact that if you had drunk three cups of coffee you would have drunk more coffee than you did drink does nothing to support the claim that you couldn't have drunk three cups of coffee.

Note that the contrast between qualitative identity and numerical identity plays no role in this diagnosis: contrary to what is often taught to students, the problem is not that the first premise is true only if 'identical' means 'qualitatively identical' while the second step is true only if 'identical' means 'numerically identical'. Whatever qualitative identity is, everything is certainly qualitatively identical to itself, and so ' x is not qualitatively identical to y ' implies ' x is not numerically identical to y '.¹²

¹¹Look (2013) seems to make something like this argument, while also apparently conflating the principle we call Leibniz's Law (that identity requires having the same properties) with the Identity of Indiscernibles (its converse): 'This is indeed an important consequence of [Leibniz's] Principle of the Identity of Indiscernibles. When we speak of the Adam who brought sin into the world and an Adam who did not, we cannot, strictly speaking, be referring to the same individual.' Arguments of this kind are quite widely discussed and taken seriously in the temporal setting, as arguments for the claim that nothing can persist through change: see e.g. Gallois 2016.

¹²For evidence that some students at least are getting this faulty training, see <https://quizlet.com/71822821/personal-identity-concepts-and-theories-flash-cards/>: a flash card on 'Numerical vs. Qualitative Identity' reads, in part, 'Two things can be numeric-

A slightly different way of generating a spurious conflict between Tolerance claims and the logic of identity avoids the grammatical blunder of ignoring mood by instead appealing to things like possible worlds or situations: ‘The table in this situation is four feet wide, but the table in that situation is not four feet wide; so, the table in this situation is not the table in that situation.’ The argument has a superficial form that resembles good arguments, like ‘The table that John thumped is big, but the table Cian thumped is not quite so big; so the table that John thumped is not the table Cian thumped.’ But it also has a superficial form that resembles bad arguments, like ‘In Boston, Van is admired, but in Oxford, Quine is not admired, so Van is not Quine’, and ‘The apartment in the brochure is luxurious, and the apartment in the video isn’t at all luxurious, so the apartment in the brochure is not the apartment in the video.’ The last of these arguments is the best guide to what is going on with the argument about situations. Being four feet wide is compatible with not being four feet wide in some non-obtaining situation, just as not being luxurious is compatible with being falsely characterized as luxurious by some misleading brochure.

In fact, the very concept of identity is something of a red herring as far as questions of tolerance are concerned. We don’t need to mention identity at all to raise the questions. For emphasis we may ask ‘Could a somewhat smaller table be identical to this one?’ or ‘Could this table have been somewhat smaller while still being one and the same table?’; but nothing important is lost if we stick to ‘Could this table have been somewhat smaller?’.¹³ If we do drag in identity, we will need to be careful not to put things in confusing ways that make some invalid arguments look like mere applications of Leibniz’s Law. By sticking to identity-free formulations, we can avoid such dangers without effort.

Apart from these altogether confused arguments, various philosophers have given other, more sophisticated reasons for rejecting Tolerance. Leibniz was led by his “conceptual containment theory of truth” to the view that Adam ‘would not have been our Adam, but another Adam, had other events happened to him’ (Leibniz 1989: 73). Quine (1953*b,c*) is sceptical of the very *intelligibility* of all claims about de re modality, motivated apparently by the thought that such claims would have to be explained somehow

ally identical but qualitatively non-identical (a penny minted in 1914 and that same penny found in 2014).’ Similarly, in the context of discussing change, Gallois (2016) says that ‘it seems that *a* and *b* can be numerically identical without being qualitatively identical by having different qualities at different times.’

¹³Lewis (1986: 193) makes the same point: ‘We *do* state plenty of genuine problems in terms of identity. But we needn’t state them so.’

in terms of analyticity (a property of sentences), where such an explanation would require making arbitrary choices amongst the various expressions that denote some object. And Lewis (1968) is sometimes erroneously interpreted as holding that literally speaking no ordinary thing could have different properties from those it in fact has, on the grounds that this would require being part of many different “worlds” whereas in fact every ordinary thing is part of only one “world”.¹⁴ While interesting, these arguments are too idiosyncratic, or otherwise alien, for us to engage with here.

