


MODAL ARGUMENTS, POSSIBLE EVIDENCE AND
CONTINGENT METAPHYSICS

Michael Thomas Traynor

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2017

Full metadata for this item is available in
St Andrews Research Repository
at:

<http://research-repository.st-andrews.ac.uk/>


Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/15595>

This item is protected by original copyright

Modal Arguments, Possible Evidence and Contingent Metaphysics.

Michael Thomas Traynor



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of PhD
at the
University of St Andrews

Date of submission: 23/09/2016

Abstract.

The present work explores various ways in which contingent evidence can impact metaphysics, while advocating that, just as a scientific realist allows for ampliative inferences to the unobservable, ampliative inferences from possible evidence can warrant possibility claims that lie beyond the reach of sensorial imagination. In slogan form: possible evidence is a guide to possibility. Drawing on Shoemaker's (1969) argument for the possibility of time without change, I advocate the following principle: If there is a possible world at which the observable facts make it objectively reasonable to conclude that *p*, then we should conclude that *p* is possibly true. This provides a route to contingentism in metaphysics, for, if one considers that there are worlds in which the observable facts make it objectively reasonable to conclude that *p*, *and* worlds in which the observable facts make it objectively reasonable to conclude that not-*p*, then my principle tells us that *we* should conclude that possibly-*p* and possibly not-*p*, i.e. that *p* is contingent. This contingency in what is reasonable to conclude, I suggest, occurs most saliently in debates where evidence of phenomenal experience and empirical science are marshalled to support one theory over another. I also explore some consequences of taking possible evidence to be a guide to possibility in this way, among them being an interesting modal analogue of the lottery paradox.

1. Candidate's declarations:

MICHAEL THOMAS TRAYNOR 30,000

I, .../..., hereby certify that this thesis, which is approximately words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in [month, year] and as a candidate for the degree of ~~PhD~~ in [month, year]; the higher study for which this is a record was carried out in the University of St Andrews between [year] and [year].

(If you received assistance in writing from anyone other than your supervisor/s):

I,, received assistance in the writing of this thesis in respect of [language, grammar, spelling or syntax], which was provided by

Date 21/9/2016 signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of ~~PhD~~ in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date 21/9/2016 signature of supervisor

3. Permission for publication: (to be signed by both candidate and supervisor)

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

a) No embargo on print copy

ELECTRONIC COPY

a) No embargo on electronic copy

ABSTRACT AND TITLE EMBARGOES

An embargo on the full text copy of your thesis in the electronic and printed formats will be granted automatically in the first instance. This embargo includes the abstract and title except that the title will be used in the graduation booklet.

If you have selected an embargo option indicate below if you wish to allow the thesis abstract and/or title to be published. If you do not complete the section below the title and abstract will remain embargoed along with the text of the thesis.

- a) I agree to the title and abstract being published
- b) I require an embargo on abstract
- c) I require an embargo on title

YES/NO
YES/NO
YES/NO

21/9/2016

Date signature of candidate

signature of supervisor

Please note initial embargos can be requested for a maximum of five years. An embargo on a thesis submitted to the Faculty of Science or Medicine is rarely granted for more than two years in the first instance, without good justification. The Library will not lift an embargo before confirming with the student and supervisor that they do not intend to request a continuation. In the absence of an agreed response from both student and supervisor, the Head of School will be consulted. Please note that the total period of an embargo, including any continuation, is not expected to exceed ten years. Where part of a thesis is to be embargoed, please specify the part and the reason.

Acknowledgements.

Thanks to everyone at the Arche Philosophical Research Centre, and everyone who contributed to the Metaphysics Research Group, particularly those who came to my talks and provided comments and criticism. Some of the material of Chapter 1 arose out of early conversations with Martin Lipman and was subsequently presented at the Joint Session of the Aristotelian Society and Mind Association and published in the *Proceedings* (114, Issue 3 part 3, 2014). I am grateful to everyone who commented during the talk and those who afterwards provided challenging and constructive criticism. Much of Chapter 2 was presented at the 7th Annual CMM Graduate Conference at Leeds in 2013; I am grateful for the feedback, particularly from Ross Cameron, who responded to the paper. Some of the central ideas of Chapter 2 were also published in *Thought* (Vol. 2, Issue 3, 2013), and I am grateful to Nikk Effingham for providing a response. I have incorporated Cameron's and Effingham's responses into the chapter. Thanks to my parents, Tom and Androulla Traynor. Thanks to the AHRC and the Royal Institute of Philosophy (Jacobsen Trust) for funding my research. Lastly, thanks to my supervisors, Katherine Hawley and Aaron Cotnoir, for innumerable conversations and guidance throughout.

Contents.

0. Introduction. (p.7)
 1. Cameron's negative route to contingentism.
 2. Positive routes to contingentism.
 3. Science as a guide to metaphysics.
 4. Impossible experiments and unobservables.
 5. Chapter summaries.
1. Contingent Scientific and Phenomenal Evidence. (p.31)
 1. Experience as evidence for the passage of time.
 2. Could we experience the passage of time?
 3. Experience of ordinary objects.
 4. Prosser's argument adapted to perception of ordinary objects.
 5. Non-metaphysical explanations, and their contingency
 6. Assuming necessity to refuse experiential evidence.
 7. The conceptual and empirical case against A theory.
 8. Concluding remarks.
2. Actual Time and Possible Change: On Modal Motivations for Mereological Harmony. (p.68)
 1. Grounding Possibilities in Parts: Sider's Desk, Hawley's Homogeneous Object.
 2. Actual Time and Possible Change: The Problem Stated.
 3. Potential Objections.
 4. The contingent explanatory power of temporal parts.
 5. Cameron's modified version of the Sider-Hawley modal argument.
 6. Effingham: An Alternative Harmony.
 7. Structure as a guide to possibility?
 8. Extending the argument to make a case against temporal parts.
 9. A similar issue for reductionism about time.
 10. Concluding remarks.
3. The Vanishing of Almost-Indiscernibility in the Case Against PII. (p.106)
 1. Introduction.
 2. Almost Black's spheres.
 3. Attempting to bridge the 'surely-gap' with causal independence.
 4. Adams's as a Subtraction Argument?
 5. An Argument from Continuity?
 6. Explaining away the intuition: Nature doesn't make jumps, but modal space might.
 7. Distinct indiscernibles as figuring in the best explanation of contingent phenomena: the case from quantum physics.
4. Part I: Imaginative resistance, possible evidence and genuinely new ideas. (p.131)
 1. Genuinely new ideas.
 2. The dependency thesis and the limits of experience.
 3. Imaginative resistance.
 4. Shoemaker's argument from the possibility of partial freezes.
 5. Beyond potential experience, via possible inductions.
 6. Warmbrod's Criticism.
 7. Beyond imagination, via imagined evidence

4. Part II: Possible evidence as a guide to possibility. (p.158)
 8. An implied modal epistemology.
 9. Comparing Hanrahan's Modal Explanationism.
 10. The Possible Evidence Principle, and Potential Objections.
 11. Further examples of the principle at work.
 12. Concluding remarks.
5. Abduction, Agglomeration, A Priority, and Necessity. (p.191)
 1. Introduction.
 2. Modal echoes of the lottery paradox.
 3. A priority, necessity, and agglomeration.
 4. Should we reject agglomeration of a priori justified beliefs?
 5. Contingentism, abduction, and agglomeration.
 6. Abduction and Modality: Choosing a necessity claim from a range of potential explanations.
6. Final Thoughts. (p.223)

0. Introduction.

“Empirical science at most tells us what *is* the case, not what *must* or *may* be (but happens not to be) the case. Metaphysics deals in *possibilities*.” (Lowe, *The Possibility of Metaphysics: Substance, Identity, and Time* (1998, p. 5))

“If we had to travel to other possible worlds to learn what goes on there, we could know nothing about uninstantiated possibilities or necessities; for possible worlds are not foreign countries or planets. Rather, they are alternative complete universes and so are physically inaccessible. Nevertheless, there are standard patterns of inference entitling us to infer facts about ways things could be – about other possible worlds.” (Sorensen, *Thought Experiments* (1992, p. 77))

It is common in philosophy to appeal to imagination for evidence concerning what is possible. Hart (1988) and Yablo (1993) have suggested a compelling analogy to make sense of this: Imagination is to the possible as perception is to the actual. There are limits to the justificatory force of imagination, however, particularly where there is disagreement about what we are able to imagine. I argue that appealing to possible evidence can justify certain possibility claims even if the possibility in question is beyond confirmation by either experience or imagination. In slogan form: possible evidence is a guide to possibility. I will defend the principle that if there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then it is objectively reasonable for us to conclude

that p is possibly true. Let's call this the *Possible Evidence Principle: PEP*. PEP is valuable in cases where *direct* experience could not, in principle, confirm a claim whose possibility is being debated; for instance, the claim that time passes without change (since any observation of the passage of time plausibly requires changes in the observed), that there are distinct indiscernible objects, that the universe has doubled in size, that space is absolute, or that machines think. Just as scientific realism allows that ampliative inferences can lead us to conclusions beyond what can be directly confirmed by observation, so possible ampliative inferences can lead us to possibility claims beyond what can be confirmed by sensory imagination.

The present work thus explores various ways in which contingent, underdetermining evidence can impact metaphysics. One of the central claims is that, where the case for a certain metaphysical thesis depends significantly on contingent matters, this should lead us to conclude that the metaphysical thesis is also a matter of contingency. The observable facts in different possible worlds may make it objectively reasonable to infer different theses; where the different theses in question are rival metaphysical claims, the PEP tells us that the debate in question is a contingent matter. Thus my principle suggests a positive route to contingentism in certain matters of metaphysics, building on Cameron's (2007) argument against the necessitarian presumption.

This introductory chapter introduces some of the main themes underlying the thesis – contingency and necessity in metaphysics, scientific realism, the bearing of science (and experience more generally) on metaphysical questions, unobservables in science and metaphysics, modal knowledge and the limits of imagination – and attempts to give a sense of the ways in which I intend to bring them together. I begin by describing some of the main

points of Cameron's (2007) case for contingentism. I consider Cameron's to be largely a negative argument for contingentism, as it proceeds by arguing against various reasons one might have for necessitarianism. While my intention is to contrast his negative route with what I consider to be a more positive route to contingentism, the negative and positive routes are nonetheless complementary.

1. Cameron's negative route to contingentism.

Van Inwagen's special composition question (SCQ) asks:

Under what conditions do the xs compose an object? (van Inwagen 1990, 30-31).

Several answers have been proposed, each falling under one of three types: always, sometimes, or never. As Cameron (2007) notes,

even though there is widespread disagreement as to what the answer to that question is, there is widespread agreement that whatever the answer is, it is a necessary truth. (Cameron 2007, 100)

However, Cameron argues against this consensus. Cameron's argument is largely negative; surveying several reasons one might give for taking the necessitarian stance, he argues that each is unsatisfactory.

1.1. Composition facts as non-trivial.

Cameron begins by arguing that composition facts are non-trivial, not a matter of definition. Answers to the SCQ give necessary and sufficient conditions for composition to occur, so can be stated in the form:

There is a y such that the x s compose y iff the x s satisfy conditions C .

Answers to the SCQ are not analytic, since the biconditionals expressing such answers assert the existence of an object on one side – the left – only. Assuming, plausibly, that existence claims are never analytic, the only way for a conditional expressing the necessary conditions for something's existence to be analytic would be for the antecedent to contain an existence claim, too (Cameron 2007, 102). And this is just as well, for conditionals expressing answers to the SCQ are meant to be informative.

But perhaps the necessitarian view is not that the answer to the SCQ is necessary, *whatever it happens to be*. Perhaps instead it is that some particular answer to the SCQ is favoured because it follows analytically from another claim, which is supposed to be necessary. An example along these lines is the view that, necessarily, composition is identity, and that, from composition as identity, mereological universalism follows trivially. One might think this because, according to composition as identity, the sum of the x s *just is* the x s; the x s, whatever they are, are identical to their sum. Hence it follows trivially that the sum of the x s exists wherever the x s exist.

However, the composition as identity theorist does not merely hold that every collection of objects is identical to *its sum*. Rather, the view is that, wherever some things compose *an object*, that object is identical to the things that it is composed by. But, Cameron points out, this does not mean that wherever there are some things, there is an object that is composed of those things. Hence composition as identity does not entail universalism about composite objects; it provides us with an identity statement about any composite objects that there happen to be, but does not tell us that there *are* any such composites. Therefore, composition

as identity is compatible with the falsity of mereological universalism (Cameron 2007, 103-104)¹.

1.2. Necessity not entailed by a priority.

Next, Cameron tackles the claim that composition facts are necessary because they are a priori. He does this, not by denying that composition facts are a priori (he grants that they are a priori for the sake of argument)², but rather by attempting to block the general legitimacy of the inference of necessity from a priority.

There is something compelling about the idea that a priori claims are necessarily true, if true at all; but this idea fell into disrepute in the latter half of the previous century, in no small part due to Kripke's seminal work in *Naming and Necessity* (1981). It is compelling, it seems, for epistemological reasons. To believe a proposition on a priori grounds is to believe it on the basis of justification that does not come from experience. Assuming that a priori justification is empirically infeasible (controversial – we'll return to this in Chapter 5), this seems to imply that whether one is a priori justified in believing something has nothing to do with how the world is. So if one is a priori justified in believing something of the actual world, one should equally believe it of all worlds, for there is no way the world could be that could diminish that justification.

¹ This has more recently been argued in McDaniel (2010), Cameron (2012). Bohn (2014) defends the entailment against McDaniel (2010) and Cameron (2012); Calosi (2016) argues that a version of composition as identity actually entails mereological nihilism (the view that there are no conditions under which composition occurs). I will not attempt to weigh in on this extended debate – my aim here is primarily to illustrate Cameron's general strategy and give a sense of the plausibility of his position.

² Though, in passing, Cameron does say that he is "not convinced that the facts about when a collection of objects composes some further object are a priori", and in a footnote says: "It is an empirical matter, for example, whether a collection of objects is fastened together, whether they contrast to their surroundings etc. And even if it is no part of what it is for a collection of objects to compose that the members of the collection be fastened together etc., such facts might still constitute *evidence* for the claim that they compose" (Cameron 2007, 106-7; footnote 21). That such contingent facts may help to provide positive justification for an answer to the SCQ is very much in the spirit of the present work.

That is to draw a route from epistemological claims to modal metaphysical claims. Kripke, however, sought to pull these apart.

We ask whether something might have been true, or might have been false . . . If it is true, might it have been otherwise? Is it possible that, in this respect, the world should have been different from the way it is? If the answer is ‘no’, then this fact about the world is a necessary one. If the answer is ‘yes’, then this fact about the world is a contingent one. This in and of itself has nothing to do with anyone’s knowledge of anything. . . [The concepts of a priority and necessity] are dealing with two different domains, two different areas, the epistemological and the metaphysical. (Kripke 1981, 36)

In a sense, this gives expression to a realist stance: modal properties are mind independent, and so are not *tied* to what we can be justified in believing, let alone what we do know.

(However, I’ll later argue that what we *could* be justified in believing can nonetheless give us reason for believing certain modal claims; this is compatible with the realist stance expressed here.)

In this spirit one might look to putative examples of contingent a priori truths to provide counterexamples to the claim that a priori claims are necessarily true, if true at all, thereby rendering the inference of necessity from a priority invalid. Cameron however argues that such purported counterexamples “seem to turn on a linguistic trick”. Kripke cites the example of the claim “Jack the Ripper committed the East End Murders (if anyone did)”, where the referent of ‘Jack the Ripper’ is fixed by the definite description, ‘the person, whoever it is, who committed all or most of the East End murders (if anyone did)’ (Kripke 1980, 71, 94). Cameron objects that all that is a priori here is that such sentences express truths. That “Jack the Ripper committed the East End Murders (if anyone did)” expresses a

truth can be known a priori – but what is known here is also necessary, for it is true in every world that the person who committed all or most of the East End murders (if anyone did) committed all or most of the East End murders (if anyone did). We do not, however, know of some particular person that they committed the East End murders – such knowledge would be of a contingent fact, but in order to know anything like that, we would have to go out into the world and investigate (Cameron 2007, 107).³

Cameron provides a different reason for rejecting the inference of necessity from a priority, but one that is no less in the spirit of Kripke's separation of the metaphysical from the epistemological. A proposition is a priori when it can be known on the basis of justification that is not empirical. But justification, of course, does not imply truth:

So the argument from a priority to necessity fails because it assumes that the way the world is plays no part in my coming to know that p. What is true is just that it plays no part in my coming to have a *justified belief* that p; but it *does* play a part in my coming to know that p, because my justified belief only amounts to knowledge if p is in fact true" (Cameron 2007, 108).

A priori justification, though it may be *actually* reliable, cannot be *assumed* to be *necessarily* reliable without begging the question. Cameron draws an analogy with appeal to theoretical virtues, such as a principle according to which simpler theories are more likely to be true. Many of us believe that simplicity is a guide to truth. We may be justified in believing one theory over another, therefore, if it trumps that theory in terms of its simplicity. But

³ The makings of this objection are found in *Naming and Necessity*: p. 63, footnote 26: "If someone fixes a meter as 'the length of stick *S* at *t*₀, then in some sense he knows a priori that the length of stick *S* at *t*₀ is one meter, even though he uses this statement to express a contingent truth. But, merely by fixing a system of measurement, has he thereby *learned* some (contingent) *information* about the world, some new *fact* that he did not know before? It seems plausible that in some sense he did not, even though it is undeniably a contingent fact that *S* is one meter long. So there may be a case for reformulating the thesis that everything a priori is necessary so as to save it from this type of counterexample" (Kripke 1980, 63, footnote 26).

presumably simplicity is only a contingently reliable guide – there may be worlds, or even possible situations, in which simplicity leads us astray. We may be justified in believing a theory on the basis of simplicity, but our belief can only amount to knowledge in those worlds in which the theory is true (Cameron 2007, 108-109).

We can draw a general principle from Cameron’s discussion of simplicity: where the route to *p* is only contingently reliable, we should stop short of concluding that *p* is metaphysically necessary. But with regard to the present question – whether necessity can be inferred from a priority – this last analogy can only take us so far. Sure, *if* a priori justification, like justification based on a principle of parsimony, is only contingently reliable, then we cannot infer necessity from a priority. But the antecedent of that last conditional is not to be assumed here. The question is, what shall we say about beliefs based on justification that (we believe) would be strong enough to warrant belief in *any* possible world? It is true that the world needs to cooperate in order for a justified belief to amount to knowledge. But what if one believes that there is no way the world could be that would diminish one’s warrant for believing *p*, and that this is what provides grounds for believing that *p* is necessarily true? If this is what the necessitarian has in mind, then Cameron’s response misses the mark. The necessitarian can agree that the world needs to cooperate in order for their belief to be true, and that *all* worlds need to cooperate in order for that belief to be necessarily true. They can also agree that *generally speaking* justification does not entail truth. Indeed, these are trivial points. But it is also trivial that justification can provide reason for belief; hence, if one has a reason to believe something such that one would be equally warranted in that belief in *any* possible world, one might think that one’s justification amounts to justification in a necessity claim. I will return to this issue in Chapter 5.

There is more to Cameron's argument, but this much perhaps gives us enough of a sense of the negative case in favour of contingentism, being the case against the presumption in favour of necessitarianism. We now turn to some positive reasons one might have for holding that certain metaphysical debates are contingent matters.

2. Positive routes to contingentism.

2.1. Combination of conceptual and empirical claims.

Rosen's (2006) proposal is that something is possible just in case it is *correctly conceivable*. For a proposition to be correctly conceivable is just for it harbour no absurdity "given a full specification of the natures of the items it concerns" (Rosen 2006, 24). (Thus, on this view, metaphysical possibility is the default status for propositions: a proposition is presumed to be metaphysically possible unless some reason is given – specifically, in the form of an absurdity – why it should not be (Rosen 2006, 23).) This preserves the link between conceivability and possibility in the face of a posteriori necessities, the contraries of which are in *some sense* conceivable, but not *correctly* conceivable. The usual example is the case of water. It is a common view that water is essentially, and therefore necessarily, H₂O, and that this is something that has been discovered empirically. Prior to this discovery we would surely have found it conceivable, in a sense, that water had a different composition. However, now that we know its essential nature, this should no longer be conceivable – that is, it is not *correctly conceivable* that water is something other than H₂O, for we find that such a claim harbours an absurdity once we see that water *is* H₂O.

This suggests a simple route to contingentism. Correct conceivability is a guide to possibility. If we believe we can correctly conceive a world in which p is true, and correctly conceive a world in which p is not true, we have grounds for concluding that p is contingent. One

problem with this is that it is often just as contentious whether something is conceivable as whether it is possible; in the case of correct conceivability in particular, we might wonder: how do we know when a full specification of the relevant natures are such as to render p contingent?

Miller (2010) suggests that the case of water provides a clue. We take it that empirical investigation has revealed that water is H_2O necessarily, in part because we take it that 'water' rigidly designates the watery (clear, drinkable, odourless, etc.) stuff of the actual world. But what if 'water' functioned differently, to pick out a functional kind, referring to whatever plays the water role. This would be to understand $\langle \text{water} \rangle$ as a functional concept, such that different things constitute water in different worlds, since different things could satisfy such a role; in which case, correct conceivability would be a route to contingency, for it would be correctly conceivable that water is H_2O *and* correctly conceivable that water is XYZ ($\neq H_2O$). This is just how Craig Bourne treats time, arguing that it is contingently A-theoretical (Miller 2010, 968-969; Bourne 2006, 220-230). In this way, a combination of conceptual and empirical claims can lead one to contingent metaphysical claims.

2.2. *Synthetic metaphysics.*

However, Miller notes a problem for the general validity of this route to contingency. The problem is that it can come into conflict with what she has called "synthetic" metaphysical necessities. To hold that there are such necessities is to "reject the idea that a priori reason can only yield necessary truths that are either conceptual [(analytic)] necessities or a posteriori necessities" (Miller 2010, 969). This means that one cannot infer that p is contingent from the fact that p is correctly conceivable and $\neg p$ is correctly conceivable. At

least, not from that alone – one should also be confident that neither necessarily- p nor necessarily- $\neg p$ is a synthetic necessary truth (Miller 2010, 969).

How does one arrive at a synthetic metaphysical thesis?

[B]y determining which of competing theses best meet a complex set of desiderata that tend to include the extent to which a theory preserves or explains our folk intuitions (perceptions and phenomenologies) and folk semantics, the extent to which it is simple, parsimonious, explanatory and consistent with our best science. Reasoning regarding which theory is preferable is a priori only against a background of accepted facts about folk semantics, and empirical facts. (Miller 2010, 970)

This is, in effect, to employ inference to the best explanation, a mode of inference whereby we shortlist the best theories and determine, by way of desiderata such as those mentioned by Miller, which provides the best account of the relevant phenomena. We then infer to whichever theory does so. This in turn suggests another route to contingentism, for different theories might best satisfy the desiderata at different worlds.

2.3. Contingency of evidence for a metaphysical thesis.

In order for inference to the best explanation to direct us to the truth of different theories at different worlds, there must be some observable difference among these worlds, some contingent *information* relevant to theory choice. Good candidates for such differences are phenomenal experience and empirical science. We'll discuss several cases in the course of this work but for now let's take Miller's example for illustration: the issue of whether and how general relativity provides evidence for substantivalism concerning space-time. This is a curious case because the metaphysical upshot is far from clear: much philosophical work is needed to determine how the science bears on the metaphysics, for arguments citing general

relativity have been advanced both for and against substantivalism. On the substantialist's side, for example, Nerlich (1994) argues that general relativity is best understood in substantival space, for the substantialist has a simple explanation for why objects should behave as if they occupy space of variable curvature such as is posited by general relativity (Miller 2010, 970; Dainton 2010, Chapters 14, 20 and 21). Against the substantialist are proponents of the 'hole' argument (Earman and Norton 1987), which begins from the premise that, once we accept substantival space, we should allow for the possibility of a permutation of any two points in the spacetime manifold, such that what happens where would be altered, even though the resulting world and the actual world would be observationally equivalent. General relativity allows for this, since it "is a *generally covariant* theory, which (roughly) means that the mathematical form of the laws of physics are the same no matter which system of coordinates are used to map spacetime points" (Dainton 2010, 372). In a nutshell, the problem is that the possibility of such a permutation of points would amount to a failure of determinism, for unduly metaphysical reasons. As Earman has put it,

The argument ... does not rest on the assumption that determinism is true, much less on the assumption that it is true a priori, but only on the assumption that it be given a fighting chance. To put the matter slightly differently, the demand is that if determinism fails, it should fail for a reason of physics.
(Earman 1989, 180; quoted in Dainton 2010, 374)

By rejecting substantivalism, we can adopt Leibnizian equivalence, according to which observationally equivalent worlds are identical. The permutation of points would then be thought of as analogous to Leibnizian shifts (where every object in the universe moves, say, fifty metres in one direction, preserving all spatial relations between objects), the possibility of which Leibniz denied.