Of course, one could motivate the rejection of Tolerance simply from the Tolerance Puzzles themselves. If one found Hypertolerance repugnant enough, and the cases for Iteration and Non-contingency compelling enough, then one might feel compelled to reject Tolerance, even if one was initially well-disposed towards it. However such a reaction seems wrong-headed to us, since as we have already said, the case for Tolerance looks dramatically stronger than the case against Hypertolerance—and far stronger still than the combined case for Iteration, Non-contingency, and the denial of Hypertolerance. And the more time we have spent in the company of the puzzles, doing our best to articulate interesting arguments for elements of that trio, the more we have been confirmed in our sense that none of the arguments in the offing have any prospect of rising to the heady level that would be required to justify anything as revolutionary as a wholesale abandonment of Tolerance.

Philosophers who are more open than we are to the denial of Tolerance often try to make it seem more liveable by appealing to the distinction between strict and loose talk. The thought is that even if one holds that *strictly speaking* a certain object couldn't have existed without having all the properties it in fact has, one can still grant that various sentences whose *literal* truth would conflict with that claim—e.g. ‘Each of these tables could have been less than four feet wide’—are true in some loose, nonliteral sense.¹⁵ If such

¹⁴See Chapter 10 for a discussion of what we take to be Lewis's actual views.

¹⁵A canonical source for such contrasts is the discussion of change over time in Butler 1736. According to Butler, remarks like ‘The same tree has stood fifty years in this place’, while false ‘in a strict and philosophical manner of speech’, are nevertheless appropriate ‘in a loose and popular sense’. Similar ideas are in play in Chisholm (1976). In the modal case, in discussing Leibniz's view that any object having a certain property could not have existed without that property, Mondadori (1975) suggests that ‘clearly, if one is a super-essentialist, one will not be in a position to interpret de re modal predications and most counterfactual conditionals in a literal way... [they] should accordingly be interpreted, and made sense of, in such a way that they turn out not to mean what they actually say’. Mondadori suggests that counterpart theory be used to provide an account of the relevant non-literal use of these sentences (see also Cover and O'Leary-Hawthorne 1999: 115ff.).

a view could be sustained, it would indeed take some of the sting out of the denial of Tolerance. However, it is quite difficult to sustain. The challenge that must be met is not just that of describing a certain mapping from sentences to the propositions that they “loosely express”. One must also explain what stops the propositions in question from counting as being *literally* expressed by the sentences in question. And in the case at hand it is quite hard to imagine how this could go. To illustrate the difficulty, imagine that the mapping in question is specified in part by associating words like ‘could’ and ‘possible’ with a certain property of propositions which the proposition that this table is less than four feet wide does have, despite the fact that it’s (strictly speaking) not possible for it to be true. What grounds could there be for denying that the use of ‘could’ to attribute this status counts as a literal one? There is already good reason to think that ‘could’ admits a wide range of literal interpretations. For example, the ‘could’ in ‘I couldn’t have come to your party in any case since I had to visit my aunt’ is naturally interpreted differently from the one in ‘I could have come to your party but only at the cost of letting down my aunt.’¹⁶ It seems *prima facie* arbitrary and uninteresting to decree that among all these ordinary uses of ‘could’ some wide range of them don’t count as “literal”.¹⁷

In any case, an appeal to loose talk does not really make the puzzle go away, since the relevant mode of loose interpretation can be applied uniformly to the whole argument. Loose talk doesn’t prohibit deductive argumentation; standard deductive argument forms like Modus Ponens remain equally compelling when they are applied to derive loosely-used conclu-

The idea of rejecting the literal truth of Tolerance while employing counterpart theory to account for its nonliteral truth is also suggested by a discussion in Kripke 1972 (51, n. 18). Kripke suggests that there is some kind of puzzle arising from the fact that certain Tolerance claims about originating matter seem true while corresponding Hypertolerance claims seem false, and entertains a response on which “strict identity” applies only to fundamental objects, with “some sort of ‘counterpart’ notion” taking its place when the modal properties of objects like tables are in question. In Chapter 10 we will discuss the potential relevance of counterpart theory to Tolerance Puzzles, though it seems to us that it is more promising to combine it with the denial of Iteration than with the denial of the literal truth of Tolerance.

¹⁶The canonical defence of this view is Kratzer 1977.

¹⁷If you were focused on Tolerance claims expressed using words like ‘identical’ and ‘same’, you might think that you could characterize the loose use simply by positing a special loose interpretation for those words (following Butler 1736). By contrast with the case of ‘could’, the challenge to explain why a use of ‘identical’ to express a relation less demanding than the one philosophers normally use it to express would count as a nonliteral use does not seem so daunting. The answer might be analogous to the reasons why ordinary uses of ‘They are exactly the same height’ are plausibly classified as nonliteral. But as we emphasized above, there is no need to mention identity at all in stating Tolerance.

sions from loosely-used premises, so long as the same modes of loose interpretation remain in play throughout. So if we can validly deduce a Hypertolerance claim from loosely true premises, we would expect the Hypertolerance claim to be loosely true. But the repugnance of Hypertolerance is by no means driven by a sudden insistence on a strict and literal interpretation. 'Loosely speaking, the Great Pyramid could have been thimble-sized' is no more appealing than the plain 'The Great Pyramid could have been thimble-sized.' Of course, one could block such a deduction by claiming that some of the inference rules employed are not such as to preserve loose truth: but such logical revisionism will have similar costs whether we say that the propositions in play are literally or merely loosely expressed. In what follows, we won't have much more to say about the concept of loose use, which we have not found helpful in sharpening any of the interesting questions raised by Tolerance Puzzles.

Despite the appeal that Tolerance-denial holds for many of those coming across Tolerance Puzzles for the first time, the kind of super-essentialism that would be required to justify Tolerance-denial as a general strategy has very little going for it. As we will see, other approaches to the puzzles offer far richer philosophical rewards.

The Structure of This Book

Our goal is to provide a careful articulation of the whole family of Tolerance Puzzles, with an eye to their general underlying form, and to examine in detail various strategies for resolving them. As explained above, we have found the level of rigour afforded by formalization in higher-order logic to be helpful in navigating this difficult terrain. So we will begin, in Chapter 1, with a systematic presentation of the formal tools that we will be relying on throughout the rest of the book. Some readers may prefer to skip this chapter on a first reading; they may come back to it later if it at some point they find themselves concerned by a potential ambiguity in one of our informal statements, or drawn to a position that accepts the premises but not the conclusion of some argument we are treating as valid. We hope that Chapter 1 will also be able to stand independently of the rest of the book as a primer on higher-order modal logic, written with an eye to using it in metaphysics, as opposed to targeting it as an object of mathematical study.

Chapter 2 lays out in detail the structure of the Tolerance Arguments that are our central topic, presenting a general schema under which various examples can be subsumed. The schema can be instantiated not only with many different parameters of variation but with many different inter-

pretations of the relevant notion of possibility. Interpretations of interest range from mere metaphysical possibility to more demanding statuses like *having a nonzero objective chance* and *being true at some time*. Chapter 3 says more about the motivations for the crucial Non-contingency premises, contrasting certain forceful motivations with others which, while tempting, are shown to be untenable by the Sorites Paradox. Chapter 4 lays out a different family of “Coincidence Puzzles”, which have sometimes been discussed alongside Tolerance Puzzles. We agree it is helpful to consider the two families of puzzles side by side, since some strategies for resolving Tolerance Puzzles generalize fairly easily to Coincidence Puzzles while others do not.