While it is true that each side in this debate invokes general relativity to reach a conclusion contrary to the other, this does not in itself make for a standoff – there may nonetheless be philosophical grounds for favouring one theory over the other, taking general relativity as a premise. I will not attempt to determine which, but suppose that one side does have a stronger case. Insofar as that case was contingent on the laws of nature, and insofar as the laws of nature are themselves contingent, it is a short step to contingentism in that debate:

If the laws of nature are contingent, then it is an open question whether in worlds with different laws, a different theory of the nature of space-time will be preferable. It does not seem at all implausible that this could be so and therefore that we could have reason to think that substantivalism and relationism are contingent metaphysical theses. (Miller 2010, 970)

There is an implicit premise here, and it represents an intriguing route to modal claims.

The general idea is that in each world, we should suppose that metaphysical thesis to be true which is the best theory of that world. This opens up the possibility that some synthetic metaphysical claims will be necessarily true: namely where such a thesis is the best theory of every world, and also opens up the possibility that some metaphysical claims will be contingent truths: namely where one thesis is the best theory of some worlds, and a competitor thesis the best theory of a different set of worlds. More generally, a synthetic thesis is necessary just in case it is preferable to its competitor theses in all worlds, and contingent where it is preferable to its competitors in only some worlds. (Miller 2010, 970)

This idea – that we should endorse a possibility claim if it forms an essential part of the best theory at some possible world – is central to the present work. In some cases, this will only provide a route to contingentism if it is allowed that the laws of nature are contingent. By and large, I will simply be assuming that the laws of nature are contingent, but I will return to this assumption in Chapter 4.

Why is this route to modal claims of interest? When combined with certain other ideas, it has the potential to bring the epistemology of the possible into line with that of the actual. It is a familiar idea – and one that I will largely take for granted in the present work – that imagination is to some extent a guide to possibility. This is sometimes understood by way of an analogy with perception: imagination is to the possible as perception is to the actual (Hart 1988, 10; Yablo 1993, 7; Hanrahan 2007, 2009). But we do not limit ourselves to perception when investigating the actual world: we can also construct theories, and believe some of these theories when they do a good job of accounting for what we perceive. Extending the perception-imagination analogy, then, we can also construct theories about imagined phenomena. Might theorising about *possibilities* provide justification for certain possibility claims, where those claims do a good job of accounting for possible phenomena?

One problem with imagination as a guide to possibility is that people disagree about what is imaginable. The idea that imagination is to the possible as perception is to the actual fits quite well with a certain characterisation of sensory imagination, namely that when we sensorily imagine something, we are imagining “from the inside” an experience of that thing (Gallois 1974; Peacocke 1985; Martin 2002). But tying imagination to possible experience in this way pushes anything beyond experience beyond imagination, too. Those who characterise imagination in this way may well disagree with certain claims to have imagined what lies beyond possible experience. This leads to a further interesting feature of the proposed route to modal claims: if valid, it allows us to infer to certain possibilities even if they are sensorily unimaginable. The idea is that, just as a scientific realist allows for ampliative inferences to the unobservable, ampliative inferences from possible evidence can warrant possibility claims that are – perhaps inherently – beyond the reach of sensorial

imagination. In slogan form: possible evidence is a guide to possibility. As the slogan suggests, this is not meant as a way of generating modal knowledge from scratch: in order to proceed from possible evidence to a possibility claim, we need to be in a position to affirm the possibility of that evidence. Thus I assume that there are other ways of justifying modal claims; but thinking about what would be objectively reasonable to conclude given merely possible evidence can, I contend, justify possibility claims where sensory imagination, conceivability, and (of course) mere stipulation fall short.

What about necessity? Here, I will argue that there is an asymmetry. Possible evidence can be a guide to possibility, and a route to contingentism when the relevant evidence varies across modal space. However, the route to necessities that Miller suggested is problematic. In Chapter 5 I show how a modal analogue of the lottery paradox arises when we allow for possible evidence to tell us what is true at a possible world. A prominent solution to the lottery paradox is to deny that warranted beliefs about *each* ticket (that it will lose) should be agglomerated into a belief about *all* tickets. Analogously, a certain belief can be warranted about *each* among a given set of possible worlds, perhaps through inference to the best explanation at each relevantly similar world, perhaps because the warrant is a priori. But can we actually agglomerate these beliefs about each world into a belief about *all* such worlds? If so, we seem to have a new route to the necessity of a priori truths; if not, the case for contingentism in metaphysics is strengthened in the face of the claim that metaphysics is a realm of armchair reasoning. In this way, surprising connections are drawn between modal metaphysics and solutions to paradoxes of rationality.

3. Science as a guide to metaphysics.

In order for this route to modal claims to lead us to contingentism, we need to feed some contingent facts into it, so that what is most reasonable to conclude differs in different worlds. I have suggested that differences in phenomenal experience and empirical science provide us with such differences. But can these things be guides to metaphysics?

Hawley (2006) is instructive on this question. Hawley characterises the following three positions.

(Radical Pessimism) The involvement of a metaphysical claim in an empirically successful scientific theory can never provide any reason to think that the claim is true.

(Moderate Pessimism) There is a kind of involvement in theory which, were a metaphysical claim to achieve this involvement, would provide some reason to think the claim is true; but there are no cases of metaphysical claims being involved in theory in this way.

(Optimism) There are actual cases in which the involvement of a metaphysical claim in an empirically successful scientific theory provides some reason to think that the claim is true.

(Hawley 2006, 456)

Here “the involvement in question is the kind of involvement which, according to scientific realists at least, gives us reason to believe a claim about unobservable entities” (Hawley 2006, 456), and scientific realism is understood as “the view that there are cases where the involvement of a claim about an unobservable entity in an empirically successful scientific theory provides reason to think that the claim is true” (458). In order to avoid begging the question at issue, it is necessary to avoid characterising metaphysics in such a way that its involvement in science, or science’s involvement in it, is either demanded or precluded. Hence metaphysics is characterised in terms of the kinds of questions it deals with, rather than the ways in which it answers them (so rather than, say, claiming that metaphysics is the

province of the a priori): “questions about the nature of time, causation, properties, numbers, persistence, possible worlds and so on will all count as metaphysical questions” (Hawley 2006, 452).

Hawley goes on to argue that scientific realists should reject Radical Pessimism. The latter would commit them to some principled difference between science and metaphysics, such that involvement of some ‘scientific’ entity in a scientific theory can provide justification for believing that entity exists (even if unobservable), whereas involvement of some ‘metaphysical’ claim could not provide such justification. However, it is hard to see how such a distinction could be motivated (Hawley 2006, 459-460).

For my project to get off the ground, I also need to reject Radical Pessimism. I will be advocating the view that possible and contingent evidence can be evidence for possibility claims, where these possibility claims can concern unobservable entities. As such, I take it that scientific realism is correct – indeed, by claiming that possible evidence is a route to the possibility of certain kinds of unobservables, I am effectively advocating scientific realism about what goes on at other possible worlds. This is not to endorse Lewis’s (1986) view that possible worlds are concrete; rather, the idea is that, when we consider that there is a possible world in which the observable facts make it objectively reasonable to conclude that p, we should *actually* conclude that p is possibly true (this will be spelled out more fully in Chapter 4). Since it is beyond the scope of the present work to argue in favour of scientific realism, I hereby designate it as an assumption.

This leaves us with Moderate Pessimism and Optimism. In the course of this work, I will attempt to provide some examples where the involvement of metaphysics in empirical theory

provides some reason for endorsing the metaphysics (we have already seen one potential case of this: whether substantivalism or relationism are is favoured by general relativity). In addition, in Chapter 2 I will attempt to provide a case in which contingent facts concerning the structure of space and time – arrived at through empirical investigation – can constrain metaphysics, by placing constraints on the kinds of metaphysical explanations that are available to the metaphysician. If that's right, then the idea that possible evidence is a guide to possibility is the other side of the coin, for when the viability of a certain metaphysical theory hinges on certain contingent facts, the obtaining of certain contrary contingent facts would provide a boost to certain contrary metaphysical theories.

4. Impossible experiments and unobservables.

What I find most intriguing about using possible evidence to justify possibility claims is the potential it has to expand metaphysics beyond the limits of other, perhaps more familiar (and I dare say, mundane) routes to modal claims. There are limits to actual experience and actual science. Scientific realists are of the view that, to a certain extent, our theorising need not be bound by these limits, for they believe that actual observations can provide justification for belief in certain unobservables. But for the metaphysician who wants to explore the modal realm, actual evidence is of limited interest, for the possibility claims so arrived at will only emerge trivially, because actuality implies possibility. Once the empirical work is done, the metaphysician need only notice it, and the work is done as far as the relevant possibility claim is concerned. (However, this is not to say that such possibility claims must be uninteresting for the metaphysician – as I argue in the Chapter 1, section 6). There are also limits to the capacity of imagination to justify possibility claims (not least because disagreements about what is imaginable are not uncommon), and appeal to possible evidence,

I claim, can provide us with the means to justify possibility claims whether the possibilities in question are imaginable or not.

In this last respect, consideration of possible evidence does for imagination what certain thought experiments can be said to do for experience, which is to provide justification for a claim that lies beyond it. Galileo's famous refutation of Aristotle's theory of falling bodies is arguably such a thought experiment. Aristotle's idea was the apparently common-sense one that heavier bodies fall faster than lighter ones, and that the difference in rate of descent is proportional to their weights. To counter this, Galileo has us imagine such a scenario.

According to Aristotle, if dropped from the same height, a cannonball and a musket ball will hit the ground at different times – the former before the latter, because the former is heavier. But suppose that they are attached to one another. According to Aristotle's theory, the two bodies together should fall faster than either one would separately, for the weight of both taken together is greater than the weight of either one taken separately. However, also according to Aristotle's theory, the two taken together should fall more slowly than the cannonball alone would fall, for the musket ball, being lighter than the cannonball, should drag on the latter, slowing it down. This amounts to a *reductio* of Aristotle's theory, showing it to be not only false, but necessarily false (Brown 2004, 24-25). This leaves us with only one alternative. By refuting Aristotle, the thought experiment establishes that objects in a vacuum fall at the same rate, regardless of size, and it does so without having to look into the world (Brown 2004, 30).

Newton's two-sphere thought experiment is another case in point. Here it is, as described by Brown:

We image the universe to be completely empty except for two spheres connected by a cord. The spheres are of a material such that they neither attract nor repel one another. There is a tension in the cord joining them, but the spheres are not moving toward one another under the force in the cord. Why not? Newton offers an explanation: they are rotating with respect to space itself; their inertial motion keeps them apart. And so, Newton concludes, absolute space must exist. (Brown 2004, 28-29)

One reason this experiment cannot be carried out in the actual world is that we are unlikely to find two such spheres, “made of a material such that they neither attract nor repel one another” (Brown 2004, 28). Another, perhaps more interesting reason, is that it is essential to the case that no object other than the two spheres exists. This is crucial for Newton’s purposes because his intended conclusion is that absolute space exists, and he reaches this conclusion by way of the idea that the spheres must be rotating with respect to something. Since there is nothing in the universe besides the spheres and the cord connecting them, that something has to be absolute space.

Newton seems to move from the claim that absolute space must exist in order for us to make sense of what happens in the two-sphere world, to that claim it must exist in the actual world. However, this seems unwarranted. Since the inference to the existence of absolute space at the two-sphere world depends on features of that world which do not pertain to the actual world, it is hard to see how consideration of the two-sphere world provides us with any grounds for thinking that absolute space *actually* exists.

However, what does seem reasonable is that, if indeed Newton has shown that absolute space is part of the best explanation of the two-sphere world, then we should grant that it is possible for space to be absolute. As I’ll be arguing, this is an instance of a kind of argument that can be particularly useful in metaphysics, in particular in contexts where the dispute is about

some possibility claim, and where an appeal to imagination, for one reason or another, is ineffective. In Newton's case, this may well be why thought experiments like the two-sphere one are appealed to: it is not at all clear that absolute space can be imagined, so appeal to imagination is ineffective against someone who doubts its possibility. This is also in line with Brown's characterisation of the thought experiment.

Newton's spheres do not so much give us a new result but, rather, give us a remarkable phenomenon, something that needs to be explained. The thought experiment establishes a phenomenon; the explanation comes later. And the best explanation, according to Newton, is the existence of absolute space. This way of looking at it is confirmed, it seems to me, by the way in which Berkeley and Mach reacted to the thought experiment. They didn't deny that rotation with respect to absolute space is the best explanation for the tension in the chord. Instead, they denied there would be any tension in the first place, if the spheres are not moving together. Or the two spheres would move together, if there were any tension. That is, they didn't bother to challenge the explanation of the phenomena that Newton posited; they challenged the alleged phenomenon itself. (Brown 2004, 29-30)

To these remarks about Berkeley and Mach, we should also add that, had their strategy been to challenge the explanation of the phenomena that Newton posited – say, by providing an alternative explanation – this would amount to tacit acceptance of the general method of argumentation. In Chapter 4 I give further examples of arguments that draw on the best explanation of possible phenomena, and defend a principle, (PEP, mentioned above) behind such appeals: That, if there is a possible world in which the observable facts make it objectively reasonable to conclude that p , then it is objectively reasonable for us to conclude that p is possibly true.

5. Chapter summaries.

To recap, the thesis will be an investigation into various ways in which contingent and underdetermining evidence might count in metaphysics, with particular interest in debates about composition and persistence, the nature and structure of space and time, and arguments from possibility – and the broader question of how this might bear on the modal status of metaphysical claims

Chapter 1) This chapter is about whether phenomenal experience and empirical science – indeed, contingent evidence more generally – can count as evidence in metaphysics. If it can, we have a potential route to contingentism in the relevant metaphysical debates. If the character of experience can have a bearing on which theory (e.g. of composition) we should take to be the correct one, then we might have contingent evidence for a metaphysical theory. In worlds in which that evidence is lacking, the case for that theory might not be strong enough to warrant belief in it at those worlds, and this constitutes a *prima facie* reason for its being considered at best contingently correct. I make the same suggestion about evidence from empirical science: e.g. special relativity has been invoked against presentism, but most philosophers accept that special relativity is at best contingently correct. Thinking about worlds in which special relativity is false encourages us to re-evaluate the metaphysics at such worlds.

Chapter 2) This chapter looks at particular modal arguments that aim to settle the spatial and temporal structure of objects. These are basically arguments from explanation: in order to account for the mere possibility of (change and) inhomogeneity, things must actually have (temporal) parts that are as fine-grained as possible (change and) inhomogeneity. I identify a potential problem for this idea: if we take it that the structure of objects can be settled in this

way, but leave open the possible structure of space and time, then we might end up having to conclude that objects are more fine-grained than the regions at which they are located. The solution I advocate is that we harmonise the investigations: if it is a contingent matter what the structure of space and time are, then the metaphysics of persistence is also a contingent matter.

Chapter 3) In this chapter, I consider three potential construals of Adams's (1979) argument for the possibility of distinct indiscernibles, against the identity of indiscernibles (PII). I'll argue that the first two bridge the gap in Adams's argument at the cost of making almost-indiscernibility – the most salient feature of Adams's case – irrelevant to the argument; as such, Adams's case is rendered ineffective for the purpose of convincing anyone apt to doubt the possibility of Black's case. The third construal takes Adams's argument to be one from modal continuity; however, while this would make almost-indiscernibility relevant to the argument, I'll suggest that it too fails to get around Hacking's (1975) critique of supposed counterexamples to PII such as Black's – a critique which, after all, Adams's argument was originally intended to overcome. I close the chapter by briefly highlighting an alternative way in which some have sought to refute PII: claiming that it conflicts with a certain interpretation of certain quantum physical phenomena.

Chapter 4) This chapter discusses whether possible evidence for p can be interpreted as being actual evidence for possibly p – taking Shoemaker's "Time Without Change" as a case study. Shoemaker considers an imaginary world in which the inhabitants have inductive evidence that suggests that time passes in the absence of change. Understood as a positive argument for the possibility of time without change, the idea is that, since the observable facts make it objectively reasonable for the inhabitants of Shoemaker's world to believe that time passes in

the absence of change, we should conclude that it is possible for time to pass without change. This assumes that there is a route from what is possibly probable to what is (probably) possible, in the sense that the former is a guide to the latter. Abstracting from the particulars of Shoemaker's argument, we arrive at a principle that lies behind it, what I have called the Possible Evidence Principle (PEP): If there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then it is objectively reasonable for *us* to believe that *p* is possibly true. If this principle can be successfully defended, it could represent a new route to possibility claims, and can be invoked to make progress in debates over possibility claims where appeals to imagination are ineffective.

5) This chapter looks at how the principle PEP, discussed in the preceding chapter, leads to what Faraci (2013) has called a "modal echo of the lottery paradox". In keeping with this principle, we should say that, for each world in which it would be reasonable to conclude that *p*, we should say that *p* obtains in that world. But if that's true, then it is natural to think that *p* obtains in *every* such world. Applying this reasoning to a priori justification, we can end up concluding that all a priori claims should be considered necessity claims (subject to a certain understanding of a priori justification, such that if one is a priori justified in believing *p* at any world, one is equally justified in believing *p* at all worlds). I'll suggest, however, that we should follow a prominent solution to the lottery paradox, and reject the agglomeration step that this argument from a priority to necessity would require. I'll end the chapter with a discussion of Biggs's (2011) case for the claim that modality is centrally abductive, in which case both possibility and necessity claims could be justified by abduction.

1. Contingent Scientific and Phenomenal Evidence.

This chapter explores the potential for (contingent) scientific and phenomenal evidence to bear on particular metaphysical debates. One area of metaphysics in which both empirical science and phenomenal experience have been brought to bear is the debate concerning whether time is dynamic or static (i.e. the debate concerning whether time really passes). After describing an argument from Prosser (2007) for the conclusion that we cannot experience the passage of time, I investigate whether the argument can be adapted to other areas of metaphysics to reach a similar conclusion: namely, that we should not take the character of experience to count as evidence in the debates concerning how things persist and whether there are composite objects. Along the way I suggest some ways one might go about maintaining that the character of experience has a role to play in these debates in spite of the arguments against doing so. I also highlight that, where results of empirical science help to make the case for a certain metaphysical thesis, and where such results could have been otherwise, the force of the case for that thesis will vary depending on what possible world you are in. This helps to see how taking possible evidence as a guide to possibility can lead to contingentism in a range of metaphysical debates.

1. Experience as evidence for the passage of time.

Following McTaggart's famous (1908) discussion, theories of the nature of time are typically categorised as either dynamic or static theories, "A theories" and "B theories", respectively. A theories can be roughly characterised as positing that time flows; objects and events go from being future to present to past, and can be considered as having corresponding, mutually exclusive properties in their turn, first of futurity, then of presentness, then of pastness. B theories deny time this dynamism: events are ordered in time, standing in earlier-than/later-

than relations to other events, but there is no *passage* of time, if ‘passage’ is understood in the A theorist’s terms. The position of events in the time series never changes: if an event is earlier/later than some other event, it is *always* earlier/later than it, and the standing of events in these earlier/later relations is all there is to the time series – there is no “objective becoming”.

A key point of contention between theorists of each category is the role of our experience. It is hard to escape the sense that time really does pass, and there is something altogether more intuitive, more natural about the A theorist’s picture. What once was future is now present, then past, and, pre-theoretically at least, we might go as far as to say that we *perceive* (perhaps through any of the five senses, or perhaps through the “mind’s eye”) the passage of time – the passing of events from the future to the present to the past. Indeed, it is a common suggestion that the character of our experience supports A theory over B theory.

Does our impression of the flow of time, or the division of time into past, present and future, tell us nothing at all about how time *is* as opposed to how it merely appears to us muddle-headed humans? ...As a human being, I find it impossible to relinquish the sensation of a flowing time and a moving present moment. It is something so basic to my experience of the world that I am repelled by the claim that it is only an illusion or misperception. It seems to me that there is an aspect of time of great significance that we have so far overlooked in our description of the physical universe. (Davies 1995, 275)

Let me begin this inquiry with the simple but fundamental fact that the flow of time, or passage, as it is known, is given in experience, that it is as indubitable an aspect of our perception of the world as the sights and sounds that come in upon us, even though it is not the peculiar property of a special sense. (Schuster 1986, 695)

[Experience is] a defeater-defeater that overwhelms any B-theoretic arguments against the reality of tense. (Craig 2000, 138)

I cannot survey all the motivations philosophers have had for the moving spotlight theory. But the motivation that I like best appeals to the nature of our conscious experience. (Skow 2009, 677)

Let's consider two lines of response open to the B theorist. First, the B theorist might question whether it really does seem to us that time passes along the lines A theorists propose. Compare the suggestion that, pre-theoretically at least, it appears to us that that the Sun moves around the Earth, rises and sets – rather than that the Earth rotates on its axis. Wittgenstein has been quoted as responding to this suggestion: “Well ... what would it have looked like if it had *looked* as if the earth turned on its axis?” (Quoted in Anscombe 1959, 151).

Wittgenstein's response is powerful, but what does it amount to when adapted to the B theorist's cause? The rhetorical force of Wittgenstein's remark is in no small part due to the fact that heliocentrism has been so roundly discredited. But in this respect the cases are disanalogous. Any context in which experience is invoked in favour of A theory is of course one in which A theory is still a contender. Thus we cannot simply say, “*This* is of course what it would look like if it *looked* like B theory were true (because B theory *is* true, and this is how it *does* look)”. Heliocentrism has been roundly rejected, but A theory has not.

Perhaps the Wittgenstein-style remark holds some other significance, but the alternatives that come to mind are of limited help to the B theorist. Here are two potential implications of that remark.

First, the suggestion might be that there is nothing in particular that it would be like to experience the world as if B theory were correct. But this is simply to suggest that B theory is necessarily without support from experience – there is nothing our experience could be like that would suggest that B theory is true.

Second, the suggestion might be that, just as it turned out that our experience is consistent with the Earth rotating on its axis, so our experience is consistent with the truth of B theory. Here, however, a caution is in order. As Bradford Skow has pointed out, “Even if it is true that each of the theories is *consistent* with my experience being as it is, it does not follow that the fact that my experience is that way fails to *favor* one of the theories over the other. This is an instance of a general epistemological truth: a body of evidence can support some hypothesis without entailing that hypothesis” (2011, 361). This point is crucial. Somehow it is easy to suppose that it is enough for the B theorist to show that B theory is consistent with the appearance that time passes (the B theorist might claim, for example, that this is merely an illusory appearance). But that would be to defend against a straw man; experience might be evidence *against* B theory, in *favour* of A theory, while nonetheless being *consistent* with both.

Skow’s general point should be familiar from the history of philosophical scepticism: my experience is consistent with the possibility that I am a brain in a vat, stimulated to have just those experiences that I in fact have and take (in general) to be veridical; but many of us still prefer the theory that we are not in such scenarios, and take it that our experiences nonetheless support claims such as “I have hands”. To take a temporal example, Paul Davies (1995) describes a scenario taken from Fred Hoyle’s sci-fi novel *October the First is Too Late*:

We are invited to think about events in the universe in terms of an unusual metaphor: a series of numbered pigeonholes containing messages about neighboring pigeonholes. But there is no “flow” of time. Instead there is a metaphorical clerk who inspects the pigeonholes one by one. Each act of inspection creates a moment of consciousness in the world. (Davies 1995, 41)

Davies goes on to note:

The curious thing about this imagery is that the clerk doesn’t need to sample the pigeonholes in numerical sequence. He could capriciously hop about all over the place, even at random, and we wouldn’t notice; we would all still have the impression of time as a continuous, ever-rolling stream. (Davies 1995, 41-2)

Such a random ordering of events is consistent with actual experience; nonetheless, most of us – presumably A theorists and B theorists alike – like to think that events are ordered in the sequence they appear to be ordered in, and that our experience as of a ‘linear’ ordering is evidence for it.

2. Could we experience the passage of time?

So we cannot so easily dismiss the A theorist’s appeal to experience. There is, however, a compelling argument against counting our experience as evidence, which I’ll recount here since, if sound, it has implications for the role of experience in metaphysics more generally. Simon Prosser (2007) ultimately aims to show that time does not pass, but crucial to his argument is a sub-argument, for the conclusion that we cannot experience the passage of time. This sub-argument is ostensibly rather simple, depending on just two premises.

P1. The conscious experiences of each individual subject nomologically supervene on

the physical state of the world.

P2. The (putative) passage of time would be epiphenomenal with respect to the physical state of the world.

C. From P1 and P2, we can conclude that the passage of time would be epiphenomenal with respect to experience, and hence that *we cannot experience the passage of time*. (Prosser 2007, pp. 78-9).

As Prosser suggests, these days, P1 doesn't require very much support. This is partly because, as supervenience theses go, P1 is rather weak. It says only that the mental is nomically dependent on the physical: as far as P1 is concerned, there may be *metaphysically* possible worlds in which the mental is independent from the physical. P1 also says nothing about whether the mental can be *explained* by the physical: as far as P1 is concerned, the supervenience of the mental on the physical may (or may not) be insufficient to say that the former can be adequately explained in terms of the latter. Finally, P1 says nothing about whether the mental can be *reduced* to the physical: it is hence open to property dualists to accept P1. Weak as it is, then, P1 should be acceptable to many, A-theorists and B-theorists alike. Indeed, it seems the only theorists who might reject P1 are interactionist dualists. (Prosser 2007, 78-79.)

P2 also seems plausible, though it perhaps requires a little more cashing out to see this. Prosser offers several ways of understanding what is behind P2. Perhaps the most central has to do with what is required to *account for* physical events.

By 'epiphenomenal' I mean that the passage of time neither causes nor in any sense influences or determines physical events. Insofar as physical events can be accounted for, the account is in terms of physics or at least in terms of what supervenes on the physical, and no appeal to the passage of time

plays a role in any such account. (Prosser 2007, 79.)

Another, closely related way of putting this is to say that all information relevant to accounting for physical events can be given in B-series terms.

One can describe the physical world and the way in which its state at one time depends on its state at other times in B-series terms; given an arrangement of matter at one time the nomologically possible arrangements of matter at earlier and later times are constrained only by the laws of physics and not in any way by real temporal passage. (Prosser 2007, 81.)

To put this another way, since temporal passage is a matter of metaphysics rather than physics, the laws of physics can be expressed in a way that makes no commitment regarding temporal passage. (Ibid. 82.)

Since we can account for physical events without invoking the A-series properties of *pastness*, *presentness* and *futurity*, we should, Prosser suggests, grant that these A-series properties have no role to play in *shaping* the physical world: hence, the passage of time would be epiphenomenal with respect to the physical state of the world (ibid. 81.)

So how does granting P1 and P2 force us to conclude that we cannot experience the passage of time? Well, since the passage of time makes no difference to what happens in the physical world, nor could it make any difference to our experience: for, given that our experience nomologically supervenes on physical states (P1), there can be no change in experience without a corresponding change in physical states. Thus, given P1, if the passage of time made a difference to our experience, this would contradict P2. Assuming that in order to have an experience of the passage of time (as opposed to an experience merely *as* of the passage of time) would require that real passage plays some role in shaping our experiences,

and that, by P1 and P2, passage cannot in fact play such a role, we are thus led to the conclusion that we cannot experience the passage of time.

Prosser goes on to argue that, since our experience as of the passage of time provides practically the only reason for us to think that time does in fact pass, we should conclude that time does not pass. I will not say much here about this latter part of Prosser's case against A-theory. However, it is interesting to note that, if we accept the argument outlined above without going so far as to say that time does not in fact pass, we can see that Prosser has given us an argument that effects a division between appearance and reality. According to P1 and P2, even if it is clear that we do in fact have experience *as* of the passage of time, and even if time does in fact pass, the former cannot in any way be *because of* the latter.

3. Experience of ordinary objects.

Other metaphysical debates in which experience might be thought relevant are those surrounding theories of composition and persistence. In the case of composition, there are theories that deny existence to some or all 'ordinary' objects. According to mereological nihilism, for example, there are no composite objects at all, only 'simples' (objects which have no proper parts). For the nihilist (e.g. Sider 2013), there are no tables, chairs, etc.; there are only *non-composite* objects – particles which themselves have no proper parts – arranged 'tablewise' and 'chairwise'. One might think that experience stands in opposition to mereological nihilism. Common sense favours the view that there are tables and chairs, and this, it might be said, is in no small part due the fact that we seem to *perceive* them.

What would it have looked like, if it had looked like there were no composite objects? One might think of a disembodied consciousness observing particles moving about independently

from one another, but even if this is a world in which it appears that there are no composite objects, it is not clear whether this would amount to an appearance of the truth of nihilism: had the particles been arranged tablewise, presumably it would appear as though there were composite objects; the only reason it appears that there are not composite objects in the world described is that there are no particles suitably arranged. To put it another way, the truth of nihilism is not needed to account for appearances at the world described; we have a preferable explanation in the fact that there are no simples suitably arranged.

In the case of persistence, there are theories according to which three-dimensional objects never exist at more than one time. According to perdurantism, for example, persisting objects are four-dimensional worms, composed of three-dimensional temporal parts. As Lewis has defined it, “Something *perdures* iff it persists by having different temporal parts, or stages, at different times, though no one part of it is wholly present at more than one time” (1986, 202). For an intuitive, if flawed, picture, we can think of temporal parts as analogous to still frames in movies, and perduring objects as the collection of frames arranged in order (provided that the narrative of the movie is linear). Change, on this picture, is thought of as qualitative variation across temporal parts.

Among perdurantism’s chief rivals is *endurantism*. For the endurantist, things persist by enduring, and as Lewis defines it, something *endures* “iff it persists by being wholly present at more than one time” (Lewis 1986, 202). Endurantism is typically thought of as the more commonsensical position. As Michael Rea put it, “It has fallen to the perdurantists rather than the endurantists to motivate their view, to provide reasons for accepting it that override whatever initial presumption there is against it” (Rea 1998, 225). Now, arguably, part of the reason that common sense stands on the side of endurantism comes from our experience:

things are represented to us as enduring. This is suggested by Prosser his 2012 paper, “Why Does Time Seem to Pass?”:

[C]hange is experienced as dynamic because the experience involves the representation of something *enduring* through the change. It is this notion of a *single entity* passing ‘through’ a change that captures at least a very important element of the experience of temporal passage (Prosser 2012, 16).

Prosser invokes this idea to help provide an alternative answer to the A theorist’s postulation of dynamic time. Thus the suggestion that objects are represented to us as enduring, rather than perduring, can be seen as B theorist-friendly, for it allows us to provide an account of why time seems to pass consistent with the truth of B theory.

This perhaps fits well with D.H. Mellor’s advocacy of both B theory and endurantism. Here is Mellor expressing what has been referred to as the “no change” objection to temporal parts.

A change ... is a thing having incompatible properties at different times ... That is, it needs a single particular for the difference to be a change *in* ... In short, differences between temporal parts fail to be changes for the same reason that differences between spatial parts do ... [The perdurantist account of change] can be ruled out by an identity condition (Mellor 1998, 89-90).

One might reasonably conjecture that Mellor’s rejection of the perdurantist notion of change has to do with the character of our phenomenal experience. After all, the perdurantist *does* claim that the persisting object is *identical* the sum of its temporal parts, but somehow, for those who argue as Mellor does, this is not identity of the right things.

Even if we disagree with the suggestion that things are represented to us as enduring, the

perdurantist faces another challenge mounted on the basis of the character of experience, in the form of an argument from the character of experience to perdurantism's four-dimensionalist rival, *stage theory* (influential expositions of which are Hawley (2001) and Sider (2001)). Both Yuri Balashov (2005) and, more recently, Josh Parsons (2015), suggest that stage theory *best explains* the (apparent) presentness of experience. Both authors suggest that the perdurantist makes the character of my experience puzzling, raising the following question: if perdurantism is true – if, that is, I am temporally located at every time of my life – then why is my experience as of a single time? By contrast, if stage theory is true, then I am not identical to the four-dimensional worm of my stages: instead, I am (now) identical to a momentary stage. The suggestion is that, simply by assuming the truth of stage theory, we have a ready explanation for the presentness of experience, and so we should opt for stage theory over perdurantism.⁴

Is there a certain way things *could* appear, which would – prima facie – positively support perdurantism? Whereas it seems that there is no way the world could be that would make it appear that B theory is true, and it is hard to imagine a world in which it would appear that nihilism is true, perdurantism perhaps shows a little more promise in this regard. By way of an account of temporal phenomenology along the lines developed (separately) by C.D Broad and Edmund Husserl, I would like to (tentatively) suggest that there are possible worlds in which objects are represented to us as perduring rather than enduring, such that the initial presumption arising from experience would be that objects perdure.

Crucial to the Broad/Husserl account is the notion of *presentedness*, a quality that the

⁴ Actually Balashov's infers to stage theory over both endurantism and perdurantism; Parsons's paper only considers the two four-dimensionalist theories.

contents of our experiences have in greater or lesser degree, depending on when they are experienced: “When a content is initially experienced it possesses this quality to the highest degree, as it is represented with subsequent contents it possesses less and less of this quality, until in the end it drops out of awareness altogether” (Dainton 2010, 109). Thus each act of awareness can contain a number of contents that vary in degree of presentedness. With each momentary act of awareness, the degree of presentedness a content has can be seen as a marker, showing when it occurred relative to other contents; contents with less presentedness will seem to occur before contents with more.

Dainton notes three major objections to the Broad/Husserl account; however, each is based in contingent claims about the character of our experience. The first objection is that the account is phenomenologically inaccurate. If the account was correct, we could expect experiences to never end abruptly, always lingering for some time after they first occur. We could also expect to be aware not only of our experiences, but “of a constantly shifting complex of representations of recent experiences, and so our consciousness would be *choked* with the residues of recent experiences” (2010, 111). As far as I can tell, Dainton is right that this seems to misrepresent the phenomenology; but it doesn’t seem to *necessarily* misrepresent it: there is a possible world in which these predictions of the Broad/Husserl account are confirmed in experience. Thus Dainton’s first objection is to the *actual* accuracy of that account – but it leaves open its possibility.

The second objection Dainton raises is that the notion of presentedness is problematic. We might think of degrees of presentedness as variations in phenomenal intensity – what Hume referred to as differences in “force and vivacity” (Dainton 2010, 112). But it is hard to see how this amounts to temporal variation. Colours, sounds, etc. can be more or less intense or

vivacious, but this has little to do with temporal phenomenology, for such variation can easily obtain across simultaneous contents – contents that are represented to us *as* simultaneous. In his 2006 book *Stream of Consciousness*, Dainton suggests that presentedness is best understood to be a *sui generis* property, but notes that “when we hear a sound while seeing a colour, we are aware of the auditory and visual characteristics of these contents, but we are not aware of any additional phenomenal characteristic that is common to both ... So the problem is there just does not seem to be any such property” (Dainton 2006, 150). This seems right – but again, it leaves open the *possibility* of such a phenomenal characteristic.

Thirdly, Dainton notes that, on the Broad/Husserl account, “successive acts [of awareness] apprehend numerically distinct contents,” but objects: “Since it is obvious that our experience is *not* confined to isolated momentary capsules, we must look elsewhere for a realistic account of time-consciousness” (Dainton 2010, 112). But again, even if this is right, it says nothing against the possible truth of the account, for there is a possible world in which it is far from obvious that our experience is not so confined.

Since each of these main objections to the Broad/Husserl account are based in certain contingent features of temporal experience, the possible correctness of that account is not ruled out. While this is of little interest to those who seek to account for our actual experience of change and succession, it can be relevant to the modal status of the phenomenological case against perdurantism. For a world in which the Broad/Husserl account is correct might be a world in which objects and experiences are represented as perduring, such that, *prima facie*, experience favours perdurantism.

As we saw, Dainton objected that the Broad/Husserl account seems to involve an *atomised*

account of experience that is not *actually* borne out; but *had* it been, it would perhaps be suggestive of a more temporally atomised world, containing temporally atomised objects; that is, such a world might be one in which objects are represented to us as perduring, rather than enduring. We might imagine a world in which the inhabitants are incredibly perceptive of change and temporal distance, to the extent that they are able to register their own instantaneous experiences, and to have a sense of their distinctness from those preceding and succeeding ones, in virtue of the varying degrees of presentedness attached to each. Insofar as it is possible for such fine-grained degrees of presentedness to be attached to experiences, it is plausible that similar degrees of presentedness could be attached to objects and occurrences. In such a world, varying degrees of presentedness might therefore be suggestive of the existence of numerically distinct temporal parts of an object which nonetheless exhibits all the usual markers of persistence: relative qualitative similarity over time, continuity of location, sameness of function, etc. (We might also stipulate that this world is one in which time is discrete rather than dense, in order that instantaneous parts of experiences precede and succeed each other in contiguity, and so that they would be able to perceive the ordering in time of temporal parts without having to process infinite contents. The next chapter discusses some ways in which assumptions about the micro-structure of time can have a bearing on the question of how things persist, and the question of whether time is substantival or reducible to the relations among objects/events.)

4. Prosser's argument adapted to perception of ordinary objects.

We have seen that there is an initial presumption against B theory, against mereological nihilism, and against perdurantism. Moreover, in the case of B theory and nihilism, it is hard to see how our experience *could* represent things so as to favour the truth of these theories. In the case of perdurantism, however, the case from experience against that theory may be a

contingent one – perhaps we can imagine a possible world in which our experience is atomised, such that it might stand in favour of perdurantism.

But should experience play *any* role in these debates? We might simply adapt Prosser’s argument against the possibility of our experiencing passage as follows.

P1*. The conscious experiences of each individual subject nomologically supervene on the physical state of the world.

P2*. The (putative) existence of ordinary composite objects (e.g. chairs) would be epiphenomenal with respect to the physical state of the world.

C*. Therefore, even if the character of our experience is such that it does in fact *seem* that there are composite objects, and even if there *are* in fact composite objects, the former cannot be in any way *because of the* latter. That is, while we can, and probably do, have experiences *as* of composite objects, we cannot experience ordinary composite objects.⁵

At first sight, this argument depends on a false premise. Would the existence of composite objects really be epiphenomenal with respect to the physical state of the world? Surely the existence of my desk has some role to play in shaping the physical world – after all, if it

⁵ It may sound a little awkward to talk of experiencing objects – perhaps only events can be experienced. In that case, simply replace ‘experienced’ with ‘perceived’. This is a substitution that Prosser is also happy to make in the case of passage: “I have said that time could not be *experienced*, but I am happy to use the word ‘perceived’” (2007, 77).

wasn't here, there would be nothing for my elbows to lean on.

But recall now the support Prosser gave for his second premise. One aspect of that support was that we can account for physical events without invoking A-series properties: everything about the shape of the physical world over time can be stated in B-series terms (ibid. 82).

Similarly, it might be said that everything about the shape of the physical world over time can be stated in the terms of the mereological nihilist, employing the familiar paraphrase strategy if need be: any statements usually containing apparent reference to composite objects are to be paraphrased in terms of arrangements of particles; so for example, my earlier reference to my desk can be paraphrased in terms of particles arranged desk-wise. If one can accept that no information relevant to physics is lost after such a paraphrase has been made, and accepts Prosser's support for the second premise of his argument about the putative experience of time, then, by parity of reasoning, one should also grant P2*, and hence conclude – just as Prosser concluded that we cannot experience the passage of time – that we cannot have experiences of composite objects.

The upshot is that accepting Prosser's argument, and the availability of a suitable paraphrase strategy, seems to commit one to a strange picture of experience: while we have experiences *as of* the existence of chairs etc., what we really experience is rather close to what the nihilist would say we experience: simples arranged chairwise. Indeed, Trenton Merricks gestures towards the same argument as part of his defence of eliminativism about ordinary macroscopic objects, in his book *Objects and Persons*: “Atoms arranged statuewise cause the visual and other sensations the folk ontologist thinks are caused by statues. Indeed, atoms arranged statuewise can do just about anything normally attributed to statues” (Merricks 2001, 9). We might think that, just as our experience as of passage is the main reason cited

for us to believe in the passage of time, likewise my belief in the existence of chairs etc. is grounded in my experience as of there being chairs etc. Hence, without this justification from experience, the case against nihilism (i.e. for the existence of composite objects) becomes far weaker.

This may cause some B theorists to think twice before endorsing Prosser's argument. Presumably B theorists will welcome Prosser's conclusion that we could not experience/perceive the passage of time, even if there were such a thing; but presumably, most B theorists – as most people – suppose that we *can* perceive ordinary composite objects.

There is, however, a way of maintaining Prosser's argument against perception of passage, while rejecting the argument adapted to the case of object perception, by disputing the support for P2*. Parsons has suggested that it is not simply the question of whether we *can* translate our scientific and proto-scientific theories into the eliminativist's terms that is important; we must also ask whether the eliminativist's proposed translation would leave us with an equally good theory.⁶ To illustrate, Parsons uses the example of a rock hitting a window, and the explanandum of why the window breaks.

A bunch of quarks and leptons (even the very ones of which the rock is composed) might as easily pass straight through a window as break it. What someone like Merricks has in mind, of course, is that if the particles arrive at the window arranged as they in fact are—arranged rock-wise, we might say—then they would break the window. But to turn this into a real explanation of the window's breaking, we need to turn 'arranged rock-wise' into a real microphysical description. Even if we could do that, we might not get any useful prediction or explanation out of physical theory, for the system of the rock and

⁶ I refer here to the eliminativist's "proposed" translation because, on the assumption that carrying it out would leave us with a worse theory, it would arguably leave us with a *different* theory, and as such would arguably fail as a translation.

window would be turned into an intractable many-body problem. ... For reductionistic explanations to render macroscopic ones redundant, they would have to be better explanations in every case. But they are not—in most cases they are more complex, even to the point of being intractable. (Parsons 2013, 333)⁷

The suggestion is that the eliminativist version of events is near-useless for the purposes of explanation and prediction, and so, insofar as the capacity of a theory to explain and predict is suggestive of its truth, we should opt for non-eliminative theories of such phenomena. This may represent a difference between the case of whether time passes and the case of whether there are macroscopic objects. In the latter case, the description makes a difference to the utility of our scientific theories about the physical world; in the former case, it seems, there is no such difference. Hence one can endorse Prosser's argument against the possibility of perceiving passage, without necessarily being committed to the parallel argument against the possibility of perceiving ordinary macroscopic objects.

How about the debate concerning how things persist? Again, we can adapt Prosser's argument to reach the conclusion that the character of our experience of objects is irrelevant to how things persist. The argument is as follows.

P1**. The conscious experiences of each individual subject nomologically supervene on the physical state of the world.

P2**. For any thing, how (metaphysically speaking⁸) that thing persists is epiphenomenal with respect to the physical state of the world.

⁷ I return to Parsons's argument in Chapter 5, where I discuss the potential for inference to the best explanation to lead us to necessity claims.

⁸ In the sense, that is, of whether it persists by enduring, perduring, or exdures, where something is said to exdure if it persists in accordance with stage theory (Balashov 2010).

C**. For any thing, how (metaphysically speaking) that thing persists is epiphenomenal with respect to one's conscious experiences.

Corollary. Even if objects do *in fact* endure; and even if objects do *seem* to endure, the latter cannot be *because of* – or in any other way be a truth-tracking representation of – the former. Thus any explanation of the representation of endurance in experience must be independent of the truth about how things persist.

We saw above that Parsons has provided a potential way of avoiding the conclusion that we cannot perceive ordinary objects. That was by rejecting the support for P2*, by suggesting that paraphrasing away sentences about rocks hitting windows (so as to mention only particles and their arrangements) would leave us with an intractable many-body problem; that is, physical theory is not indifferent to the metaphysics of objects, at least not at the level of description.

In the present case, however, it does not seem that there is an analogous way of blocking P2**. Whether we describe an object as existing wholly at different times with certain time-relative properties, or as having distinct temporal parts at different times with certain intrinsic properties, or as having temporal counterparts at certain times with certain intrinsic properties, it is hard to see how this would make a difference to the ease with which one can make explanations and predictions concerning that object.

So the Prosser-style argument remains forceful, against counting experience as evidence in favour of endurantism. Even if objects are represented to us as enduring, as was suggested

above, this should not be counted as favouring endurantism.⁹

What about the suggestion, made by Balashov (2005) and Parsons (2015), that experience actually favours stage theory over perdurantism? Recall that both Balashov and Parsons suggest that stage theory best accounts for the *presentness* of experience: the fact that our experiences, when they occur, are perceived as being present, not past or future (Balashov 2005, 295).

In my (2014), I briefly suggested that P1** and P2** provide a response to the Balashov/Parsons argument, as follows.

However, if Prosser's argument can be adapted to any of the rival accounts of persistence, then it seems that there is no argument to be had from the character of our experience here, either. For, if both premises are true, it must be that how an object persists has no bearing on the character of experience. Crucially, this would mean that, if things persist by *perduring*, this should not be considered to be *in spite of* appearances. For it is only when we expect that appearances should be somehow responsive to the metaphysics that the perdurantist presents us with something puzzling – something that demands an explanation. (Traynor 2014, 387)

I now suspect that my response fails to counter the Balashov/Parsons argument that experience favours stage theory, on two grounds. First, the *Corollary* noted above stands against the claim that the fact that objects persist by enduring could somehow help to make sense of the fact (supposing it is a fact) that objects are represented to us as enduring.

⁹ I use 'favour' here with Skow's point, mentioned in section 1 above, in mind: "Even if it is true that each of the theories is *consistent* with my experience being as it is, it does not follow that the fact that my experience is that way fails to *favor* one of the theories over the other. This is an instance of a general epistemological truth: a body of evidence can support some hypothesis without entailing that hypothesis" (2011, 361). What I have just suggested is not merely that alternatives to endurantism are *consistent* with our experience; the character of our experience of objects does not favour endurantism, even if objects do appear to us as enduring. The same applies to the following two paragraphs concerning stage theory.

Adapting *Corollary* to address stage theory (exdurantism) instead of endurantism results in

*Corollary**:

*Corollary**. Even if objects do *in fact* exdure¹⁰; and even if objects do *seem* to exdure, the latter cannot be *because* of – or in any other way be a truth-tracking representation of – the former. Thus any explanation of the representation of exdurance in experience must be independent of the truth about how things persist.

But *Corollary** addresses a different (and arguably false) explanandum – namely, that things are *represented* to us as exduring – than the one that concerns Balashov and Parsons: the presentness of experience. Roughly,¹¹ what both Balashov (2005) and Parsons (2015) are looking for is an answer to the question: why is my experience as of a single time? For Parsons, this question is especially pressing for the perdurantist, according to whom I am a four-dimensional worm. For the worm theorist, Parsons notes, “[l]ast year is just as real as, and I am located there just as much as this year. Why then are my experiences of events of this year and not of last?” (Parsons 2015, 239). This is a question not so much about how objects seem to persist, as a question about a feature of all experience, what I have referred to as its ‘presentness’, that is perhaps made puzzling by the perdurantist’s picture.

Secondly, my (2014) attempt to counter the Parsons/Balashov argument is based on a version of what Lipton (2004) refers to as the “familiarity model of explanation”, according to which,

a phenomenon is unfamiliar when, although we may know that it occurs, it remains surprising because it is in tension with other beliefs we hold. A good explanation shows how the phenomenon arises in a way that eliminates the tension and so the surprise. (Lipton 2004, 25)

¹⁰ See footnote 8.

¹¹ There are differences between Balashov’s (2005) Parsons’s (2015) formulations of the question, but I don’t think these are significant for my purposes.

In the present case, the tension in question is between perdurantism and our experience of presentness. In my (2014), I aimed to show that, given P1** and P2**, there is in fact there no tension, for the metaphysics has no bearing on appearances. But as Lipton goes on to note, the familiarity model of explanation does not give an adequate description of how we give and judge explanations. In order for our search for an explanation to be worthwhile, there need not be any evident tension between the phenomena to be explained and other beliefs we hold. Lipton illustrates with the simple example of the rattle sound in his car: “you [do not] have to convince me that, in fact, it is somehow surprising that there should be a rattle, in order to get me interested in explaining it. Surprise is often a precursor to the search for explanation, but it is not the only motivation” (Lipton 2004, 27). One can explain the familiar by way of an explanation that is itself unfamiliar, and one need not thereby be eliminating a tension in a way that the familiarity model seems to suggest. Thus the stage theory-based explanation of the presentness of experience may nonetheless be of merit, even if perdurantism does not render presentness surprising. (The same applies to Paul’s (2010) non-metaphysical explanation of our experience as of dynamic time, which we’ll look at in the next section. That B theoretic explanations of our temporal experience are worth pursuing does not hang on the idea that there is a tension between B theory and temporal experience that B theorists need to eliminate.)

However, even granting that my criticism of the Balashov/Parsons argument was based in an inadequate model of explanation, one can still hold that the comparative advantage of stage theory is smaller if perdurantism does not make the presentness of experience puzzling. Hence one would still like to know whether it could be that perdurantism makes the presentness of experience puzzling, given that C** is true. Recall that C** states that for any

thing, how (metaphysically speaking) that thing persists is epiphenomenal with respect to one's conscious experiences. How could perdurantism make the presentness of experience puzzling, if the way that things persist (whether by enduring, perduring, or exduring) is epiphenomenal with respect to conscious experience?

Perhaps the answer lies with the distinction between *how*, and *whether*, things persist. It may well be that how something persists is epiphenomenal with respect to the physical state of the world (which is what P2** says); but *whether* something persists is of course not epiphenomenal (e.g. if my chair stopped persisting now, I would be writing this from the floor). According to stage theory, I am a momentary stage, and it is momentary stages, not four-dimensional worms, that have experiences. Strictly speaking, these momentary stages do not persist. But P2** concerns persisting things, not momentary stages.¹² As such, P1** and P2** taken together cannot tell us that perdurantism and stage theory do not have different implications for my experience.