Our exploration of the strategies for solving Tolerance Puzzles begins in chapters 5 and 6, which explore the option of accepting Hypertolerance in many or all of the puzzles. Chapter 5 discusses some arguments against Hypertolerance which have been influential in the literature but which we don’t find compelling, while Chapter 6 focuses on what seems to us to be the most forceful challenge to Hypertolerance, based on certain physicalistic supervenience principles. Chapters 7 and 8 turn to the strategy of denying the Iteration premise. Chapter 7 argues that denying Iteration for metaphysical possibility does not provide a sufficiently general solution to the full range of Tolerance Puzzles; Chapter 8 argues that because metaphysical possibility is the *broadest* form of possibility, Iteration holds for it. Chapter 9 focuses on one particular kind of Tolerance Puzzle where there are special barriers to both Iteration-denial and Hypertolerance, namely those in which the operative modality is something explained in terms of *objective chance*. Chapter 10 takes up a cluster of ideas from the literature that goes under the name of “counterpart theory” and has been widely thought of as offering some distinctive help with Tolerance Puzzles (often via Iteration-denial). Our view is that counterpart theory is a red herring.

Chapter 11 finally unveils what we take to be the most promising general strategy for resolving Tolerance Puzzles (as well as Coincidence Puzzles), one that combines a plenitudinous ontology of material objects with semantic shiftiness in various expressions that figure in the puzzles (e.g. ‘The Great Pyramid’, ‘chess’, and ‘table’). We explain how this picture undermines the central motivation for Non-contingency that emerged from Chapter 3. Chapter 12 refines the strategy, arguing that a mixed approach that sometimes accepts Hypertolerance and sometimes rejects Non-contingency is better than a uniform treatment; we also survey some interesting further choice points. Chapter 13 considers some nearby alternatives that appeal to shiftiness in other relevant expressions (e.g. ‘possible’), and also discusses what we see as the most pressing challenges to our favoured

approach.

Finally, chapters 14 and 15 focus on a special class of Tolerance Puzzles which are new to the literature and raise special challenges that are not straightforwardly addressed by the previous discussion. These puzzles turn on a narrow notion of “indiscernible possibility” on which *qualitative* truths are automatically necessary; this makes for a distinctive new motivation for Non-contingency, based on the qualitiveness of certain properties (like *being a table*). Chapter 14 lays out these puzzles and explores the options left open when the relevant qualitiveness premises are accepted, and Chapter 15 presents our favoured approach to the new puzzles, which involves denying that properties like *being a table* are qualitative, and explores some further questions for it using the ideology of “plural aboutness”.

There are several paths through this book in addition to the most obvious approach of reading it from cover to cover. Chapters 2, 3, and 11 are the core chapters: readers who just want to know what the central puzzles are and how we propose to solve them can just read these three chapters. We should warn such readers that, as with all workable strategies for resolving these puzzles, our approach has some uncomfortable and disconcerting features. To fully appreciate why we have made our peace with these surprises, one will need to absorb the costs of alternative strategies. So we recommend that readers not initially taken with our positive view supplement the three core chapters with those chapters that engage with whichever alternative approach they find most attractive. On the other hand, those who are inclined to regard our view as obviously correct should have a look at the discussion of objections in Chapter 13. Other readers may be particularly interested in our background agenda of demonstrating the utility of higher-order logic for bringing rigour to metaphysics. These readers should certainly not skip Chapter 1; in addition to the core chapters, they will find material of special interest in chapters 7 and 8 (dealing with Iteration) and in chapters 14 and 15 (on qualitiveness).

References

- Butler, Joseph (1736). 'The Analogy of Religion, Natural and Revealed, to the Constitution and Course of Nature', in *The Works of Bishop Butler*, ed. David E. White (new edition, Woodbridge: Boydell and Brewer, 2006): 147–314.
- Carnap, Rudolf (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic* (Chicago: University of Chicago Press).
- Chisholm, Roderick (1976). *Person and Object: A Metaphysical Study* (London: Routledge).
- Cover, J. A. and O'Leary-Hawthorne, John (= Hawthorne, John) (1999). *Substance and Individuation in Leibniz* (Cambridge: Cambridge University Press).
- Frege, Gottlob (1879). 'Begriffsschrift: A Formula Language, Modeled Upon That of Arithmetic, for Pure Thought', in *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931*, ed. and trans. Jean van Heijenoort (1967): 1–82 [Ger. orig., *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens* (Halle: Verlag von Louis Nebert, 1879)].
- Gallois, Andre (2016). 'Identity Over Time', *The Stanford Encyclopedia of Philosophy*, Winter 2016, ed. Edward N. Zalta (Metaphysics Research Lab, Stanford University), <https://plato.stanford.edu/archives/win2016/entries/identity-time/>.
- Kratzer, Angelika (1977). 'What 'Must' and 'Can' Must and Can Mean', *Linguistics and Philosophy*, 1: 337–55.
- Kripke, Saul A. (1959). 'A Completeness Theorem in Modal Logic', *Journal of Symbolic Logic*: 1–14.
- Kripke, Saul A. (1963). 'Semantical Considerations on Modal Logic', *Acta Philosophica Fennica*, 16: 83–94.
- Kripke, Saul A. (1972). *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980). From Donald Davidson and Gilbert Harman (eds.), *Semantics of Natural Language* (Dordrecht: D. Reidel, 1972): 253–355.

- Leibniz, Gottfried Wilhelm (1710). *Theodicy: Essays on the Goodness of God, the Freedom on Man and the Origin of Evil*, trans. E. M. Huggard (La Salle, IL: Open Court, 1985).
- Leibniz, Gottfried Wilhelm (1956). *Philosophical Papers and Letters*, ed. and trans. Leroy E. Loemker (Chicago, IL: University of Chicago Press).
- Leibniz, Gottfried Wilhelm (1989). *Philosophical Essays*, ed. and trans. Roger Ariew and Daniel Garber (Indianapolis: Hackett).
- Lewis, David K. (1968). 'Counterpart Theory and Quantified Modal Logic', in *Philosophical Papers*, i (Oxford: Oxford University Press, 1983): 26–46. From *Journal of Philosophy*, 65 (1968): 113–26.
- Lewis, David K. (1986). *On the Plurality of Worlds* (Oxford: Blackwell).
- Look, Brandon C. (2013). 'Leibniz's Modal Metaphysics', *The Stanford Encyclopedia of Philosophy*, Spring 2013, ed. Edward N. Zalta (Metaphysics Research Lab, Stanford University), <https://plato.stanford.edu/archives/spr2013/entries/leibniz-modal/>.
- Mondadori, Fabrizio (1975). 'Leibniz and the Doctrine of Inter-World Identity', *Studia Leibnitiana*, 7/1: 21–57.
- Moore, G. E. (1959). 'Four Forms of Scepticism', in *Philosophical Papers* (London: Allen and Unwin): 196–226.
- Quine, Willard van Orman (1953a). 'On What There Is', in *From a Logical Point of View* (2nd edn., Cambridge, MA: Harvard University Press): 1–19. From *Review of Metaphysics*, 2 (1948): 21–38.
- Quine, Willard van Orman (1953b). 'Reference and Modality', in *From a Logical Point of View* (2nd edn., Cambridge, MA: Harvard University Press): 139–59.
- Quine, Willard van Orman (1953c). 'Three Grades of Modal Involvement', *Proceedings of the XIth International Congress of Philosophy*, 14: 65–81.
- Quine, Willard van Orman (1970). *Philosophy of Logic* (2nd edn., Cambridge, MA: Harvard University Press, 1986).
- Sider, Theodore (2011). *Writing the Book of the World* (Oxford: Oxford University Press).
- Sider, Theodore (2020). *The Tools of Metaphysics and the Metaphysics of Science* (Oxford: Oxford University Press).
- Yli-Vakkuri, Juhani (unpublished). 'Rigid Designation'.