5. Non-metaphysical explanations, and their contingency

Another way of arguing against taking the character of experience as evidence for certain metaphysical claims is to find non-metaphysical explanations of the phenomena in question, thus blocking arguments from explanation to the truth of such a metaphysical thesis (along the lines of: metaphysical claim p would, if true, account for feature f of experience; therefore p is true). Drawing on cognitive science, Paul (2010) provides a potential explanation of the apparent 'flow' of time. A dot is shown on a computer screen, then vanishes; shortly after, another dot appears to the left of where the now vanished dot appeared on the screen, before

¹² Thus P2** and C** perhaps should have begun: "For any *persisting* thing, ..." I have left them as they were in Traynor (2014) in order to reconsider here the point I made there against Balashov (2005) and Parsons (2015).

vanishing; shortly after, a dot appears where the first one appeared, and the series continues in this way – left dot, right dot, left dot... Anyone observing this process will experience what cognitive scientists have called “apparent motion”, the appearance that a single dot is moving from left to right, rather than that distinct dots are appearing in quick succession. This phenomenon is familiar from the experience of watching movies. As Paul notes, the appearance of motion persists nonetheless when we are aware that what is in front of us is a series of static images, “as long as the succession is rapid enough and spatiotemporally close enough” (Paul 2010, 348).

Paul goes on to argue that this data helps provide a potential explanation for why it *seems* to us that time passes, an explanation that is compatible with B theory, and that rivals the metaphysical explanation that there *are* irreducibly temporal properties like *nowness* and *passage*.

The reductionist can then argue that, if the brain can create the illusion of flow in cases of apparent motion, then it can create the illusion of flow in cases of experiences as of passage. In other words, the reductionist can use the experimental facts involving apparent motion, apparent change, and apparent persistence to argue that, even though all she endorses is the existence of a static universe of a series of stages, this is sufficient for the brain to produce the illusion of motion and flow involved in the experience as of change. (Paul 2010, 353)

[W]hen we have an experience as of passage, we can interpret this as an experience that is the result of the brain producing a neural state that represents inputs from earlier and later temporal stages and simply “fills in” the representation of motion or of changes. (Paul 2010, 352)

In cases like that of the series of dots, what we find is that an appearance of motion is created whenever objects appear to us and are appropriately related. The idea is that the cognitive

system that assimilates these successive images to produce an appearance of a single object persisting through change (where in fact all there is is a succession of distinct images) works to produce the appearance of change and passage more generally. This blocks the A theorist from claiming that A theory provides the only (potential, reasonable) explanation for the appearance of passage.

Similarly, some empirical findings can be used to make trouble for arguments from experience to endurantism. In multiple object tracking (MOT) experiments, subjects are presented with several objects of the same size and shape. Some of the objects are identified as “targets”, before moving about in unpredictable ways. When they stop moving, the subjects have to identify which are the targets. Then the process is repeated, except occluders are added to the picture. What is found is that subjects are very good at tracking objects, even when occluders are in place.

As Scholl (2007, p. 567) puts it “[the principle of spatio-temporal continuity has] become wired into our minds, and helps to control our experience of objects as persisting in the world”. This is not only so when various kinds of occluders are introduced, but also even when observers *know* that the targets they are supposed to follow and identify as being one and the same object are in fact *not* one and the same object” (Benovsky 2015, 692).

So, just as we have a propensity to see movement where in fact there is only a succession of static images, we also have a propensity to see distinct objects as being the same. So long as the latter exhibit spatio-temporal continuity, they will *appear* to be one and the same object, even if we *know* otherwise. As in Paul’s discussion, then, these empirical findings suggest a non-metaphysical explanation for why it seems to us that things persist by enduring: not that apparently enduring things do in fact persist by enduring, but rather that they exhibit spatio-

temporal continuity. Since things satisfying the spatio-temporal continuity condition appear to be the very same object even when they are not, the explanation for the appearance of endurance in the MOT experiments can be applied to explain why things appear to endure, even if they in fact perdure.

Insofar as such evidence helps to block claims that A theory and endurantism provide the only reasonable explanations for the appearance of passage and the appearance of endurance, respectively, and helps provide plausible explanations for the B theorist and perdurantist, this evidence strengthens the case for each of the latter theories.

However, it is all too easy to overlook the *contingency* of the evidence that Paul and Scholl have appealed to. In order for these experiments to be worth conducting, it must be conceivable for them to have turned out otherwise. In Paul's case of the series of dots, alternating from left to right to left, and so on, we can imagine subjects not experiencing any motion, regardless of how fast the succession of images. In such a world, it would be less clear whether being rapid enough and spatio-temporally close enough suffices to engender an appearance of motion from a series of static images; thus *metaphysical* explanations for the appearance that time passes – perhaps explanations that involve irreducibly temporal properties, would remain strong contenders.¹³

¹³ The failure of such experiments to engender B theoretic explanations of appearances would perhaps, in turn, cause one to doubt one or other premise of Prosser's (2007) argument that we cannot experience the passage of time. As Le Poidevin notes, one potential response to that argument arises from an A theoretic explanation for why we do not perceive the distant past or the future: by claiming that "the presentness of an event is a necessary condition of its causal efficacy", the A theorist can simultaneously explain the temporal limits of perception while rejecting Prosser's premise (P2 above) that the passage of time would be epiphenomenal with respect to the physical state of the world (Le Poidevin 2007, 84-85). The idea is that, "[o]nce an event ceases to be present, it loses the power to set anything in motion, and that is why we do not perceive the distant past" (Le Poidevin 2007, 84). This is a metaphysical (as opposed to, say, cognitive scientific) explanation for temporal limits of perception, to rival the one we met earlier, from Balashov (2005) and Parsons (2015), invoking stage theory.

Indeed, the very effort to find alternative, non-metaphysical explanations seems to imply this. That is, citing the evidence from cognitive science seems to imply that the following assumption is in play: had the experiments turned out otherwise, the A theoretic and endurantist explanations of appearances would have remained strong contenders; and moreover, that the B theorist and perdurantist would have faced challenges that their rival A theorist and endurantist theorists did not – namely, to explain away appearances that are somewhat at odds with the picture presented by the former two theorists.

The point goes both ways. What should we say about the possible world at which the experiments cited by Paul and Scholl turned out differently, so that B theoretic and perdurantist explanations would not be supported by the data? I have said that, at such a world, arguments from experience, against B theory and perdurantism, in favour of A theory and endurantism, would have more weight. But what should they say about our world, in which the data is as Paul and Scholl describe? Insofar as the inhabitants of such a world can conceive of the experiments (or some along those line), and grant – as they should – that they could have produced the results that they in fact did, they should say that the argument from experience against B theory and perdurantism, in favour of A theory and endurantism, could have been significantly weaker than it is at their world.

6. Assuming necessity to refuse experiential evidence.

A final point on the potential for experience to bear on metaphysical issues. Merrick's briefly suggests another argument against taking experience as evidence for the claim that there are composite objects like tables and chairs.

Here is more support for that claim [that whether atoms arranged statuewise compose a statue is not straightforwardly empirical]. The fundamental question is not so much whether some particular alleged statue exists. That question might—sceptical scenarios aside—seem to be a matter of just looking and seeing. The issue is rather whether, in general, atoms arranged statuewise compose a statue. And whether or not they do so will be a matter of necessity. But it seems that this question of metaphysical necessity cannot be decided, one way or the other, simply by a trip to the museum or a ride down Monument Avenue. It must be decided on philosophical grounds. (Merricks 2001, 9)

The idea that matters of necessity are beyond the reach of experience is fairly common. Here is Stephen Biggs, for example:

Even if modal properties are wholly present in our world (cf. Armstrong, 1997; Kment, 2006), however, necessities/possibilities do not reveal themselves by causally affecting us. In order to reveal themselves causally, necessity and mere-contingent-actual-invariance would need to affect our sensory apparatus differently from one another. Of course, they do not: whether cordates are necessarily renates (perhaps per impossible) or merely contingently renates, the world will appear the same to vision, audition, etc. (Biggs 2011, 286).

[An account of our knowledge of necessities and possibilities] cannot hold that we find necessity in sensory experience without violating Hume’s famous observation that such experience cannot distinguish necessity from mere actual invariance: “When we look about us towards external objects,” writes Hume, “we are never able, in a single instance, to discover any power or necessary connection . . . We only find that the one does actually, in fact, follow the other” (Hume, 1748/1961, p. 351). (Biggs 2011, 289)

However, there is an exception to what Biggs says – there are some possibilities that reveal themselves by causally affecting us, it’s just that they emerge rather trivially, from what we know to be true of the actual world, and the fact that actuality implies possibility. (Biggs’s

discussion side-lines knowledge of these kinds of possibility, “[b]ecause the epistemology of actualized possibilities may be relatively straightforward” (Biggs 2011, footnote 1.)

Bearing this in mind, what Merricks said becomes problematic. Suppose our question was whether or not some particular alleged statue exists. As Merricks said, “That question might—sceptical scenarios aside—seem to be a matter of just looking and seeing” (Merricks 2001, 9). So suppose that question *is* a matter of just looking and seeing. If our trip to the museum does give us evidence that some particular statue actually exists, that amounts to evidence that it is not necessarily the case that atoms arrange statuewise do not compose a statue. That is because, more generally, if *p* is necessarily true if true, and we have evidence that not-*p*, then we have evidence that not-necessarily-*p*. Thus, Merricks’s argument will not work unless the question of whether some particular statue exists is *not* a matter of looking and seeing.

Merricks will not be too concerned here – as we saw above, Merricks provides an argument earlier on, against the idea that experience tells us that there are statues. It’s just unclear whether the argument we are considering in this section – from the claim that the debate about composition is a matter of necessity, to the claim that experience has no bearing on the issue – adds anything to Merricks’s case.

Perhaps more significantly for my purposes, there are dialectical constraints on the applicability of such arguments from the supposed necessity of the matter. Suppose that experience can count as evidence for a particular metaphysical claim, such as that objects persist by enduring, but that the character of our experience could have been otherwise. This might then be used to argue for the contingency of the matter: there are worlds in which the

experiential evidence counts against endurantism, and hence the case for endurantism varies depending on what world you are in. This might be taken to support the contingency of the matter. Roughly, the idea is that, when the case for p (over rival claims) varies across modal space, we should say that it is a contingent matter whether or not p. That is to suggest, there may be a route from the claim that experience is evidence on a certain metaphysical question, to the contingency of the answer to that question. Any argument from cross-world variation in the experiential evidence on a certain matter, to the claim that that matter is one of contingency cannot be legitimately countered by an argument from the supposed necessity of that matter, to the claim that experience does not count as evidence with regard to it. This is not a point against Merricks, because Merricks is not writing in the context of a debate over whether questions surrounding composition have necessary or contingent answers. I just wanted to point out that an argument along the lines that we are discussing here must be discounted if there is a line to be drawn from the assumption that experience does count as evidence in a certain matter of dispute, to the claim that that matter is a contingent one.

7. The conceptual and empirical case against A theory.

In this section I look at the impact that special relativity has had on the debate between A and B theorists, and what the B theorist's appeal to special relativity suggests about the shape of the dialectic. In particular, I suggest that, the idea that special relativity makes a significant dent in the prospects of A theory seems to presuppose that the conceptual case against A theory is not sufficient – so, those whose case against A theory is primarily based on physics should grant that worlds where the data do not favour special relativity will be worlds in which the prospects for A theory are significantly stronger; which, on my view, is ultimately suggestive of contingentism in this debate. (The interplay between conceptual arguments and empirical evidence is also an important theme in the next chapter.) The section

concludes with a brief discussion of underdetermination in physics and metaphysics.

McTaggart (1908) famously argued that the A theory of time is incoherent. The properties of futurity, presentness, and pastness that events are mutually exclusive, yet somehow each event must have at least two such properties if it is to count as moving through time (or perhaps rather, if time is to count as passing).¹⁴ The A theorist will naturally respond that, of course, no event has any two of these temporal properties at any *one* time; rather, present events *were* future, are *now* present, and *will* be past. That is, any present event is future in the past, present in the present, and past in the future.

The trouble is, when McTaggart contends that, in order to be counted as moving through time, each event must be past, present and future, he is not just talking about present events when they are present. Take some present event E, and suppose that it has existed in the past, and will exist in the future. At all points of its existence, the A theorist contends, E is moving through time. As such, McTaggart contends, E must have had the properties of being past, present and future at *every time* at which it exists; as such, the alleged incoherence must be removed from the past, present and future of E. This means that the present strategy amounts to applying all nine of the following combinations of predicates to each event:

Is past		past
Is present	in the	present
Is future		future.

¹⁴ We might just as well talk about temporal properties of times, instead of events: a time can have the properties of being past, present or future.

The trouble is, this only leaves us with further incoherence. Since iterations of the predicate ‘in the present’ yield equivalent statements, among the nine statements that result are three incompatible ones: E is past in the present, is present in the present, is future in the present. That is equivalent to saying that E is past, present, and future. Crucially, the fact that iterations of ‘in the present’ yield equivalent statements means that ascending to a higher level of predication will not remove the contradiction.

This is an example of a conceptual argument against A theory. Far from refuting A theory, however, McTaggart’s argument sparked a debate that is ongoing. As mentioned above, Craig claims that “[Experience is] a defeater-defeater that overwhelms any B-theoretic arguments against the reality of tense” (Craig 2000, 138). E.J Lowe has argued that McTaggart’s argument commits an “indexical fallacy”: while it does make sense for someone in the future to say “Event E is happening now”, it does not make sense for us to say “E is present in the future” – to do so would be analogous to saying “that an event *e*, which is not happening *here*, but is happening over *there*, is happening *here over there*” (Lowe 1987, 66). Craig Bourne contends that McTaggart has shown “not that tense is unreal but that only one of the three tenses can have instances” (2006, 76) – thus paving the way for presentism, a version of A theory that denies the existence of past or future objects and times.

The special theory of relativity (STR) is often thought to present a fundamental problem for presentism. For the presentist, the present is ontologically very significant: only present objects exist. According to STR, whether two events are simultaneous is frame-relative; whether two events both count as present depends on an arbitrary choice of reference frame. This means that, according to presentism, whether an object exists also depends on an arbitrary choice of reference frame. This unpalatable consequence has led many to claim that

STR is incompatible with presentism, and thus that presentism is actually false (to take just a few: Putnam (1967); Weingard (1972); Mellor (1974); Shimony (1993); Callender (2000); and Saunders (2002)).

The impact that STR has had on the debate is interesting for several reasons. It provides an excellent case for investigating the extent to which science can bear on metaphysical questions, and vice versa, and the assumptions that underpin that interplay, but in order to focus on modal issues, I'll assume that science can bear on the truth or falsity of metaphysical claims (as Hawley notes – as we saw in the introduction – this assumption is best understood as being made from a standpoint of scientific realism (Hawley 2006, 460)).

We have seen McTaggart's conceptual argument against A theory, which can be read as an argument against the *possibility* of time passing. But we have also seen that McTaggart has not had the last word (no one has). What impact does STR have on the conceptual case against A theory?

Think of a theorist who endorses McTaggart's argument. What does STR add to this theorist's case? It is generally thought that theories like STR are contingently true, if true.¹⁵ Thus, we can approach our question a little more easily by raising another: Would this theorist's case be significantly weaker had STR been false?

¹⁵ Presumably, this is partly due to the conceivability of a world in which the predictions of the theory are not borne out; which is to say, it is due to a consideration of possible evidence. As we'll see in Chapter 4, consideration of imagined evidence can do little to convince dispositional essentialists – who argue for the necessity of laws of nature – that the actual laws are contingent; but this does not render a principle according to which possible evidence is a guide to possibility impotent, for such a principle can be used to convince those who do grant the contingency of the laws of nature of further metaphysical contingencies that they have not already assented to.

If the answer is ‘no’, then STR adds little to the case against presentism. So suppose that the case against presentism *would* be significantly weaker had STR been false. Then it is not so clear that McTaggart’s argument should be taken to show the impossibility of presentism, for there may be worlds in which other evidence for presentism – evidence from experience, for example – outweighs arguments against it. (After all, reliance on STR to make the case against presentism would seem to suggest that the conceptual case against presentism is not that strong. Mellor seems to recognise this when he says that “[t]he support that relativity and modern cosmology give to the B-theory of time, while very welcome, is by no means necessary” (Mellor 1998, 57).)

If STR does have *modal* implications for the debate, they are more modest. Suppose STR entails the actual falsity of presentism. This means that our evidence in favour of STR is also evidence that presentism is not necessarily true. This would be a blow to any presentist who contends that presentism is necessarily true.

Things are complicated further, however, by the suggestion that STR does *not* entail the falsity of presentism. In relation to whether quantum physics can bear on the metaphysics of identity, Stephen French has observed that “[there is] a problem for this programme of ‘reading metaphysics off current physics’, to put it crudely, which arises from what might be called the ‘underdetermination’ of metaphysics by physics” (French 1998, 93). The underdetermination of metaphysics by physics applies no less to the case of presentism and STR. Partly this is due to the underdetermination of physical theory by the data, which incidentally means that it is open to the presentist to adopt an alternative to STR, such as Hendrik Lorentz’s (undetected) aether theory, on which there is a privileged reference frame (stationary in the aether) – thus choosing the physics to fit the metaphysics.

But the metaphysics can also be made to fit the physics. As Sklar has said, “The science can change the philosophy and put the dispute in a new perspective, but it cannot resolve the dispute in any ultimate sense” (Sklar 1974, 275). The presentist is not forced to reject relativity, but nor is the adherent of STR forced to reject presentism wholesale. Putnam (1967) argues that, given the truth of STR and certain other plausible assumptions, we can prove that B theory is true – which is to say that all STR worlds are B theory worlds. However, one of these assumptions is that the relation R, where Rxy is ‘y is real with respect to x’, is transitive. This is an assumption, however, that a presentist might reasonably reject. As Dainton notes, “anyone who believes that time and existence are intimately related might be very tempted to suppose that if simultaneity is frame-relative, existence must be too” (Dainton 2010, 331). For example, Dainton describes a version of presentism – “cone presentism” – according to which the real is relative to a single spacetime point, and only the surface of the past light cone of an object is regarded as delineating that object’s present. In one respect, this version coheres with common sense: the fact that the speed of light is finite does not entail that everything one sees is in the past; rather, given cone presentism, everything that one sees is to be counted as present (Dainton 2010, 335-336).

Unless there is an argument to convince everyone that the relation R must be transitive, it thus seems that STR is not incompatible – in the sense of being *inconsistent* with – presentism, provided the latter theory is appropriately construed. It may of course be argued that B theory is nonetheless *favoured* by STR. Perhaps the point of coherence with common sense noted in the previous paragraph will be thought offset by the fact that cone presentism renders temporally distant events as present (indeed, that is why everything one sees is counted as present on cone presentism – the notion of the present is stretched to include as

present all events on the surface of one's light cone). In that case, it might remain that evidence for STR is evidence against presentism; it's just that we should grant that there are some STR worlds at which presentism is true. Suppose that STR favours the actual falsity of A theory – what should we say about other worlds in which STR is true? Presumably it would be just as reasonable for any inhabitant of an STR world to deny presentism. But if STR is compatible with a coherent version of presentism, then there must be some world at which STR and presentism are both true. So even if every STR world is one in which it is reasonable to deny presentism, it cannot be that all STR worlds are ones in which presentism is false.¹⁶

8. Concluding remarks.

In Chapter 4 I'll defend what I have called the Possible Evidence Principle (PEP): If there is a possible world at which the observable facts make it objectively reasonable to conclude that p, then it is objectively reasonable for us to conclude that p is possible. It is a consequence of this principle that, if the observable facts vary from one world to another in such a way that they make it objectively reasonable to conclude that p at some worlds, and objectively reasonable to conclude that $\neg p$ at others, then we should say that that claim is contingent – true at some worlds, false at others. If there are worlds in which the character of experience can be counted in favour of a particular position in the debates concerning time, persistence or composition, and worlds in which the character of experience can be counted in favour of some contrary position, or ought to be discounted altogether, then on my view, this provides us with some reason for thinking that these debates are contingent matters.

¹⁶ Relatedly, we'll see in Chapter 5 that a theory can be the most virtuous in every world (of a certain type) and still be contingent, where I discuss a modal analogue of the lottery paradox.

I have suggested that we need only appeal to empirical evidence where the conceptual case is not sufficient for what we wish to argue for (as may be the case for the B theorist's arguments against A theory). Where appeal to empirical evidence is meant to make a significant contribution to the case for a certain metaphysical thesis, this can be seen as tacit admission that the conceptual case for that thesis needs help. But while a conceptual argument might lead one to a necessity claim, it is hard to see how appealing to contingent facts can *add* to the case for the *necessity* of some claim; at best, actual contingent facts can suggest contingent actual truth. As such, it is interesting to consider what the possible absence of such empirical evidence means for the modal status of a metaphysical thesis whose plausibility hinges on such evidence.

As well as supplementing conceptual arguments, (possible) empirical evidence can also be cause for doubt about a conceptual argument. We'll see this in how evidence pertaining to the structure of time can be problematic for temporal parts explanations of possible change, and for relationism about time (Chapter 2); how quantum mechanics can be problematic for the identity of indiscernibles, thus giving us cause to doubt the bundle theoretic argument for it (Chapter 3); and how possible inductive inferences can lead us to the conclusion that time passes without change (Chapter 4).

2. Actual Time and Possible Change: On Modal Motivations for Mereological Harmony.

The main purpose of this chapter is to present a further instance in which a metaphysical theory can vary in plausibility, depending on contingent facts. I do this by presenting a problem for modal arguments for the view that spatial and temporal parts must be as fine-grained as space and time. Various solutions to the problem are discussed, and each has its problems, but ultimately I suggest that taking the metaphysics of persistence – and for analogous reasons, the question of whether time is relational or substantival – to be a contingent matter is the best way to accommodate the contingency of evidence concerning the structure of time at different worlds.

1. Grounding Possibilities in Parts: Sider’s Desk, Hawley’s Homogeneous Object.

Anyone who finds herself initially drawn to a temporal parts picture of reality will sooner or later be faced with the question of how to best formulate a temporal parts theory. One issue concerns how fine-grained objects are. Two of the primary motivators of temporal parts are the problem of change and the paradoxes of coincidence; temporal parts provide neat solutions to these problems, and, to provide for these solutions, it initially seems that we need only posit temporal parts as fine-grained as *actual change*. However, there is something of a consensus among advocates of the four-dimensionalist picture (Lewis 1986; Hawley 2001; Sider 2001) that the temporal parts of persisting objects are as fine-grained as *time* – that is, any object has as many temporal parts as there are times in the interval through which it persists.¹⁷

¹⁷ In this chapter, I will largely be ignoring the distinction between stage theorists (Sider (2001) and Hawley (2001)) and perdurantists (Lewis (1986)); the following discussion applies equally to both kinds of temporal parts theorist. For an alternative formulation of temporal parts theory, where the briefest temporal parts are non-instantaneous, see Butterfield (2006).

Four-dimensionalist reasoning is traditionally driven by analogies with the spatial case (as such, everything said here about the temporal case will be taken to carry over to the spatial case, and vice versa). One such argument suggested by Sider, on the issue of fine-grainedness, is rather direct: Since Sider's desk is spatially extended, we have, by way of analogy, a *prima facie* case for the claim that objects are likewise temporally extended; and if spatial extension implies the existence of spatial parts, then, by analogy, we have a *prima facie* case for the claim that temporal extension implies the existence of temporal parts (Sider 2001, 88).¹⁸ However, the antecedent of that last conditional will not be warmly received by friends of extended simples. In particular, if the suggestion is that it is a conceptual truth that to be extended is to have proper parts, the suggestion is refuted by anyone able to conceive of extended simples (Parsons 2007, 211-212).

More promising for the fine-grained-as-time view are certain modal arguments, which Sider goes on to advance, and which can also be found in Hawley (2001). Considering his desk again, Sider says,

Suppose the desk is divided into halves. ... Since the halves are spatially separated, they are distinct objects. These objects were presumably not created by division, so they must have existed before division. But surely it does not matter whether the division ever *actually* occurs, so anything *potentially* divisible must actually have parts. (Sider 2001, 90).

¹⁸ This idea, that objects have spatial and temporal parts that mirror the structure of the spatio-temporal region at which they are located, I refer to as 'mereological harmony'.

Continuing with the analogy between space and time, Sider suggests that anything potentially *changeable* must actually have *temporal* parts. Hawley argues similarly. Considering a spatially homogeneous object, she says,

If the homogeneous object does not actually have spatial parts, we face a problem. The object could have become inhomogeneous midway through its career – someone might have painted stripes on it. Then, presumably, it would have had parts. But would it have had parts before the painting, because of the later painting, or does painting the object create its parts? Neither option is attractive, and both can be avoided if we accept that the homogeneous object actually has spatial parts. Moreover, if the object actually has spatial parts, this explains why it *could* have been inhomogeneous – its different parts *could* have had different properties. Similarly, if stages are as fine-grained as possible change, not merely as fine-grained as actual change, this explains how an actually-unchanging object could have changed – its different stages could have had different properties. (2001, 49).

In order to account for possible change, temporal parts (or stages) must be as fine-grained as possible change. Possible change is as fine-grained as time. Therefore, temporal parts are as fine-grained as time.

How fine-grained is time? Following Hawley, let's consider three potential candidates for the micro-structure of time – three alternative orderings of the instants: discrete, dense, and continuous.

If the ordering is discrete, then it is like the ordering of the integers, and every instant is followed by a unique next instant. In that case, finite stretches of time are not infinitely divisible, but are made up of discrete smallest units. If the ordering is dense, it is like that of the rational numbers (the fractions), and instants are not followed by unique next instants or successors. Between every two instants there is another instant. If the ordering is continuous, then it is like that of the real numbers (the rationals together with the irrationals), and, again, between every two instants there is another instant. (Ibid. 51).

For any given interval of time, each of these possible orderings, as enumerated, posits more points on the time series than the last. As Hawley notes, it is an empirical matter which of these orderings actually corresponds to the micro-structure of time;¹⁹ hence, if the structure of objects mirrors that of space and time as Hawley and Sider have both suggested, then the structure of objects – their fine-grainedness – is likewise an empirical matter (Hawley 2001, 52).

2. Actual Time and Possible Change: The Problem Stated.

Now let's take a closer look at the argument that emerged from the modal considerations in Sider and Hawley.

P1. In order to account for possible change, temporal parts must be as fine-grained as possible change.

P2. Possible change is as fine-grained as time.

C. Therefore, in order to account for possible change, temporal parts must be as fine-grained as time.

As we have seen in the above quoted passages, part of the motivation for P1 comes from an appeal to intuitive claims regarding explanatory priority. In order to make sense of actual change, one might think it is enough to posit temporal parts as fine-grained as actual change – if a persisting object does not change at all, then we need only say that it has one temporal part: itself. But if that same object *had* changed, (the temporal parts theorist must say) it

¹⁹ See Newton-Smith (1980) and Forrest (1995) for arguments to that effect; Dainton (2010), Chapters 16, 17, and Section 22.7, provides an accessible survey of relevant considerations from physics, in the midst of metaphysical discussion.

would have had *proper* temporal parts. Yet this seems to involve positing what are intuitively rather strange acts of creation: on this picture, it sounds like an act of changing an object (say, painting stripes on it) *creates* temporal parts, as though things have temporal parts because they change. Instead, what we wanted was to invoke temporal parts to explain change: change is *possible* because things have temporal parts. Positing temporal parts as fine-grained as possible change seems to get things in the right order of explanatory priority.

This thought can be further fleshed out in a way that reveals just how central P1 is to the temporal parts theorist's motivations. In the previous paragraph we took it that the mere possibility of change is something in need of explanation. But why would that be? The answer has to do with the indiscernibility of identicals. We take it as a datum that, whether or not persisting objects do in fact change, they nevertheless *can* change; that is, much change is, of course, possible, even if non-actual. But this needs explaining, because the eminently plausible principle of the indiscernibility of identicals suggests otherwise. Merricks, for example, notes that endurance combined with the indiscernibility of identicals seems to make change *impossible* (Merricks 1994, 166). Again, considering Leibniz's Law, $\forall x \forall y (x=y \rightarrow \forall F (Fx \leftrightarrow Fy))$, Andre Gallois says "The trouble with (LL) is that it rules out change" (Gallois 1998, 36). (As Gallois goes on to note (ibid. 37), the endurantist might remedy this problem by offering some suitable account of property instantiation, such as the relations-to-times account; however, it is usually by way of a certain distaste for such accounts that a temporal parts ontology comes to be entertained in the first place (Lewis 1986, 204).) The principle does not merely tell us that identicals *are* not discernible; rather, it tells us that identicals *cannot* be discernible. The suggestion is that the principle has modal import, and one sometimes sees this suggested in the very statement of the principle: "if x and y are

numerically identical then x and y *must* possess the same genuine properties” (Dainton 2010, 260; the emphasis is my own).

P1 thus seems to be a very natural premise for any temporal parts theorist to adopt. However, turning now to P2 (“Possible change is as fine-grained as time”), an interesting problem arises for the temporal parts theorist. The trouble begins with the observation that P2 is at best contingently correct. To see this, suppose, for the sake of argument, that ours is a world of discrete time. Also, let’s grant the claim that temporal parts are as fine-grained as time, mirroring the actual micro-structure of time in its discreteness. Now, plausibly, ours could have been a world in which time is dense – I, for example, could have existed in a world in which time is dense. Of course, it may be said that it is not *physically* possible for me to exist in a world in which time is dense, but, I suggest, it is at least *metaphysically* possible. But, if it is the case that I could have existed in a world in which time is dense, then I could have instantiated change more fine-grained than actual (discrete) time. But if *that* is true – if possible change is more fine-grained than actual time – then what P1 is really demanding is that temporal parts be more fine-grained than time itself.

Somewhat more formally, I am suggesting the following reductio:

P1*. Temporal parts are as fine-grained as possible change.

P2*. If actual time is discrete, then possible change is more fine-grained than actual time.

P3*. Actual time is discrete.

C*. Therefore, temporal parts are more fine-grained than actual time.

Evidently P1* is endorsed by Sider (2001, 90) and Hawley (2001, Section 2.4). P2* is supported by the thought that, even if an object actually has temporal parts that mirror discrete time, it *could* have existed in a world in which time is either dense or continuous; that is, it could have existed in a world in which time is more fine-grained than time actually is. If that much is granted, it is but a short step to what P2* affirms; for if the object could have existed in a world in which time is dense, then it could have instantiated change more fine-grained than discrete-time. Meanwhile, P3* is of course not obviously true – indeed, it may well be false. But the four-dimensionalists surely did not intend to foreclose any of the candidate orderings of the micro-structure of time as being actualised – and certainly did not intend to rule out all but the *most* fine-grained structures as being actualised.²⁰ And yet we see in this argument that, by assuming in P3* that time is discrete (along with P1* and P2*), we are led to a rather odd conclusion (C*). Even if this conclusion is intelligible, it is certainly counterintuitive, and surely one that few temporal parts theorists would be willing to endorse.²¹

In the next section, we will consider some potential objections to the problem just raised. I take it that none of the objections, or any of the replies, is decisive. But they will serve to

²⁰ It seems natural to think that there could not be more than continuum-many times. However, if we begin with a continuous time series, then interpolate time-branches at every point, with branching at every point of every branch, this would yield 2^{\aleph_1} many times (thanks to Stephen Read for pointing this out). However, it is not clear that the possibility of such extreme branching would have any implications for change on a temporal parts picture: for the perdurantist would naturally urge that, while *time* might branch in this way, persisting objects are invariably linear. More worrying for the temporal parts theorist would be Nolan’s suggestion of “Hypergunk” (and its temporal analogue), “gunk such that for any set of its parts, there is a set of strictly greater cardinality of its parts” (Nolan 2004, 5). For the most part, I will continue to speak as though “discrete, dense, continuous” exhausts the possible orderings, but note that if Hypergunky objects are possible, this could render any attempt to account for possible variation by way of actual parts problematic.

²¹ Cf. Parsons (2007), on the spatial analogue of C*: “To say “There are things smaller than any region of space”, is to treat space like a kind of ghostly matter that happens to inter-penetrate all material things, but which might or might not have parts that match the sizes and shapes of material things. This immediately raises the question of why we should believe in such spooky stuff, since it seems objects could have all sorts of geometrical properties whether or not there are places around to be those objects’ exact locations” (210).

elicit the kinds of commitments that the temporal parts theorist must take on, or the concessions she must make, if she is to continue to advocate the modal arguments that argue from the mere possibility of change to actual parthood.

3. Potential Objections.

Objection 1: The temporal parts theorists need only account for physically possible change. While it may be metaphysically possible for an object in a discrete-time world to exist in a world in which time is more fine-grained than that (i.e. dense or continuous), it is not physically possible. Sider said that anything *potentially* divisible must actually have parts; but by ‘potentially divisible’ he meant *divisible within the bounds of the physical laws of the world in which the desk actually exists*.²² For an object to instantiate change more fine-grained than actual time it would have to exist in a world with different laws of nature, but what might happen to objects in such *metaphysically* (but not physically) possible worlds is not something that a temporal parts theory is required to account for.

Reply: As with any such attempt to scale back P1, this last claim would need to be substantiated.²³ What is the difference between metaphysically possible change, on the one hand, and physically possible change, on the other, such that the metaphysics of persistence is required to account for the latter but not the former? I will not say that there is no such

²² Cf. Markosian (1998), Section IV: Simples as Indivisibles.

²³ Another way of scaling back P1 is to say that, for any world *w*, parts in *w* are only required to account for possible inhomogeneity/change in worlds where the structure of space/time is *not more* fine-grained than it is at *w* (thanks to an anonymous referee for *Thought* for this suggestion). This is to propose an interesting asymmetry: it means that, for instance, if the actual world is one of continuous time, actual parts *are* needed to explain possible change in worlds of discrete time; yet, if the actual world is one of discrete time, actual parts are *not* needed to explain possible change in worlds in which time is continuous. To this suggestion, we should ask, why is there such an asymmetry in the demand for explanation? To avoid being ad hoc, something more would need to be said here – particularly since the original demand for an explanation of possible change came from the indiscernibility of identicals, and said nothing of possible space/time topologies. That said, I will consider a further proposal to scale back P1 below, in section 5.

difference to be found, but I cannot think of a good one. Moreover, it is unlikely that Sider intended to make such a distinction. He claimed that since the desk could be divided every which way, we should say it has parts, both spatial and temporal, as fine-grained as space and time. But is it really *physically* possible that the desk be divided arbitrarily? Perhaps not; rather, I suggest, Sider's modal claim about possible desk division is only plausible enough to support the argument if it is a claim about what is *metaphysically* possible.

Objection 2: Discrete time is metaphysically *impossible* (for proof, see Zeno's stadium paradox). Hence P3* is necessarily false, in the strongest sense.

Reply: The argument for C* would work just as well if we replace the instances of 'discrete' with 'merely dense'.²⁴ For even if time is actually merely dense, I would urge, it could have been continuous. Yet Zeno's stadium paradox does not apply to dense orderings. Moreover, as Hawley notes, it has been argued by Newton-Smith that the success of our best physical theories does not decide between the competing hypotheses that space and time are merely dense, on the one hand, and that they are continuous, on the other (Hawley 2001, 51; Newton-Smith 1980, Section VI.3).

Objection 3: Newton-Smith was wrong. As Penrose has observed, in all successful dynamical theories, "on the small end of the scale, it is the entire range of real numbers that is in principle being made use of. The ideas of calculus underlie other physical notions, such as velocity, momentum, and energy. Consequently, the real-number system enters our

²⁴ Since continuous orderings are also dense, I use "merely dense" to refer to orderings which are dense but not continuous.

successful physical theories in a fundamental way for our descriptions of these quantities also” (Penrose 2004, 61; quoted in Dainton 2010, 271).

Reply: Be that as it may, current physics hardly warrants closing the book on the structure of space and time (Dainton 2010, Section 22.7), and it would surely be a mistake to be moved to thinking otherwise simply to maintain a version of temporal parts theory that can claim to account for possible change. But more to the point, let’s suppose that we do somehow have certain knowledge that time is continuous: even so, does the temporal parts theorist only want to claim that worlds like the actual world – worlds of continuous time – are populated by temporal parts? That would certainly be an option, but it would represent a significant concession on the part of certain temporal parts theorists;²⁵ moreover, it would be an interesting concession at that, effectively being an admission that facts about composition and persistence are contingent. We will return to this potential upshot in the next section.

Objection 4: The physical laws are metaphysically necessary, as according to the dispositional essentialism of Bird (2007), so it is not really possible for any object of this world to exist in a world where the micro-structure of time is any different. While dispositional essentialism allows that there are possible worlds in which the laws of nature are different from our own, it says that objects in these alien worlds do not represent genuine possibilities for objects in the actual world. The intuition that it is metaphysically possible for

²⁵ For Sider’s part, this can be seen in his discussion of the homogeneous rotating disc argument, an argument supposed to show that temporal parts theorists cannot account for motion in certain possible worlds containing homogeneous substances. Sider says: “Someone who defended four-dimensionalism’s truth in the actual world might claim that three-dimensionalism must be true at the disk-worlds. But this would involve granting the possibility of endurance, and thus giving up on the possibility of purely conceptual arguments in favor of perdurance... A four-dimensional solution must be sought.” (2001, 226).

an object in a discrete-time world to instantiate change more fine-grained than discrete-time is merely epistemic.

Reply: This would certainly block my argument, as it is effectively a denial of P2*.

Nevertheless, adopting dispositional essentialism is a substantive commitment. Of course, one may find that theory independently plausible. But Sider and Hawley no doubt intended their formulations of temporal parts theory – positing temporal parts as fine-grained as possible change and therefore as fine-grained as time – to be plausible without having to invoke this contentious account of natural laws.

Admittedly, the options that we have just considered are not exhaustive. But they are perhaps enough to indicate that the temporal parts theorist – who would argue from modal premises to claims about actual parthood – has more explaining to do. As we saw in Objection 3 above, one option is to insist – or rather, gamble – on the claim that space and time are continuous. However, it would be implausible to suggest this as a metaphysically necessary truth – at least not without extensive argument. As such, the temporal parts advocate, according to whom temporal parts are as fine-grained as time, may have to concede that it is not necessarily the case that temporal parts can account for possible change.

4. The contingent explanatory power of temporal parts.

If it is incumbent on a temporal parts theory to explain possible change, grounding it in facts about objects as they are in the actual world, then we ought to hope for an explanation of *metaphysically* possible change. Yet we have seen that, on certain assumptions about the

nature of time (and space), the fine-grainedness of possible change (and division) outstrips that of actual time (and space). For example, if time is discrete, we cannot ground all intuitively possible change in actual parthood unless we make the highly implausible – perhaps conceptually incoherent – claim that objects have temporal parts more fine-grained than time.

Does this mean that we should abandon the hopes of Sider and Hawley to have temporal parts account for possible change? Not necessarily. Notice that the argument leading to C* essentially involved an assumption that ours is a world in which time is not *actually* as fine-grained as it *could* be. For objects in such worlds (i.e. worlds of discrete or merely dense time), there is little hope of grounding facts about metaphysically possible change in facts about fine-grainedness of objects as they are at those worlds – for, in such worlds, there are not enough times.

But what if ours is a world in which time really is as fine-grained as metaphysically possible change? Plausibly, metaphysically possible change is as fine-grained as the continuum, so let's suppose ours is a world of continuous time. In that case, it is no longer clear that the desired explanation is unavailable: taking temporal parts to be as fine-grained as time in a world of continuous time *does* seem to give us enough temporal parts to ground the possibility of change, no matter what possible world is being accessed. That is to suggest, the modal facts that the temporal parts theorist can account for could well depend on metaphysically contingent facts about the actual world. If time is continuous, then it appears temporal parts can be invoked to account for metaphysically possible change; if time is merely dense, then it appears temporal parts can at most account for possible change where the relevant possible worlds are such that time is either dense (but not continuous) or discrete;

and if time is discrete, then the ability of temporal parts to account for possible change may be limited to those worlds in which time is discrete.²⁶ If such an explanation is required of a temporal parts theory, as Sider and Hawley have suggested, then it would seem that temporal parts have the greatest explanatory power in worlds in which time is as fine-grained as metaphysically possible change, and less in worlds in which time is more coarse-grained.²⁷

One might take these considerations to speak in favour of the thesis, defended in Cameron (2007) and Parsons (2013), that the facts of composition – and by extension, the metaphysics of how things persist – are contingent. Indeed, while Hawley's (2001) is a sustained defence of stage theory, she stopped far short of claiming it as clear victor, and in the very last line of the book can be found an intriguing suggestion: "It is a contingent matter how things persist" (208).

In the following sections (5-7), I'll look at three other potential responses to the problem raised; there are problems with each (including the suggestion of the present section, that the metaphysics of persistence is a contingent matter), but ultimately I'll suggest that taking the metaphysics of persistence to be a contingent matter is preferable, for it is the best way to accommodate the contingency of evidence concerning the structure of time at different worlds.

²⁶ As Newton-Smith notes, certain topologies of time can be embedded in "richer" topologies: for instance, "a discrete, linear structure with first and last elements ... can be embedded in a linear continuous structure without first and last elements in the sense that there is a one-one order-preserving map from the former onto a proper sub-set of the latter. *The converse does not hold.* In addition, any linear dense structure can be embedded in any linear continuous structure". (Newton-Smith 1980, 228-229; emphasis added.) This recalls the potential asymmetry noted above in note 23.

²⁷ And since our reasoning up to this point has relied heavily on the spatial-temporal analogy, analogous remarks would apply to the view that spatial simples are as fine-grained as space. Again, however, Nolan's (2004) suggestion of the possibility of Hypergunk might be troublesome here (see note 20 above): If the actual world is anything less than hypergunky, then actual parts will not be able to account for all metaphysically possible change (assuming hypergunky change is possible) without admitting C*.

5. Cameron's modified version of the Sider-Hawley modal argument.²⁸

I suggested above that the claim that temporal parts are as fine-grained as all possible change is problematic, for (depending on the actual structure of time) it might amount to saying that temporal parts are more fine-grained than time. Cameron suggests a way of modifying the Sider-Hawley modal argument, to motivate the claim that was originally intended: that objects are as (and no more) fine-grained as the regions they actually occupy. Cameron's suggestion is to scale back the explanatory burden of temporal parts. Yes, temporal parts are needed to explain the possibility of change, but not *all* metaphysically possible change; rather, the only merely possible changes that actual temporal parts are required to account for are those that might come about in counterfactual circumstances, changes that *would have* come about had the objects in question been acted upon in some way. It is easy to see how this would get around the problem I raised above: now the demand that temporal parts be actually as fine-grained as 'possible' change is only a demand that temporal parts are as fine grained as time in the closest counterfactual worlds in which things would have changed, had they been acted on in a certain way – but since, for instance, painting an object does not alter the structure of space or time, this does not have the problematic consequence that temporal parts are more fine-grained than actual time.

Why might it be that temporal parts are *only* required to account for possible change that is brought about by acting on objects? I gave two potential motivations for positing temporal parts as fine grained as possible change. The first was based on the dilemma (italicised below) that the temporal parts theorist would have to face if – in an effort to make sense of

²⁸ I presented much of the above at the 7th Annual CMM Graduate conference at Leeds (Oct. 2013), where Ross Cameron was my respondent. In the present section I consider some of Cameron's suggestions.

change in the face of (LL) – she only went so far as to posit temporal parts as fine-grained as *actual* change. Recall Hawley’s homogeneous object:

If the homogeneous object does not actually have spatial parts, we face a problem. The object could have become inhomogeneous midway through its career – someone might have painted stripes on it. Then, presumably, it would have had parts. *But would it have had parts before the painting, because of the later painting, or does painting the object create its parts?* Neither option is attractive, and both can be avoided if we accept that the homogeneous object actually has spatial parts. (Hawley 2001, 49).

If one only has this issue in mind, Cameron’s suggestion seems very natural, for, on the face of it, it seems Hawley’s dilemma can be avoided by positing temporal parts as fine-grained as *actual* time; and if we do this, we have effectively posited temporal parts as fine-grained as possible change – when, that is, it comes to possible change of the kind that Cameron suggested we concern ourselves with.

There are two points I would like to bring forth against taking this line.

Firstly, as I suggested above, the reason why an explanation of merely possible change is wanted is that, while we intuitively accept that any persisting object could undergo change that it doesn’t actually instantiate, the indiscernibility of identicals seems to suggest that, given the identity at different times of some object, that object couldn’t possibly have changed. But the trouble posed by the indiscernibility of identicals here is not restricted to closest counterfactual worlds. Rather, it suggests that, for *any* possible degree of fine-grainedness of change, it is not possible for anything to instantiate that degree of change unless that thing has parts as fine-grained as that degree of change.

And in that case, we have a few options. If it is possible for time to be more fine-grained than it actually is,²⁹ then the need to posit parts as fine-grained as possible change might lead us to posit parts more fine-grained than (actual) time. Another option is to apply the same Sider-Hawley modal argument to the structure of space and time, too, to conclude that, necessarily, space and time are as fine-grained as they could be (this option is suggested by Effingham – I’ll go into more detail about it below). In that case, we would avoid making the claim that objects might actually be more fine-grained than the regions at which they are located, for all actual regions would be as fine-grained as it is possible for them to be. The third option, which I will discuss below in section 7, is that which takes structure of objects to be a guide to possibility.

The other consideration that I wanted to bring forth against the line Cameron suggested is one suggested by Effingham, that there might be mixed worlds, in which the structure of space and time is variable, and probabilistic.

Given the structure of regions is to be empirically investigated, it’s reasonable to take regions to be contingently abundant/sparse³⁰. Given the structure differs interworld there’s no reason to think it can’t differ intraworld, e.g., that it’s possible for spacetime to be continuous everywhere except New York, where it’s dense. Such a possibility is exotic, but not impossible. Further, we could have it that probabilistic laws govern that structure, e.g., that some region has an 80% chance of being abundant, and 20% of being sparse. At that world it’s physically possible for an object exactly occupying a sparse region to have occupied an abundant region, and to undergo changes more finely grained than it could physically undergo given that it actually occupies a sparse region. So it follows, even if we restrict

²⁹ Talking of different possibilities for time itself may require – or at least, be best suited – to the notion that time is a functional concept. See my Introduction, section 2.1, and Bourne (2006, 220-230).

³⁰ “Say that the structure of a region is abundant iff it has as many parts as any metaphysically possible region might have (so if every metaphysically possible region is continuous, dense, or discrete, all and only continuous regions are abundant); otherwise say they are sparse. (Use similar terms for an object being sparse or abundant.)” (Effingham 2013, 190).

ourselves to physically possible changes, that there are exotic worlds where the argument from possible change demands objects have a structure more finely grained than the regions that they occupy.

(Effingham 2013, 191.)

If such a world is genuinely possible, it means that the idea I brought forward above, that we restrict the demand to explain possible change to physically possible change, would not help get around the problem I raised in section 2. But it would also spell trouble for the line that Cameron recommended. Cameron's idea was to posit, for any object, parts as fine-grained as any change that could be brought about by our acting on it in some way (say, painting stripes on it). But in the mixed world Effingham describes, there is no guarantee that the closest such counterfactual world will be one in which the object in question is located in a region that is not more fine-grained. And as such, there is no guarantee that the counterfactual world in which the object is painted will be one in which the change is no more fine-grained than the object is in the world in which it is not painted. In short, if there are mixed worlds of the kind Effingham describes, then, even if we restrict ourselves to having to account for possible change of the sort Cameron suggests we concern ourselves with, positing actual parts that are as fine-grained as *possible* (close, counterfactual) change may nevertheless mean positing actual parts that are *more* fine-grained than the regions that they are actually located at.

Again, this consequence, I take it, is something we should embrace only as a last resort, so let's consider some of the other options for dealing with the problem I raised in section 2. In the next section, we'll see how Effingham suggests we deal with it: namely, that we take possible change to determine the structure of space and time, so that, necessarily, the latter is as fine-grained as the former (for any degree of possible fine-grainedness of change). After that, in section 7, we'll look at another alternative: namely, that we take actual structure to delimit the fine-grainedness of possible change.

6. Effingham: An Alternative Harmony.

In section 4, I suggested that we should consider that how things persist is a contingent matter. As Effingham (2013) remarks in response to Traynor (2013), I am effectively advocating harmonising the investigation of the structures of space and time (or space-time), on the one hand, and objects, on the other. My suggestion was that, since the former investigation is empirical (and so, plausibly, a contingent matter), and since the capacity of temporal parts to explain possible change depends on the outcome of that investigation, we should conclude that how things persist is a contingent matter, also. Where the structure of space and time makes it difficult for temporal parts to explain all intuitively possible change, we may be tempted by rival accounts of persistence.

However, in suggesting this line I was not recommending that we treat the structure of objects as beyond the reach of any armchair reasoning whatsoever. It is just that the success of such armchair reasoning is dependent of the outcome of the empirical investigations concerning the structure of space and time. Since the reasoning leading up to my claim that we should harmonise the two investigations (by taking it that they are *both* amenable to empirical investigation, and hence are *both* contingent matters) depended on the assumption that armchair arguments like that from possible change could nonetheless in principle be informative, it would be self-defeating of me to claim that the *upshot* is that the structure of objects is *wholly* the province of empirical investigation.

As Effingham notes, there is a potential tension here. On the one hand, I have followed Hawley (2001), Newton-Smith (1980), and Forrest (1995), in accepting that the structure of space and time is to be settled empirically. On the other hand, I have allowed that armchair

arguments like that from possible change can help settle the structure of objects. These two strains of thought are in tension because it is hard to see why arguments like that from possible change are applicable to objects, but not to regions. Effingham notes: “If we think that empirical observation, and interpretations of physical theories, will settle what the structure of regions is, it is likely that similar observations and interpretations will settle the structure of objects. To think otherwise would be to maintain an odd disparity!” (2013, 193). Effingham suggests that this would leave little room for arguments like that from possible change, and so suggests an alternative harmony. Rather than claiming that both investigations are empirical, and hence their outcomes contingent, we might harmonise the two investigations in a different direction: *neither* the structure of space and time nor the structure of objects are *settled* by empirical investigation; rather, armchair arguments like that from possible change can settle the structure of space and time *as well as* that of objects. As we’ll see shortly, this would mean that both the structure of objects and the structure of space and time (‘structure’ being used, as before, in the sense of how fine-grained they are) are not contingent: instead, necessarily, regions and objects are as fine-grained as it is possible for them to be.

It is easy to see how the argument from possible change might apply to regions:

Spatial regions can warp and change in the presence of mass, so over time they can change; the same reasons that drive us to thinking that objects possibly changing demands that objects are abundant, will drive us towards thinking spatial regions are abundant. Similarly for spacetime regions: even though they don’t change over time (for how they are, they atemporally are) they are different in possible worlds with differing mass distributions. If it’s metaphysically possible for an object to change in a more finely grained way than the structure of actual regions permit, the same can be said of spacetime regions such that they

could be (atemporally) warped in a more finely grained way than their structure physically permits. (Effingham 2013, 193)

Applying the argument from possible change to regions of space and intervals of time, we arrive at the conclusion that both are abundant: regions of space and intervals of time are as fine-grained as it is metaphysically possible for them to be. Suppose that it is not possible for any region to have more than continuum many sub-regions; then, for any region that *could* have continuum many sub-regions, the argument from possible change tells us that that region *actually* has continuum many sub-regions. In that case, it is not, after all, an open question what the micro-structures of space and time are: they are as fine-grained as it is possible for them to be, and as such it is no problem for the temporal parts theorist to account for possible change. Whatever region and possible world we consider, there are enough times to accommodate as many temporal parts as are needed to account for change of any degree of fine-grainedness. The temporal parts theorist need not posit more temporal parts than there are times; any such modal argument that makes it reasonable to posit a certain number of parts makes it just as reasonable to posit the same number of regions and times at which those parts are located.

This problematizes the claim – which we have so far taken for granted – that the structure of regions is an empirical matter; however, Effingham suggests, this is not as problematic as it may initially seem.

For purpose of example imagine a simplified experiment in a simplified world where O is an example of a smallest object explicitly picked out by physics. Assume that we've conducted an experiment whereby we're justified in thinking that space is continuous if we observe O starting by exactly occupying spatial region R , then exactly occupying a region to its right composed of parts of R and R^* ,

and then coming to rest at R^* . If we observe O ‘jumping’ from R to R^* , and never exactly occupying the region composed just of parts of those regions, we are justified in thinking it’s discrete. It’s experiments like this that should lead us to think the structure of regions can be empirically settled. But the armchair philosopher can mirror what they say about objects. With objects, even if our best theory recognises extended objects as the smallest things, we should nonetheless believe that they have parts—just as science ignores the trout-turkeys and turnip-tenors, science is oblivious to, e.g., the parts of the smallest superstrings. The same can be said of regions, and that no matter what empirical investigations turn up, regions have parts and an abundant structure. (Effingham 2013, 194)

It’s not that the armchair reasoning here makes the empirical investigation redundant; rather, there is a division of labour that ensures that physicists and armchair philosophers do not step on each other’s toes. The empirical investigation speaks to what objects and regions are *like*; the armchair reasoning speaks to what regions *there are*.

Effingham’s alternative harmony has it that possible change is a guide, not only to the structure of objects, but – since regions can change, too – also to the structure of regions.

However, I am sceptical of such an extension of the argument from possible change.

Effingham’s suggestion would result in a very strong claim: that necessarily, all regions and all objects are as fine-grained as it is possible for them to be. This means that, if it is *possible* for objects, time and space to be continuous, then necessarily, they are continuous; so when physicists and philosophers have entertained the thought that time is discrete or merely dense, they have actually been entertaining metaphysical *impossibilities*.

One might have the sense that, if we follow Effingham’s proposal, we are getting things the wrong way around. Considering a similar proposal about the actual structure of time, Newton-Smith says that “rather than adopt a framework that is more or less adequate to any possible world, we want to adopt the framework that is best for the specific world we find

ourselves in” (Newton-Smith 1980, 229). Recall the experiment that Effingham suggested, in which some object is observed to jump from one region to the next, never occupying both regions at the same time. If our armchair reasoning tells us that space and time are continuous, the observed phenomenon is puzzling – why can’t we make the object occupy proper sub-regions of each region R and R^* at the same time? A hypothesis naturally suggests itself: it is because there *is* no region that comprises only proper sub-regions of R and R^* ; R and R^* have no proper sub-regions. For certain phenomena, discrete space/time may provide the best explanation.

Perhaps I should rather say, that, for certain *possible* phenomena, discrete space/time may provide the best explanation, for what we are talking about here are potential experiments, or possible evidence, which may make it reasonable to draw conclusions about the structure of space. I am here making use of a slight modification of the following principle proposed by Phillip Bricker:

We have warranted belief that a structure is logically possible if that structure plays, or has played, an explanatory role in our theorizing about the actual world. (Bricker 1991, 609)

Possible observations or experiments of the kind Effingham offers provide support for the possibility of alternative space/time structures, provided the following principle holds:

We have warranted belief that a structure is logically possible if that structure plays, or has played, *or could play*, an explanatory role in our theorizing about *some possible* world.

A principle along these lines appears to be behind Newton-Smith’s defence of the possibility of discrete time. If we think of discrete regions as being ultimately composed of extended

grains, this comes up against the intuition that such grains must in some sense also be divisible. But taking the latter intuition seriously can only get us as far as the claim that time is dense.

Why then do we tend to regard time as not merely dense but also continuous? ...[T]he answer is quite simply that the best physical theories we have in fact constructed of the physical world require in their mathematical formulation a time parameter that ranges over elements of the real number system... That is, our belief in the continuity of time does not arise from any argument relating to infinite divisibility, it arises from our projecting onto the world the richness that is present in the mathematical system which we have found to date to be essential to the construction of viable physical theories.

When it is seen that this is the source of our conception of time as continuous, we can see what might lead us to adopt a conception of time as discrete. For suppose it should transpire that the best scientific theories of the physical world represented time by a parameter which ranged not over the rationals, nor the reals, but over the integers. ... If this were our device for representing time, we could not represent any interval of time marked by a finite event as having an infinite number of finite parts. It would have a finite number of extensionless parts. Given that we project continuity onto the world because currently our best theories involve representing time by a continuous time variable, if our best theories should in the end turn out to involve representing time by a discrete time variable, we would have at least as much reason to regard time as discrete as we now have for regarding it as continuous.

(Newton-Smith 1980, 118)

Admittedly, this is speculative; but it suggests that if we adopt Effingham's suggestion to apply the modal arguments considered above to regions as well as objects, our armchair metaphysics risks being in conflict with scientific theories in certain possible worlds. When actual well-regarded scientific theories conflict with a certain metaphysical picture, this is often taken to count against the truth of the latter; it is hard to see why the same should not apply to possible scientific theories. If there is a possible world in which the best science of

that world is in conflict with a certain metaphysical claim, we should be open to the idea that that claim is false at that world.

7. Structure as a guide to possibility?

We now turn to the final proposal I want to consider, for one who would argue from possible change and inhomogeneity, on the one hand, to a claim about the actual structure of objects (mereological harmony), on the other. In section 4, we held fast to the claim that temporal parts must be as fine-grained as possible change if they are to explain possible change, and this brought us to the suggestion that how things persist is a contingent matter: A theory of persistence is required to explain why, given the indiscernibility of identicals, non-actual change is nevertheless possible; the ability of temporal parts to perform this task depends on the structure of time; the structure of time is contingent; therefore, it is contingent whether temporal parts can meet their explanatory burden. This of course does not force us to conclude that it is contingent how things persist; but the suggestion is there for those who a) agree that the utility of temporal parts varies depending on the structure of time, yet b) nonetheless find the case for four-dimensionalism compelling. In short, that line was based on the idea that *possible change is a guide to the actual structure of objects*; conjoining that view with the claim that the structure of space and time is contingent – and hence that objects could be more fine-grained than they are – is what led to the problem I identified in section 2.

The final alternative proposal I want to consider is this: that the structure of space and time *imposes limits* on what is metaphysically possible for objects, which is to say that structure (in the limited sense of fine-grainedness) is a guide to possibility. The idea here is that objects can only instantiate change that is as or less fine-grained than the object is actually. So, for instance, if an object is discretely composed, it is not metaphysically possible for it to

instantiate change more fine-grained than its actual structure – it could not instantiate change as fine-grained as the rational number line, nor the real number line. The same goes for objects that mirror the structure of merely dense space and time, and those that mirror the structure of continuous space and time. If an object is (merely) densely-structured, it *could not* instantiate change as fine-grained as the continuum, though it could instantiate change that is as or less fine-grained, e.g. change that could be instantiated by a discretely-composed object. But if an object is as fine-grained as the continuum, then it could instantiate change as fine-grained as either of the structures we have just considered, as well as change as fine-grained as the continuum.

Our starting point for this proposal, then, is that objects cannot instantiate change that is more fine-grained than they actually are. If we nevertheless allow that densely-structured and continuously-structured objects *could* have been *less* fine-grained than they are, this would mean accepting that the logic of metaphysical possibility is weaker than S5. In particular,

$$(B) A \rightarrow \Box \Diamond A$$

is an axiom of S5, but must be rejected on the present proposal. To see this, suppose that time is in fact continuous, and that all actual objects mirror the structure of time by having temporal parts that are as fine-grained as the continuum. Now take some particular one of these objects, O, that, let us suppose, instantiates change as fine-grained as the continuum. On the present proposal, while it is *possible* for O to instantiate change as fine-grained as the continuum (since it actually does, and every world accesses itself), it is not *necessarily* possible for it to instantiate such change. This is because it could have existed in a world in which time is, for instance, discrete; and the present proposal says that an object *cannot*

instantiate change more fine-grained than the region it is actually located at; as such, had O existed in a world of discrete-time, it would not have been possible for it to instantiate change as fine-grained as the continuum. In short: O instantiates change as fine-grained as the continuum, and it is possibly not possible for O to instantiate change as fine-grained as the continuum. That is,

$$(p \ \& \ \diamond \neg \diamond p),$$

where p is 'O instantiates change as fine-grained as the continuum'. This is a counter-instance to (B). On the present proposal, then, (B) does not hold, and so the accessibility relation would be asymmetric.

Short of claiming that it is impossible for objects in worlds of discrete and dense time to exist in worlds where time is more fine-grained than that, we can take it that the rejection of (B) – and hence of S5 as the logic of metaphysical modality – is an essential feature of the present proposal. Whether we must go further, to reject S4 as the logic of modality, depends on what further assumptions we build into the proposal.

So far, our present proposal has it that objects could be less fine-grained than they are.

Consistent with this, we may, or we may not, allow that objects could be *more* fine-grained than they are. We'll now see that, if we do build the latter into the proposal, our modal logic will be weaker than S4; but if we don't, then S4 should be compatible with it.

So, suppose that any object could have been *more* fine-grained than it is. To accommodate this supposition, suppose, for example, that time and space are discrete at the actual world,

@, and take some particular actual object, D, that mirrors the structure of the regions it is located at. Now, as we saw, it is crucial to the present proposal – the proposal that structure imposes limits on the possible fine-grainedness of change and variation – that it is *not* possible for D to instantiate change more fine-grained than discrete-time. For ease of exposition, let's take the case of possible continuous change: since D is *actually* located at a discrete-structured region, and hence D is itself discretely structured, it is not possible for D to instantiate change as fine-grained as the continuum. But we are presently supposing that it is nevertheless possible for D to exist at a world, w_c , of continuous time. This means that, had D existed at w_c , it would have been possible for it to instantiate change as fine-grained as the continuum. All this is just to say that, it is *not* (actually) possible for D to instantiate change more fine-grained than discrete-time, but it is *possibly* possible for it to instantiate such change. That gives us a statement of the form $(\neg\Diamond A \& \Diamond\Diamond A)$. In order to avoid contradiction, then, we must reject the inference from $\Diamond\Diamond A$ to $\Diamond A$, otherwise what we are saying is here is equivalent to a statement of the form $(\neg\Diamond A \& \Diamond A)$. So allowing objects to be possibly more fine-grained than they are on this proposal requires rejecting the transitivity of possibility; since transitivity of possibility (and of necessity) is a key feature of S4, the assumption that objects could be more fine-grained than they are, when coupled with the essential features of the proposal outlined above, would require that the logic of metaphysical modality is weaker than S4.

However, it is also open to us to say instead that, while objects could be *less* fine-grained than they are, they could *not* be *more* fine-grained than they are, and thereby maintain transitivity, and hence keep S4. Consider the claim that it is possible for object O to instantiate change as fine-grained as the continuum. On the present proposal, no possible world that makes this claim true in the actual world will be one in which O is less fine-

grained than the continuum; *only* worlds that are (at least) as fine-grained as the continuum are worlds at which that possibility claim is true. But again, on the present proposal, *all* such worlds are worlds in which the same possibility claim is true (for, on the present proposal, if time is continuous, then objects *can* instantiate change as fine-grained as the continuum); so allowing that an object could have been *less* fine-grained than it actually is poses no problem for the transitivity of possibility on the present proposal. It is only if we allow, as we did in the preceding paragraph, that objects could have been *more* fine-grained than they actually are: in that case, transitivity does not hold.

If that is correct, it is a somewhat curious fact. In particular, it is suggestive of a certain noteworthy relationship between modal asymmetry and transitivity, on the one hand, and asymmetry between certain spatio-temporal structures, on the other (the latter being the asymmetry noted in footnote 26 and footnote 23, based in the fact that there is a one-one order preserving map from certain (sparse) structures to richer structures, where the latter in a sense contain the former, but not vice versa).

We have just been considering the general position that spatio-temporal, and mereological, structure impose limits on what is possible, but we have now considered two variants of that position. In the first, the transitivity of possibility is rejected; in the second, it is not. Which is preferable? While some philosophers reject transitivity of possibility anyway (a well-known rejection is Salmon 1989), I would side with the second variant, which is compatible with S4. On this version of the present view, objects could be less, but not more fine-grained than they are. Why should we *reject* the view that objects could be *more* fine-grained than they are, while taking it that objects *could* nevertheless be less fine-grained than they are?

Recall that we began this chapter with the question of whether we can motivate the view that objects are as fine-grained in their spatial and temporal parts as possible variation and change by way of the claim that parts need to be as fine-grained as possible change and variation in order to *explain* possible change and variation – to explain, that is, why it is that change is possible, in the face of the indiscernibility of identicals, which seems to suggest otherwise. In section 2, I raised a problem for this way of motivating mereological harmony. However, recall that, in order to pose that problem, we had to assume that it is possible for an object to be more fine-grained than it actually is, and hence that it is possible for it to instantiate change more fine-grained than the temporal region that it is actually located at. So by rejecting the claim that things *could* be *more* fine-grained than they are is one way of sidestepping the problem. We can also sidestep the problem if we reject the transitivity of possibility, but to do so we need to claim that, even though it is possible for a discretely-composed object to *exist* at a more fine-grained region, it is *not* possible for it to instantiate *change* as fine-grained as that possible region. Quite simply, it is hard to see why the object cannot instantiate more fine-grained change if we allow that the claim about the object's possible locations is correct; while one might motivate a failure of transitivity in some other way (say, via a kind of modal Sorites series), the present concern of explaining possible change in the face of the indiscernibility of identicals does not seem to provide any additional motivation for the rejection of S4.

However, before ending this section, I would like to flag a potential problem for the present proposal. Recall that one consideration that I brought against Cameron's suggestion is Effingham's thought that there might be mixed worlds, where an object could be variously located at more and less fine-grained regions. This thought is in tension with the present

proposal, which says that an object could not be located at a more fine-grained region. As such, the present proposal either needs to rule out the possibility of these mixed worlds altogether, or else say that anything that seems to pass from a less to a more fine-grained region at such a world has really ceased to exist when it does so, only to be replaced by an uncannily similar object. The latter is certainly an unattractive option for the temporal parts theorist; and the former is *prima facie* implausible. I see no easy way around this issue at present.

8. Extending the argument to make a case against temporal parts.

Since the previous two sections involved the temporal parts theorist in some controversial claims, anyone who lacks such four-dimensionalist sympathies might simply make use of the above discussion to help make the case against a temporal parts ontology. With regard to the reasoning that led to the suggestion that a temporal parts theory is contingently true (if true at all), adding a single additional premise would do the trick.

1. It is incumbent on a theory of persistence to account for metaphysically possible change.
 2. Since temporal parts can only account for metaphysically possible change in worlds where time is as fine-grained as possible change, any temporal parts theory will at best be contingently true, if true.
 3. Whichever theory of persistence is correct, it is necessarily correct.
- C. Therefore, any temporal parts theory is false.

It is perhaps noteworthy that the third premise is not in itself tied to any particular side in the debate over how things persist; and yet, by simply adding the third premise, we have made

the case against temporal parts. Moreover, one might take it that the difficulties raised for the other ways of circumventing the problem of section 2 do not bode well for a temporal parts picture more generally.

One thing to note here is that the problem raised in section 2 applies equally to spatial parts: if it is a problem for a parts-based account of change, it is also a problem for a parts-based account of spatial qualitative variation. Therefore, one cannot simply jump to the conclusion that endurantism is correct, without at the same time adopting some alternative to the parts-based picture of spatial qualitative variation.

It is also worth noting that relativistic physics is often taken to favour perdurantism over endurantism. The problem I have raised for the temporal parts theorist requires that we take contingent evidence seriously. As we have seen, one way around the problem is to note that experiments of the sort that Effingham and Newton-Smith have discussed underdetermine the metaphysics, and hence allow that armchair arguments like that from possible change can teach us about the structure of space and time, as well as the structure of objects. In this way, the temporal parts theorist can retain the necessitarian assumption, and deny premise 2 of the above argument, for behind premise 2 is the claim that the structure of space and time is a contingent matter. In making this move, though, one trades in world-type-specific explanations for generality in metaphysics.

Endurantists can hold on to the claim that the structure of time is contingent, since for them, change is explained at each world by invoking properties like being red-at-t, or adverbial properties such as being red t-ly. Since the endurantist invokes properties that feature in their

very specification the times at which they are, or *would* be, instantiated, they need not worry about whether there will be enough times to accommodate possible change.

But the endurantist may have to make a similar trade, with regard to relativity, for it has been argued that relativity favours a temporal parts picture (Balashov 2010; Calosi unpublished).

If that's right, it remains for the endurantist to note that that there appears to be no *entailment* from the former to the latter, and hence that the physics underdetermines the metaphysics. If that is the state of play, then the endurantist is in a similar position to the temporal parts theorist who adopts Effingham's suggested extension of the argument from possible change to the conclusion that space and time are as fine-grained as it is metaphysically possible for them to be; which is to say, they will be making a trade between metaphysical commitments, and the drive towards necessitism, on the one hand, and respecting the weight of contingent facts of the actual world, on the other.³¹

9. A similar issue for reductionism about time.

In the case of the modal arguments for temporal parts being as fine-grained as possible change, we look to what we consider to be possible for objects and then try to figure out how objects must be if we are to best make sense of these possibilities. The main reason why we were faced with this explanatory burden was that we also assumed the truth of a principle which, crucially, we took to have *modal import*, the indiscernibility of identicals: If x and y are identical, then they *must* share all their (genuine) properties. This suggests that anything

³¹ This recalls the trade off, alluded to in Chapter 1, section 7, that presentists would have to make, were they to maintain that presentism is necessarily true in the face of the evidence in favour of special relativity.

without temporal parts would be incapable of change, and so, if we want to allow that something temporally extended could change, we should posit temporal parts as fine-grained as possible change.

We might thus expect that other principles with modal import introduce similar explanatory burdens, forcing us to draw conclusions about the actual world based on what we take to be merely possible, perhaps thereby going beyond what is suggested by empirical science. One such case, similar in this respect to the one we have been considering, can be found in the debate between substantivalists and relationists about space and time, where versions of relationism have been proposed which tie the existence of regions and times to *possible* locations of objects. (As above, I will focus on time, but similar points should apply to space.)

Roughly, substantivalists think that times are entities in their own right: the existence of times is not dependent in any way on objects or events. Relationists deny this: the existence of times is somehow dependent on events. According to one version of relationism, times are identified with actual events. On this view, where there are no events, there are no times, which is to say, there are no such things as temporal vacua, nor could there be. While Leibniz is perhaps the best-known advocate of this brand of relationism, it is perhaps best motivated by empiricist-epistemological concerns: if times could exist in the absence of any events, then there are possible worlds which are indiscriminable from our own, but are nonetheless distinct from our own merely in the respect that they have more or less times. This would mean that, for all our experience can tell us, our world could be far older than we think: perhaps a billion years has passed since I began typing this sentence. The worry would be that, if the substantivalist is correct, there is nothing in my experience which could rule

that out³², in which case we *actually* never know (indeed, *cannot* know) how much time has passed from one event to another.

Relationists, however, need not deny the possibility of temporal vacua, and may endeavour to formulate their view of times so that this possibility can be accommodated.³³ According to these relationists, the existence of times is not tied to actual events, but, instead, is tied to *possible* events. Newton-Smith offers the following schema for the temporal relationist who would like to accommodate the possibility of temporal vacua.

There is a period of time between the events E_1 and E_2 if and only if relative to these events *it is possible* for some event or events to occur between them. (Newton-Smith 1980, 44)

As Newton-Smith notes, the plausibility of this schema may depend on what kind of possibility is being invoked. If the relevant possibility is metaphysical possibility, similar issues arise as did for the temporal parts theorist. Treat the possibility invoked in the schema as metaphysical, and consider any two events. It is presumably metaphysically possible for there to be another event between those first two, but in that case, according to the schema, there is a period of time in between those two events: the *possibility* of their being separated by a third event suffices for there to *actually* be a time between them, meaning that for every two times, there is a time in between. Hence, treating the relevant possibility as metaphysical possibility, time is not in fact discrete, and nor could it be, insofar as a version of relationism conforming to the schema that Newton-Smith offers is correct.

³² Since a billion years was arbitrary, any interval-length would be indiscriminable from all others.

³³ Newton-Smith (1980, Chapter II) provides arguments for the possibility of spatial and temporal vacua, building on Shoemaker's (1969). We'll take an in-depth look at Shoemaker's argument in Chapter 4.

It may seem that these considerations rule out only that time could be discrete, without deciding between dense or continuous time. However, Le Poidevin shows that relationism, so construed, in fact entails that time is continuous.

Suppose that there exist two times, t and t' , separated by an interval of $2n$ units. It is logically possible that there should exist an event, n units after t , where n ranges over the real numbers. This, on the relationist analysis, entails that there is an actual time existing between t and t' . Again, the argument applies quite generally, and relies on no contingent assumptions, so relationism entails that time is continuous as a matter of necessity. (Le Poidevin 1990, 427)

Since n ranges over the real numbers, the schema tells us not only that between every two times there is a third, but also that there are continuum many times between any two times. Once again, we have a metaphysical claim that implies that time has a certain structure; while it is often thought that the micro-structure of time is an empirical matter, our schema, so interpreted, tells us otherwise. Provided relationism is the correct account of time, time is continuous, in spite of any apparent empirical evidence to the contrary.

This recalls Effingham's suggestion that, where empirical findings underdetermine the structure of time, we should allow armchair arguments to help settle the matter. But think back to Effingham's possible evidence – object O 'jumping' from regions R to R^* , never occupying both at once – and Newton-Smith's speculations, about possible evidence for discrete time. If I am right to suggest that possible phenomena such as this provide evidence for the possibility discrete space/time, then it would seem to provide evidence that relationism is possibly false, and so contingently true at best.³⁴

³⁴ As we saw in section 2.3 of my Introduction, general relativity has been brought to bear on the debate between relationists and substantivalists about space and time. Given that relationism entails continuous time –

Newton-Smith also suggests that this kind of relationism, which takes the relevant modality to be logical or metaphysical, “would lead us into incoherence” (Newton-Smith 1980, 45), much as the argument from possible change can lead us to the conclusion that objects are more fine-grained than the regions at which they are located.

For it is logically possible that some type of event happened, once, twice, three times etc. between E_1 and E_2 . If the logical possibility justified introducing reference to actual periods of time we could say that there was a period, one type E_x event in length between E_1 and E_2 ; and that there was a period, two type E_x events in length, and so on. (Newton-Smith 1980, 45)

The point here seems to be that this version of relationism would require that between every two times, not only are there continuum many *times* as Le Poidevin showed, but there are also intervals of every possible event-length you might consider. Because there *could* be an event of n (extended) units duration between every two events, there actually *is* an n unit interval between every two times. This stretches any interval of time beyond recognition. Hence Newton-Smith recommends that the modality involved in the formulation of relationism be interpreted as physical modality, which is analogous to the move in the case we considered earlier, that involved scaling back the demand for explanation of possible change, to a modality that would guarantee that the parts-based explanation would not require more parts than there are actual times (Newton-Smith 1980, Section II.10).

As in the case of the modal arguments for temporal parts, here, too, we have seen that problems can arise when we hold a principle which tells us that metaphysical possibilities

and assuming that both the structure of time and the laws of nature are contingent matters – this is a further point at which contingent evidence can have a bearing on the relationism/substantialism debate.

impose demands on the ontology of the actual world. In both cases, the divergence between what is possible and what is actual can lead to incoherent ontology. Wanting our theories to cover the whole of metaphysical possibility, we render them unable to account for what we know of the actual world. But it is not just divergence between the possible and actual: divergence between what is metaphysically possible and what is physically possible might also pose constraints on what arguments of this type can be allowed to succeed. This was evident from our earlier discussion of temporal parts. If we allow for the physical laws to impose constraints on the structure of time, but take it that the structure of objects is to be decided by metaphysical reflection, we can end up with an ontology of objects which is incompatible with not only our *actual* ontology of space and time, but also with the sphere of *physically possible* ontologies of space and time.

10. Concluding remarks.

While the Sider-Hawley argument for positing parts as fine-grained as possible change and variation had initial plausibility, we saw in sections 2-3 that it could be highly problematic in the absence of some further substantive claims. In the sections that followed, I explored some ways of retaining the original motivation while getting around the problem. As we saw, each of these ways had some drawbacks. Ultimately, allowing that the structure of regions is the subject of empirical investigation suggests that different empirical results would suggest different structures, in which case, it seems reasonable to conclude that the structure of regions is a contingent matter. This in turn places restraints on the resources available for a parts based explanation of possible change and variation. When the sphere of possibilities vis-à-vis mereological and topological structure is construed as being sufficiently broad, any attempt to ground these possibilities in facts about the actual concrete structure of the world is likely to run into problems unless we pay heed to the constraints given by the

contingent structure of regions. The somewhat analogous case of modal space/time reductionism suggested similar caution is needed there, too.

The next chapter (Chapter 3) provides a critique of Adams's (1979) argument for the possibility of distinct indiscernibles, which is therefore an argument against the (necessity of) the principle of the identity of indiscernibles. As we'll see, Adams's argument is designed to appeal to those who either do not grant the imaginability of distinct indiscernibles, or do not grant that their imaginability is sufficient evidence for their possibility. I'll argue that Adams's argument fails, but, in looking for a way of arguing for a possibility claim that lies beyond demonstration by imagination, he was on the right track.

In Chapter 4, I'll argue that appeal to possible evidence can justify certain possibility claims even if the worlds in question are, in principle, sensorily unimaginable. Thus an appeal to possible evidence can have probative force for a certain possibility claim where imaginative resistance forms part of one's reason for doubting that possibility.

3. The Vanishing of Almost-Indiscernibility in the Case Against PII.

1. Introduction.

In order to continue exploring the role of possible situations and thought experiments in establishing metaphysical truths, this chapter explores an argument which has attracted significant attention in the decades since it was first published: Adams's (1979) for the

possibility of distinct indiscernibles, which would be a counterexample to the necessity of the identity of indiscernibles. Adams makes his argument with Hacking's (1975) paper in mind, in which the latter objected to the use of imagination to establish such counterexamples.

In the face of such an objection, one must do more than simply ask us to imagine the world in question. Accordingly, Adams provides an argument from a possibility that all can grant: a world of *almost* indiscernible objects. However, Adams's argument is incomplete, and, examining three recent construals of that argument, I'll argue that each fails.

I'll close the chapter by touching on an alternative route to the claim that there could be distinct indiscernibles, being an argument from quantum physics, which equally does not hinge on the imaginability of distinct indiscernibles. The argument from quantum physics seems to presuppose that the best theories of actual contingent phenomena can provide a guide to possibility. This will bring us a short distance to the central idea of Chapter 4: that the best theories of merely possible phenomena should themselves be taken to be *possibly* true.

2. Almost Black's spheres.

In the following famous case, Max Black posed what many take to be a serious challenge to the principle of identity of indiscernibles (PII).

Isn't it logically possible that the universe should have contained nothing but two exactly similar spheres? We might suppose that each was made of chemically pure iron, had a diameter of one mile, that they had the same temperature, colour, and so on, and that nothing else existed. Then every quality

and relational characteristic of the one would also be a property of the other. Now if what I am describing is logically possible, it is not impossible for two things to have all their properties in common. This seems to me to *refute* the Principle. (Black 1954, 156)

But Black's spheres have not been universally accepted as a counterexample to PII (see, for example, Della Rocca 2005). In the face of Black's case, the advocate of PII has two main options. First, one might claim that Black's spheres are unimaginable/inconceivable (let's just speak just of imaginability hereafter). In that case, Black's description would fail to provide a counterexample to PII, even if it were granted that imaginability entails possibility. Second, one might think that Black's scenario is impossible, whether it is imaginable or not. In that case, even if Black's scenario is imaginable (or at least, *prima facie* imaginable), this would be an instance where imagination leads us astray. This response might be motivated by, for instance, an overriding commitment to bundle theory, or to the principle of sufficient reason, or perhaps both.

Faced with either of these strategies of resistance, can anything more be said in favour of the possibility of Black's scenario? Enter Adams's 1979 argument from almost-indiscernibility:

We may just have an intuition that there could be distinct, though indiscernible, globes in these circumstances. But there may also be an argument for this view—which will depend in turn on other intuitions, like all arguments in these matters. The argument might rest on an intuition that the possibility of there being two objects in a given spatiotemporal relation to each other is not affected by any slight changes in such features as the color or chemical composition of one or both objects. If we accept that intuition, we can infer the possibility of indiscernible twins from the uncontroversial possibility of *almost* indiscernible twins. No one doubts that there could be a universe like the universe of our example in other respects, if one of the two globes had a small chemical impurity that the other

lacked. Surely, we may think, the absence of the impurity would not make such a universe impossible.

(Adams 1979, 17)

Adams's case of almost indiscernible spheres is meant to shore up the case against PII. But how exactly does *almost* indiscernibility function in Adams's case, to help convince someone who is not already convinced by Black's scenario?

Suppose it is the imaginability of Black's case that is at issue (the first potential strategy of resistance mentioned above). In that case, Adams's case of almost-indiscernibility can provide a little rhetorical nudge towards accepting it: "Are you sure you cannot imagine Black's scenario? Just think of Adams's spheres, only without the small chemical impurity – there you have it!" In this instance, *almost*-indiscernibility plays an important part: in virtue of the apparent closeness of Adams's case to Black's case, it would seem to require no great leap of the imagination to move from the imaginability of the former to that of the latter.

But now suppose instead that it is the move from imaginability to possibility that is at issue (the second potential point of resistance mentioned above). Then Adams's suggestion of an *argument* for the *possibility* (as opposed to merely the imaginability) of symmetrical universes comes into focus.

Adams originally invoked his almost-indiscernibles as part of a response to Hacking's (1975) critique of cases like Black's – cases which, Hacking argued, fail to say anything non-question-begging against PII. In the context of Hacking's critique, imagination is not enough. Ultimately, Hacking says, "it is vain to contemplate possible spatiotemporal worlds to refute or establish the identity of indiscernibles" (ibid. 249), for "[n]o matter how vivid your imagination, it remains a question how correctly to describe the content of your

imagination” (ibid. 251). To describe Black’s spheres in non-question-begging terms, one must do so without writing their distinctness into the description. If Adams’s argument is to go beyond Black’s, it must provide us with grounds for thinking that distinct indiscernibles are possible, without simply asking us to imagine some things that are qualitatively identical, yet numerically distinct.

Hence what’s needed is an argument to say that there could be distinct indiscernibles, one that addresses those who object to Black’s scenario whether it is imaginable or not.

However, there seems to be something of a gap in Adams’s own presentation of the argument. “*Surely*”, Adams said, “the absence of the impurity would not make such a universe impossible” (Adams 1979, 17; my italics). Following Cross (2009), we may call this ‘the surely-gap’.

3. Attempting to bridge the ‘surely-gap’ with causal independence.

Cross (2009) attempts to fill the surely-gap by adding – or perhaps, making explicit – the premise that the almost-indiscernible spheres enjoy a degree of causal independence, so that the following counterfactual holds at Adams’s world: *had the slightly impure sphere existed without its impurity, the other, pure sphere would be unaffected*. Assuming the necessity of identity, any world that satisfies the antecedent of the counterfactual would be one in which the spheres are indiscernible, yet distinct.³⁵

³⁵ Suppose there is a world that satisfies the antecedent of the conditional but at which the spheres are identical. The necessity of identity tells us that they are therefore identical in the world in which the spheres are discernible. But the principle of the indiscernibility of identicals (much less controversial than the identity of indiscernibles) tells us that this is absurd. Therefore, there is no world that satisfies the antecedent of the counterfactual but at which the spheres are identical.

Since the aim is to bridge the surely-gap, a little more needs to be said here about why we should think the relevant counterfactual holds at Adams's world (for no progress would be made if we simply asserted the counterfactual with "surely" prefixed), and Cross obliges: we are free to conceive of Adams's world as being such that the counterfactual in question is *nomologically underwritten*.

Which laws might do the underwriting? Perhaps those of the actual world? After all, Cross notes, "[o]ne of the features of the laws of nature in the actual world is that they permit a great deal of *causal independence*": if my neighbour's car was involved in a traffic accident, this does nothing to alter the status of my car which, locked away in the garage, remains dent-free; that is, the fact that my own car rests intact in my garage seems to be causally independent of the denting of my neighbour's car (2009, 280). Similarly, then, might the actual laws of nature underwrite the counterfactual, *Had the impure sphere existed without its impurity, the other sphere would remain as it is?*

The trouble is, as Cross admits, it is not at all clear that the *actual* laws of nature could sustain scenarios like those described by Black and Adams:

[H]ow did the almost indiscernible spheres in Adams's world come into existence if Adams's world is governed by the actual world's laws of nature? If the two spheres are just two miles apart and obey the actual world's laws of nature, will gravity not soon pull them into a collision (perhaps giving them distinguishing dents of their own)? And can we be sure how alike the two spheres would be if the blemished one had no impurity? (ibid.)

Thus it is problematic, at the very least, to invoke the actual laws to make good Adams's possibility claim.³⁶

Nonetheless, Cross suggests, we can underwrite the relevant counterfactual by stipulating other laws.

Indeed, that world need have only one law of nature, namely, that whenever two almost indiscernible spheres are as described [by Adams], the relevant causal independence obtains. Thus, the laws of nature in the world of the almost indiscernible spheres may be somewhat exotic after all, but, exotic or not, it is completely plausible that a world with such laws is metaphysically possible. (ibid., 282)

As we saw above, if Adams's case is aimed at nudging us into granting the imaginability of Black's spheres, then almost-indiscernibility has a clear role to play. But if instead the main aim of Adams's case of almost-indiscernibles is to provide an *argument* for the *possibility* of Black's spheres, then what does almost-indiscernibility have to do with it? Suppose that Adams had presented his case a little differently:

(Sphere-Cube): No one doubts there could be a universe just like the universe in Black's example, but where some asymmetry obtains. So think of two objects, *almost* exactly alike, except that, while one is a sphere with a diameter of one mile, the other is a cube. *Surely*, one might think, had the cube been a sphere made of chemically pure iron, one mile in diameter, then Black's scenario would obtain.

³⁶ It would be particularly doubtful that our laws obtain in any world conforming to Adams's and Black's stipulations if one held a Humean/best system account of laws (given that relatively little goes on in those worlds).

If we accept Cross's suggestion for bridging the surely-gap in Adams's argument, the argument from the sphere-cube possibility should be just as good. After all, if we can posit exotic laws to underwrite the causal independence of two almost-indiscernible *spheres*, it is hard to see why we cannot also posit laws that would similarly underwrite the causal independence of the two objects in the modified case of the sphere and the cube – in which case, the degree of similarity between the two spheres in Adams's case does little for the argument.

To see this, consider that my substitution of a cube for Adams's slightly impure sphere is arbitrary: any other shape would suffice for my purpose – perhaps the shape of a train set, for instance, or a duplicate sphere embedded within a giant statue; for, given the power to stipulate exotic laws³⁷ to underwrite the causal independence of Adams's spheres, we are no less entitled to do the same *no matter what the difference in shape or size between the two objects*.³⁸ This is problematic for Cross's bridging because, had Adams originally presented the sphere-cube case, I would wager, few would have taken it as providing any additional support for the possibility of distinct indiscernibles; fewer still would have found any support for the possibility of Black's case had the slightly impure sphere been a train set; fewer still, had the slightly impure sphere been a galaxy (or for that matter, a pure iron sphere embedded in a galaxy), and so on.

³⁷ Have I missed a trick here? One might think of a simple relativistic – non-exotic – world in which two things begin their existence far enough apart, and (co-)exist briefly enough, that causation between them is and remains (physically) impossible. However, this still leaves almost-indiscernibility irrelevant – the two systems could be as asymmetric as you like, as long as they are far enough away from each other and (co-)exist briefly enough.

³⁸ Of course this may be challenged – perhaps there is some principled constraint on the extent to which two possible objects can differ while enjoying a sufficient degree of causal independence from each other; but this would require substantial argument. Note also that it would not be enough to challenge the claim that a sphere could have been an object of any arbitrary shape, for there are enough other dimensions of dissimilarity – size, colour, markings and dents – that the asymmetry could nonetheless be arbitrarily great.

Since the plausibility of the argument varies depending on the degree of asymmetry given in the case, what is doing the convincing here is the degree of similarity between the worlds in question. Yet the causal independence version of Adams's argument can be made regardless of the degree of asymmetry given in the case; thus what does the convincing is not operative in the argument so construed.

Cross's proposal bridges the gap in Adams's argument at the cost of making almost-indiscernibility – the most salient feature of Adams's case – irrelevant to the argument. As such, Adams's case is rendered dialectically ineffective for the purpose of convincing anyone apt to doubt the possibility of Black's case. If indeed the degree of qualitative similarity between the objects in Adams's case is irrelevant, then the argument amounts to nothing more than an insistence that there could be perfectly symmetrical universes of more than one object – precisely what was at issue, and in need of demonstration.

4. Adams's as a Subtraction Argument?

In lieu of causal independence, one might look to metaphysical independence, focusing on the *inessentiality* of the impurity to either sphere in Adams's scenario. Accordingly, we might think of Adams's as a subtraction argument, modelled on Baldwin's (1996) argument for the possibility of an "empty world" (that is, a world wholly lacking in concreta).

Baldwin's original argument involves a process of subtraction. Roughly: there could be a world, w , with a finite domain of concrete objects, each of which exists only contingently. The non-existence of any one of these objects would not necessitate the existence of anything

else. Hence, starting from w , we can ‘subtract’³⁹ any one of the objects that exists at that world, to arrive at world w^* , which contains everything in w except the subtracted object, and nothing more. Repeating this process of subtraction at w^* , and again at the resulting world, and so on, we will eventually arrive at a world at which no concreta exist.

One might think that a similar process of subtraction underpins Adams’s argument.⁴⁰ The impurity in the one sphere is essential to neither sphere, and the absence of that impurity would not necessitate the existence of anything else. Hence, by subtracting that impurity, we arrive at a world at which there are two duplicate spheres, and nothing else – a world, that is, at which there are distinct indiscernibles, a counterexample to PII.

How might almost-indiscernibility play a role in the argument thus construed? Perhaps the present emphasis on essential/non-essential features provides the key here. The *smallness* of the impurity might be considered important because, in order for the argument to work, we need to be sure that once the asymmetry-making feature of the scenario is subtracted, the resulting world is one in which the *very same* sphere exists. If the impurity of the one sphere is too great – suppose it is 49% copper – we cannot be sure that subtracting the impurity would leave us with a numerically identical object. If the modal subtraction does not preserve the identities of the objects, then the necessity of identity cannot be invoked to show

³⁹ “Baldwin’s phrase, ‘this procedure of subtraction’, is undoubtedly meant to be a metaphor. If it is taken literally, it suggests a physical operation on a possible world to create another possible world by taking away some of the objects in the original possible world. Because possible worlds are not the kind of thing one can create by a physical operation, such a literal interpretation of ‘this procedure of subtraction’ is thus inadvisable and was surely never intended. Instead of taking subtraction in such a literalist sense, we should try to understand subtraction in terms of a relation between the objects which exist at different worlds. In this way, we can make sense of subtracted worlds, that is, worlds whose domains stand in a certain relation to each other, a relation suggested by Baldwin’s subtraction metaphor.” (Efrid and Stoneham 2005, 305). Similar caution must be exercised when thinking of the absence of the impurity from Adams’s case: it is not that at one time, it is there, and a later time, it is not – for in that case the two spheres would remain discernible even once the impurity has vanished.

⁴⁰ Thanks to Aaron Cotnoir for pushing this idea. I am also indebted to David Ingram’s paper (ms.), in which he argues in favour of this way of construing Adams’s argument.

that there must be *two* spheres in the world that results (see footnote 35). Hence, the proponent continues, almost indiscernibility plays a crucial role in the argument so construed.

However, there is a difficulty for the subtraction-construal of Adams's argument. A crucial premise in that argument is that *the subtraction of the impurity* (or perhaps rather: the *feature of impurity*) *would not necessitate the existence of anything else*. But if that subtraction is meant to leave us with a sphere of pure iron, of the same size and shape as the other, then a certain amount of *additional* iron would need to result from that subtraction (not to mention the addition of the feature of *purity*). This additional portion of iron is clearly something other than the impurity (whether the impurity is thought of here as some portion of non-iron matter, or as an abstract feature). It is also something other than the sphere (to think otherwise is to commit the "mereological fallacy" of mistaking a thing for some proper part of it (Bennett and Hacker 2003, 68ff.)). Hence, in order for the subtraction to leave us with two qualitatively identical spheres, the premise in question must be false, for the subtraction needs to involve an *addition* of one sort or another if it is to leave us with Black's scenario.

Moreover, it is not clear that the subtraction would result in a *pure* iron sphere. Getting to the closest world that results from the subtraction of the impurity might, for example, involve adding some other element to the sphere, so as to leave the sphere impure in virtue of some other small impurity. And it is not clear that the subtraction must result in a pure iron *sphere*, at least not one of the same size as the other. For perhaps the subtraction of the impurity does not involve adding anything besides the feature of *purity* – in which case we would end up with a further asymmetry in the shape of the objects: one being a sphere of pure iron, the other being *almost* spherical – that is, spherical but for an impurity-shaped notch. Or perhaps the subtraction of the impurity does leave us with a pure iron sphere, but one which is smaller

than the other sphere, made of slightly less matter than its co-inhabitant (again, leaving us with an asymmetrical scenario).

It may therefore seem fortunate, for the subtraction-construal of Adams's argument, that Adams need not have invoked the *kind* of asymmetry-maker that he did. The version of PII under discussion covers extrinsic as well as intrinsic properties⁴¹. Adams appealed to an impurity not because of the kind of entity it is (an aspect or feature), but because it makes for an asymmetry in what otherwise would be a description of indiscernible spheres; in this way, Adams is starting from a world that all parties in the debate can agree is possible. The impurity as such – being a *feature*, or *aspect* of one of the spheres – could just as well have been replaced by a small *object* adjoined to one of the spheres and not the other. Unlike the impurity, it is easy to see how such an additional object could be subtracted to leave us with pure iron spheres of the same size, without requiring the existence of something else.

However, this brings us back to the same problem that afflicted Cross's proposal. True, in order to block Adams's move from almost-indiscernibles to distinct *indiscernibles*, it seems the advocate of PII is required to say something strange, namely that *the absence of some small qualitative asymmetry in the scenario would not have made for Black's world of two indiscernible spheres*. But why is it any more difficult for the advocate of PII to deny the possibility of Black's spheres in light of the scenario of *small* asymmetry than it is in a case of two duplicate spheres embroiled in a scenario where some *vast* asymmetry obtains?

⁴¹ "Let us therefore here reserve the title 'Identity of Indiscernibles' for the doctrine that any two distinct individuals must differ in some suchness, either relational or non-relational" (Adams 1979, 11).

Now this could go both ways, depending on what is at issue. Recall (section 2 above) that Adams's case of almost-indiscernibles could be doing either of two things: it can either be used to convince one of the imaginability of scenarios like Black's spheres, or it can be used to argue for the possibility of such scenarios *whether imaginable or not*. There may be something capricious about denying the *imaginability* of distinct indiscernibles, which the case of almost-indiscernibles brings out in a way that a case of vast asymmetry does not (though, in the next Chapter, I'll suggest principled grounds for rejecting the imaginability of distinct indiscernibles). But now suppose again that Adams's argument is addressing those who deny the possibility of Black's spheres whether they are imaginable or not. Whatever reason these theorists have (commitment to bundle theory, say) should remain unaffected by a consideration of almost-indiscernibles unless almost-indiscernibility plays some role in an argument against PII, one that does not hinge on the imaginability of a counterexample to the principle.

And almost-indiscernibility does *not* play any such role in the subtraction-based construal of Adams's argument. In the modal-subtraction version of the argument, the possibility of satisfying two conditions is meant to allow us to make the move from the scenario of distinct *discernibles* to the possibility of distinct *indiscernibles*: we need to grant a) that the thing that makes for the asymmetry in our case is essential to neither of the objects, and b) that the absence of that thing would not necessitate the existence of anything else. But these conditions could be satisfied no matter how great the degree of asymmetry in our example: we could just as well argue for the possibility of distinct indiscernible spheres via a world in which there are two duplicate spheres *and a plethora of asymmetry-making things besides* – stars, planets etc. In such a case, the plethora of things besides the duplicate spheres, just as the impurity in Adams's case, is essential to neither sphere, and its non-existence would not

necessitate the existence of anything else. Thus the problem here is essentially the same as the one I raised for Cross's proposal: once again, it seems that almost-indiscernibility has nothing to do with the argument. And yet, an argument from the possibility of duplicates in a vastly asymmetrical universe, to the possibility of distinct indiscernibles – such as Black's spheres – would be wholly unconvincing. Given that the degree of asymmetry utilised, as well as what makes for the asymmetry, is an arbitrary matter as far as the subtraction argument is concerned, that argument reduced – without loss of informativeness – to the simple statement that two duplicates could have existed in the absence of anything else, and hence says nothing more than Black's original argument did.

* * *

I have argued that the first two construals of Adams's argument render almost-indiscernibility irrelevant. The situation recalls a point made by Sorensen: "Many hypothetical claims owe their plausibility to slanted storytelling" (Sorensen 1992, 29). The point can be nicely illustrated by looking at how re-description of body-swap cases can elicit markedly different intuitions about what identity consists in.

The typical scenario features two people, A and B, who enter a machine that swaps their skills, memories and character. Now suppose that A knows that he will enter the machine. He is given a choice as to which person will receive a reward (say \$100,000) and which will be penalized (say, by torture). It seems in A's self-interest to assign the penalty to the person emerging with his old body and the reward to the person emerging with B's old body. This shows that ... (just as the dominant view says) our psychology is more important to personal identity than physical features. Williams objects by redescribing the situation as one in which amnesia is induced, skills erased, and character is blanked out, so that new "memories," skills, and character can be written into the erased person. Under this

description, it seems that A's self-interest is served by rewarding the person emerging with the A-body and penalizing the B-body person. (Sorensen 1992, 29; Williams's discussion is in his 1973, 46-63)

The re-description leaves the case essentially intact, and yet the superficial differences between the descriptions make a marked difference to the conclusion drawn; hence, what does the convincing is only superficially of relevance, having to do with how the case is described, not with the case itself. In the same way, the two construals of Adams's argument that we have just considered can be re-written so as to omit the appeal to almost-indiscernibility, leaving each version of the argument essentially intact. Yet once so re-written, the force of Adams's case is totally lost; hence, what does the convincing on either of the first two construals of the argument is merely superficial, having to do with how the argument is presented, not with the argument itself.

In section 6, I'll make a positive suggestion towards explaining away the intuition supporting Adams's case. But first, let's consider a third and final construal of the argument, according to which it is an argument from modal continuity.

5. An Argument from Continuity?

Rodriguez-Pereyra presents the continuity version of the argument as follows.

But Black's world is possible. For there is a possible world with two almost indiscernible spheres. That is, there is a possible world like Black's except that the spheres differ infinitesimally in temperature.... That world contains two particular spheres, *a* and *b*. But if *a* has a temperature *T* and a different particular *b* of the same kind as *a* has a temperature *T** infinitesimally different from *T*, then it is

possible for a to have T^* . Thus if the world with the almost indiscernible spheres is possible, so is another world in which the spheres are completely indiscernible. (Rodriguez-Pereyra 2004, 74)⁴²

As Cross notes, “So interpreted, the argument counts as a continuity argument because it relies on the assumption that *if there are possible worlds that exhibit a range of degrees of a certain characteristic, and if the degrees in that range approach some limit, then there must be a possible world at which the limit is realized*” (Cross 2009, 285-286; emphasis added). Hereafter, I’ll refer to the italicised assumption as “Rodriguez-Pereyra’s principle”. Cross prefers his own construal, since “it in no way turns on an assumption that the space of possible worlds must be continuous in the way described [by Rodriguez-Pereyra]” (ibid. 286). However, on the face of it, the continuity argument is not subject to the criticism I made of the other versions of the argument we have considered: unlike Cross’s proposal, and the proposal according to which the argument is one from modal subtraction, it seems that *almost-indiscernibility is* relevant in the continuity argument, functioning as *representative* of any difference which, however small, could be smaller. By invoking the continuity principle, we are supposed to be able to infer the possibility of Black’s world from the existence of a *series* of possible worlds displaying ever-closer symmetry.

However, Rodriguez-Pereyra’s principle does not enable us to infer the possibility of Black’s world from the existence of such a series. To see this, we must first take care not to confuse the principle behind Rodriguez-Pereyra’s argument with a similar continuity principle endorsed by Leibniz, and discussed by Graham Priest,

⁴² Note that Rodriguez-Pereyra, p.74, footnote 6, says that there are “important differences” between his argument here and Adams’s – though he doesn’t say what these are.

[which] is intended to apply to all limiting processes—not just arithmetic, but geometric, physical, temporal, and so on. In virtue of this, we might state the principle thus: given any limiting process, whatever holds up to the limit holds at the limit; or, as L'Huilier, who, like most eighteenth-century mathematicians, endorsed the principle, put it: *if a variable quantity at all stages enjoys a certain property, its limit will enjoy the same property*. (Priest 2006, 166; emphasis added)

The important difference between the Leibnizian principle and the modal principle that Rodriguez-Pereyra is appealing to is that the former tells us that the property obtaining in the *run up* to the limit obtains *at* the limit. Rodriguez-Pereyra's principle, on the other hand, does not demand this: we are first to determine what the limit case is; only then does the principle tell us what is possible.

This disanalogy between the two principles is in some respects fortunate for Rodriguez-Pereyra. First, the Leibnizian continuity principle wouldn't give Rodriguez-Pereyra the conclusion he is after. The Leibnizian principle tells us that what obtains in the run up to the limit obtains at the limit; in the series in question, there would be some asymmetry at each of the stages, and so, according to the Leibnizian principle, whichever property makes for the asymmetry in the approach to the limit would also apply at the limit. Hence, the Leibnizian continuity principle would give us a limit world at which some asymmetry obtains – not one that contains distinct indiscernibles.

Second, as Priest notes, in this most general and unqualified form, the Leibnizian continuity principle would allow us to prove too much: “using it carelessly, one could prove all sorts of undesirable things, such as that every real number is rational (since it is the limit of a sequence of rationals), that the limit of every sequence of continuous functions is continuous,

and so on” (ibid.).⁴³ To take a modal example, consider a series of worlds, each containing just a sphere and a ruler⁴⁴. The series begins with a world where the diameter of the sphere is half the length of the ruler; the next world in the series is one in which the diameter of the sphere is one quarter of the ruler; the next, one eighth – and so on. If we can assume that what applies in the series approaching the limit applies *at* the limit, then we are forced to say that the limit case is a world containing a point-sized sphere, i.e. an impossible object.

Rodriguez-Pereyra’s principle doesn’t lead us to these undesirable conclusions, for it doesn’t tell us that what applies in the series approaching the limit applies at the limit. It is neutral on the question of what the limit case is for any given series. Only once we determine what the limit case of a given series is does Rodriguez-Pereyra’s principle tell us what is possible.

This feature of Rodriguez-Pereyra’s principle, however, brings with it a problem of its own, for it allows us to ask the following question: Why should we think that the limit of the series of worlds Rodriguez-Pereyra invokes is a world containing *two* objects of the same temperature, rather than one? It may seem intuitively obvious: there were two in the series up to the limit, so there are two at the end. But consider now a different series of worlds, starting with a world containing two black cubes of the same size, some distance from one another.⁴⁵ So as not to assume the falsity of PII, let’s add a white china teapot which floats around in such a way as to break the symmetry. For the next world in the series, the two cubes overlap, sharing a part half of the area of each, which we can stipulate to be white. (To maintain that there is continuous variation across the series of worlds, suppose also that the

⁴³ Indeed, Priest himself goes on to use a similar principle to argue that change is contradictory. Staunch advocates of the law of non-contradiction might interpret this as a reductio.

⁴⁴ The ruler is there to appease the Leibnizian (relationist) about space, who would not grant that there could be two worlds, each containing nothing but a single sphere except where the sphere in one world is larger than the sphere in the other.

⁴⁵ Thanks to Katherine Hawley and Bruno Jacinto, who both independently suggested a case along these lines.

china teapot follows the same trajectory at all worlds of this series.) Across the series of worlds, the area of overlap gets progressively greater, with the white area taking up half, then $3/4$, then $7/8$, then $15/16$, then $31/32$, etc. of the area of each cube. Now, what shall we say is at the limit? It is clear that any object in the limit world will be white, and that the greatest non-teapot-shaped region occupied by an object at the limit world will be a cube-shaped one. But is there a *single* largest cube in that world, or are there two, wholly overlapping cubes? Whichever is the limit case, Rodriguez-Pereyra's principle tells us that it is possible. However, advocates of any of the following principles will naturally deny the possibility of distinct but wholly overlapping cubes:

Composition as identity: objects are identical to the sum of their parts (Lewis 1991).

No coincidence under a sortal: no two things of the same kind can occupy (precisely) the same place at the same time. (Locke 1975, Book II, Chap. 27, §1)

No co-located duplicates: "there cannot be distinct indiscernible things that occupy precisely the same location at all the same times and have all the same parts" (Della Rocca 2005, 486).

So, if it is somehow settled that the limit case is indeed one in which there are two cubes, reasons for advocating any of these principles will come into conflict with Rodriguez-Pereyra's principle *qua* way of generating possibilities.

The situation is similar in the case of PII. If it is somehow settled that the limit case in Rodriguez-Pereyra's series is indeed one in which there are distinct indiscernibles, reasons

for advocating PII will come into conflict with Rodriguez-Pereyra's principle. But how shall we settle the question of what lies at the limit? It cannot be simply by way of the idea that what applies at the approach applies at the limit: as we have seen, that's not what Rodriguez-Pereyra's principle says, and had it been, the argument from the possibility of almost-indiscernibles would not work. One wants to say that, in the case of progressively overlapping cubes, the plausibility of some principle among the above gives us reason to think that the limit case is one of just a single largest cube – but this suggests an analogous move for the advocate of PII. If Rodriguez-Pereyra's principle cannot be used to refute composition as identity, then nor can it be used to refute PII.

It might be thought that there is a crucial difference between the series of increasingly similar spheres and the series of increasingly overlapping cubes: in the former series, the spheres remain at distinct locations at all worlds in the series approaching the limit, whereas in the latter series, the locations of the cubes overlap in tandem with the overlap of the objects. However, we must remember here that, in the debate surrounding PII, absolute space is not to be taken for granted. Whatever reason we have for thinking that the limit case is one of distinct indiscernibles, it cannot be that one is *here*, the other *there* (Hacking 1975, 251-2; Hawley 2006, 302). After all, there would be no need for an Adams-style argument if distinctness of locations could be invoked to ground the distinctness of the objects. Nor can the reason be that the spheres in the series of worlds running up to (not including) the limit are at distinct locations, for we have been given no reason for thinking that distinctness of locations is something that the limit case inherits from the rest of the series.

It thus seems that Cross was right to be suspicious of the modal continuity-based construal of Adams's argument. Since Rodriguez-Pereyra's principle does not offer us any way for

deciding what the limit case is, invoking that principle would simply relocate the debate, from one about which possible world corresponds to a non-question-begging description of Black's case, to one about what we ought to take to be limit of the series of possible worlds varying in degree of asymmetry.

6. Explaining away the intuition: Nature doesn't make jumps, but modal space might.

As a general rule, whenever one contradicts a common intuition, one thereby takes on a burden: to say something as to the origin of the intuition, consistent with its falsity in spite of its prima facie appeal. I have examined three recent construals of Adams's argument and found each wanting. This, I think, constitutes a defeasible argument to the effect that the surely-gap cannot be effectively bridged. But what of the intuition upon which Adams suggested his argument might rest? Recall, this was "an intuition that the possibility of there being two objects in a given spatiotemporal relation to each other is not affected by any slight changes in such features as the color or chemical composition of one or both objects" (Adams 1979, 17). Indeed this is a natural intuition to have, and there can be little doubt that it has prima facie plausibility. Ultimately, however, this naturalness can be employed to explain away the feeling many of us have, that Adams's case somehow supports the possibility of Black's spheres. Adams's argument trades on the feeling we have that the degree of similarity between the two spheres is somehow relevant in the move to distinct indiscernibles. My suggestion is that the relevance of similarity is drawn from our knowledge of what goes on in the actual world – and misapplied to the modal realm at large.

To illustrate, compare an uncannily similar thought experiment, offered by Leibniz, to refute the actuality of Cartesian laws of collision.⁴⁶ According to the Cartesian laws, when a moving object collides with a larger, stationary object, the smaller one will rebound with the same speed with which it approached; but when a moving object collides with a *smaller* stationary object, both will continue in the direction of the original impetus in such a way that the total quantity of motion is conserved. Leibniz objected to this theory with a thought experiment. First, think of two such objects (say, two spheres), where the moving one is imperceptibly smaller than the stationary one; the Cartesian laws tell us that it will ricochet off the stationary one with the same speed with which it hit. But now, think of the same objects, where the roles are reversed: the moving one is imperceptibly larger than the stationary one which it hits. In this case, the Cartesian laws tell us that that both objects will move off in a way that conserves the total quantity of motion. But this seems absurd: when applying the Cartesian laws of collision to Leibniz's two cases, we find that an *imperceptible* difference in size makes for *strikingly* different outcomes. In this way, Leibniz brings out the fact that the Cartesian laws of collision violate an intuitive judgement: *nature doesn't make jumps*.

It seems that Adams's argument rests on a similar intuition: surely, a tiny difference – the lack of a small chemical impurity – would not affect the scenario to such an extent as would be required to sustain PII (whether by necessitating the non-existence of one of the spheres, the identity of the two spheres, or the presence of some other asymmetry-maker). However, in Leibniz's case against the Cartesian laws of collision, the intuition appealed to has its basis in what we know of the actual world, and hence it is easy to see how that intuition would be

⁴⁶ The following description is drawn from Sorensen 1992, 10.

relevant to the question of which laws *actually* obtain; but, as we saw in section 3, the natural laws of the *actual* world cannot be appealed to in order to sustain the relevant counterfactual in Adams's case. Given the exotic nature of Adams's and Black's worlds, we cannot straightforwardly transpose the relevance of similarity from the actual world onto Adams's case; and, hence, we should be cautious of letting intuitions garnered from our knowledge of what goes on at the actual world influence our judgements about what would have been the case had Adams's impure sphere lacked its impurity – small though that impurity may be.

Given any of the discussed construals of Adams's argument, it is reasonable for the advocate of PII to hold their ground. I have argued that essentially the same arguments as Cross's and the subtraction argument can be cast without invoking almost-indiscernibility, meaning that almost-indiscernibility is not relevant to these arguments. Yet almost-indiscernibility must be somehow relevant if it is to play more than a rhetorical role in an argument from the possibility of almost-indiscernibles. Rodriguez-Pereyra's proposal provided the most promising construal on that score: on that construal, Adams's small impurity can be seen as representative of some difference that, no matter how small, could have been smaller, hence illustrating at a stroke that there is an infinite series of possible worlds varying in degree of asymmetry. But we saw that this does not tell us that the possible world at the *limit* of that series is one in which there are two indiscernible spheres. Hence, even if modal continuity tells us that the limit is possible, it does not tell us that it is possible for there to be distinct indiscernibles.

It is perhaps because we expect variation in modal space, as (macroscopic) change in time, to lack abrupt variation that Adams's case of almost-indiscernibles inclines us to think that it is somehow capricious to deny the possibility of distinct indiscernibles, as if it means positing a gap in modal space. But having examined these three different attempts to bridge the surely-gap, I am sceptical of the prospect of doing so, and suspect that the intuitive appeal of the argument sketched by Adams is a reflection of our intuitions about the natural course of events, as opposed to an intuition of an entailment from the possibility of almost-indiscernibles to the possibility of distinct indiscernibles.

7. Distinct indiscernibles as figuring in the best explanation of contingent phenomena: the case from quantum physics.

In a survey article on the identity of indiscernibles, Forrest (2012) suggests that the identity of indiscernibles appeals to empiricists, for it seems we could never have any evidence that indiscernible items are distinct. We may have evidence that intrinsic duplicates exist, but in order to recognise that they are distinct, they must somehow be differently related to us: in order for someone to be able to count two of them, one must have a property that the other lacks, such as being to the left of the person doing the counting. Coupling this with the premise "that there are no two things which are not empirically distinguishable", we reach the conclusion that the identity of indiscernibles holds (Forrest 2012, Section 3 (i)).

In the following chapter I will defend the principle that, if we consider that there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then we should say that *p* is possible. Interestingly, Forrest also seems to rely on such a principle in his discussion of the empiricist argument against the actual truth of the identity of indiscernibles. Concerning the empiricist's premise "that there are no things which are not

empirically distinguishable” (Forrest 2012, Section 3(i)), Forrest goes on to say the following:

Presumably the premiss would not be proposed as anything more than contingently true. For there are possible situations in which there would be theoretical reasons for believing in indiscernible items as a consequence of a theory which best explains the empirical data. Thus we might come to hold a theory of the origins of the physical universe which had large amounts of empirical support, and which implied that, in addition to our enormously complicated universe, various simpler ones had been generated. For some of the simplest universes this theory might imply that there were exact replicas. (Forrest 2012, Section 3(i))

Forrest is here appealing to hypothetical evidence to cast doubt on the empiricist’s premise; but we may have actual evidence to do the same job, for it has been argued that certain phenomena of quantum physics are best understood in such a way that implies the existence of distinct but indiscernible objects. French (1989) shows that, if we are allowed to identify state-dependent properties of particles “with all the monadic and relational properties which can be expressed in terms of physical magnitudes that can be defined for the individual particles,” a suitably weak version of PII is violated in quantum mechanics, for “two bosons or two fermions in the appropriate superposition state have the same monadic properties and the same relational properties one to another” (French 1989, 158).

This has not gone unchallenged. Saunders (2006), for example, points out that, while it is indeterminate which of a pair of maximally entangled fermions is spinning in which direction, it is determinate that each is spinning in the opposite direction to the other, and that neither is spinning in the opposite direction to itself. This means maximally entangled fermions are weakly indiscernible, meaning that they stand in some irreflexive relation. If we

allow PII to incorporate such relations, the result is that maximally entangled fermions, while distinct, do not violate the principle – they are, after all, (weakly) discernible. As for entangled bosons, Saunders maintains PII by claiming that “we would do better to say that there is only a single particle present (with proportionately greater mass)” (Saunders 2006, 60). Others have offered alternative – albeit non-standard – interpretations of quantum mechanics, which allow PII to be retained (van Fraassen 1985 and 1991).

Thus the dispute over the significance of quantum physics for PII is a matter of some controversy, and not one that I am able to adjudicate. But at the very least, we can note a presupposition of that debate: best theories of actual contingent phenomena can be a guide to possibility. Indeed, the very effort to show that the account of quantum phenomena that entails the falsity of PII is not the *best* account seems to presuppose that, *had* it been the best account, we would have had grounds for denying PII.⁴⁷ So what about best theories of merely possible but contingent phenomena? In the next chapter I explore the idea that the best accounts of possible phenomena should themselves be taken to be possibly true. In this way, possible evidence can be a guide to possibility. More precisely, when we consider that there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, we should conclude that *p* is possibly true. This seems to be the principle behind Shoemaker’s (1979) argument for the possibility of time without change. In the next chapter, I’ll use that argument to show how the principle can provide a way of arguing for a certain possibility claim where there are dialectical constraints on the probative force of imagination.

⁴⁷ Matters are complicated, of course, by the fact that some will take a theory’s implication of the existence of distinct indiscernibles as a reason for rejecting that theory (one person’s *modus ponens* is another’s *modus tollens*). Given his long-term commitment to an empiricist/anti-scientific realism stance, and the fact that an empiricist stance can be taken to favour PII, this is presumably at least part of what drives van Fraassen’s alternative interpretation of quantum physics. However, as noted in my Introduction, I am taking scientific realism for granted.

4.

Part I: Imaginative resistance, possible evidence and genuinely new ideas.

1. Genuinely new ideas.

In the previous chapter we encountered some difficulties for demonstrating to the advocate of PII the possibility of distinct indiscernibles. I only briefly touched on the imaginability of such a scenario, and did not question the notion that Adams's (1979) appeal to almost-indiscernibles might help show that distinct indiscernibles, such as Black's (1954) spheres, are imaginable: it seems to require no great leap of the imagination to proceed from a mental image of a scenario where some small asymmetry obtains, to one where no asymmetry obtains.

But perhaps things aren't as simple. Gendler and Hawthorne raise the following question: "can imagination/conception enable us to gain access to genuinely new ideas? Or do they merely *re-present*, albeit perhaps in a modified form, what we are already acquainted with through other means, such as perception?" (Gendler and Hawthorne 2002, 9).

One plausible and powerful principle for generating modal claims is the principle of recombination: "that anything can coexist with anything else, at least if they occupy distinct spatiotemporal positions. Likewise, anything can fail to coexist with anything else" (Lewis 1986, 88). Behind this lies another plausible principle, the Humean dictum, that there are no necessary connections between distinct existences. However, often modal disputes are

primarily about whether certain things really are distinct, where the possibility of separation cannot be presupposed without question-begging (we see this in the debates over whether time=change, or whether identity *just is* indiscernibility). In such cases, where the identity or distinctness of some things is what is at issue, the principle of recombination offers little help; for example, recombination cannot tell us whether time could exist without change until we find out whether time and change are the same or distinct. In some cases, we might appeal to imagination or conceivability to help us decide whether some things are distinct, and then let the principle of recombination do its work. As Schaffer has noted, “[r]ecombination and conceivability are interrelated. Part of the justification for recombination is the conceivability of what results” (Schaffer 2005, 12). But one will only grant the conceivability of what results if one grants the conceivability of the elements to be recombined. Thus there is a certain limit to the capacity of recombination to resolve disputes about what is possible. It seems fair to say that recombination counts as re-presenting, in a modified form, what we are already acquainted with.

In this chapter I’ll be looking at a potential route to modal claims, a route which takes *possible evidence* in favour of a proposition to count in favour of the *possible truth* of that proposition. I’ll provide two cases for which this route to modal claims might be appropriate and useful: first, the case of Black’s (1952) spheres, as a counterexample to PII; and second, the cases of whether time can pass without change (with a focus on Shoemaker’s (1969) classic paper on the question). Owing to its appeal to possible ampliative inferences, this route to modal claims can also be a route to genuinely new ideas. It is this feature that helps carve out a niche for possible evidence qua guide to possibility, for, if such appeals are legitimate, they bring with them the potential for arguing for a possibility without having to depend on the imaginability of the possibility being argued for, which can be useful where

imaginative resistance is a factor. To illustrate, I'll close the chapter with some further examples where appeal to possible evidence can be fruitful.

2. The dependency thesis and the limits of experience.

In chapters 1 and 2 I suggested some ways in which experience can be brought to bear on metaphysical questions. However, there is also much of interest in metaphysics that lies beyond experience. As Hume famously argued, experience cannot reveal necessary connections; and possibility claims that are the subjects of metaphysical speculation are typically non-actualised possibilities, and hence not to be justified by perception.

There are also some metaphysically significant scenarios which by their very nature lie beyond experience. Black's case of two indiscernible spheres is arguably one such a scenario. In order for one thing to be distinguished in experience from another, something has to be true of one that is not true of the other, for example, that one is to one's left, the other to one's right, or that one is observed before the other. That is, in order for two things to be experienced and counted as distinct things, they must be somehow differently-related to the observer, hence discernible.

With this in mind, the move from the imaginability of *almost*-indiscernibles to the imaginability of distinct indiscernibles cannot be taken for granted; whether cases of the latter are imaginable is contingent on the relation between imagination and experience.

The Dependency Thesis, as M.G.F Martin has named it, proposes a partial answer to the question of how sensory experience and imagination are related. Here is Martin's statement of it:

to imagine sensorily a φ is to imagine experiencing a φ .

In other words, “we visualise objects by imagining visually perceiving them” (Martin 2002, 404). Martin’s statement of the principle itself does not say anything about how imagination and experience relate to *possibility*. One might think that a sensory imagining of something *entails* the possibility of that thing; in that case, the Dependency Thesis suggests that, when we sensorily imagine a φ , we know, not only that φ is a possible object or event, but also that it is possible for us to have an experience of a φ . One might, of course, hold that the link between imagination and possibility, while epistemologically significant, is weaker than that: imagination is a guide to possibility, but a fallible one. In that case, when one imagines a φ , one has reason to believe that it is possible for there to be a φ , and that it is possible for one to have an experience of a φ , but also that one might be mistaken about this.

On an earlier statement of the Dependency thesis, by Andre Gallois, however, it seems that something more on the relationship between imagination, experience, and possibility is built into the thesis itself. Here is Gallois’s version:

Whenever one can be said to image, one can be said to image a potential experience and moreover an experience which is potentially one’s own. (Gallois 1974, 62)

Here ‘to image’ can presumably be understood as meaning to visually imagine. Extending the claim to cover all sensory imaginings, and if we read ‘potential’ here as implying ‘possible’, we can read Gallois’s version as saying

One sensorily imagines a φ only if φ is possibly experienced.

Then the contraposition of this tells us that, if it is not possible to experience a φ , then it is not possible to sensorily imagine a φ ; and if we take it that one can only have an experience of what is possible, then we have it that one cannot imagine impossibilities – if one can sensorily imagine a thing (or event/scenario), then that thing (event/scenario) is possible.

I am offering this reading as a suggestion of a potential link for the link between experience, imagination, and possibility (though of course, nothing forces us to read Gallois's statement of the thesis in this way). However, even if this account is correct, one might wonder how helpful this would be. Gendler and Hawthorne voice the worry for potential accounts of the link between possibility and *conceivability*:

[T]he challenge is to come up with characterizations that are both sufficient to the task at hand and non-circular. For, of course, one might define 'conceivable' in such a way that P is not *really* conceivable unless P is possible. While this will resolve the problem of reliability, unless there is an independent way of determining that we are conceiving in the relevant sense (and not merely seeming to conceive), the practical significance of the link will be negligible. (Gendler and Hawthorne 2002, 11)

One might wonder about the practical significance of the link drawn by the present reading of Dependency Thesis, between imaginability and possibility. As we saw, my reading of Gallois's dependency thesis does seem to force the claim that imaginability implies possibility (if to imagine something is to imagine a potential experience of that thing (which is what the Dependency Thesis says), and every potential experience is a possible experience, then imaginability implies possibility). But even if we grant that, we still have Gendler and Hawthorne's worry to deal with: *unless there is an independent way of determining that we*

are imagining in the right sense (and not merely seeming to imagine), the practical significance of the link will be negligible.

A natural move in response to Gendler's and Hawthorne's worry is to brand it a sceptical one, and respond accordingly: the mere possibility of my *merely seeming* to conceive/imagine does nothing to prevent me from knowing that I am in fact imagining (when I *am* in fact conceiving/imagining). This is an instance of a familiar kind of response to a familiar kind of objection. The *possibility* of one's evidence leading us astray, or *merely seeming* to be evidence, is suggested to count against any claims made on the basis of that evidence; but any such suggestion is then found guilty by its association to external world scepticism of being fallacious, or at least of being superfluous, invoking the reasoning of the external world sceptic but adding nothing to the latter sceptic's conclusion.⁴⁸

Can we rightly consider Gendler's and Hawthorne's worry to be analogous to more familiar and general scepticism? Yablo considers whether the scepticism with regard to conceiving (the discussion would apply no less to the case of imagining) might be disanalogous to general external world scepticism in the following way:

Granted the unavailability of any philosophically satisfying *reason* to think that perception is adequate to its task, we see at least how it *could* be. In fact perception itself brings word of sensory mechanisms seemingly hard at work monitoring external conditions. By contrast "we do not understand our own must-detecting faculty." [(Blackburn 1986, 119)] Not only are we *aware* of no bodily mechanism attuned to reality's modal aspects, it is unclear how such a mechanism could work even in principle. (Yablo 1993, 3-4).

⁴⁸ This kind of response can be found in Schaffer (2005) and Hawthorne (2001), in defence of a quiddistic account of properties, and Stanford (2001), against certain arguments from underdetermination to scientific anti-realism.

Yablo responds, that “the same could be said about various other faculties, notably logical and mathematical intuition; and to judge by our reaction there, they constitute a reason less for mistrusting the faculty than for reconsidering either the nature of the target facts or the nature of our access to them” (Yablo 1993, 4).

My aim is not to provide a full-fledged defence against general scepticism towards the justification of modal claims. As mentioned in the Introduction, I take it for granted that some possibility claims can be justified by imaginings and conceivings. But one can still argue that in some cases what is being imagined is not what the imaginer thinks it is. As Peacocke suggests, imagery alone does not settle what is being imagined: whether one is imagining a suitcase, or a suitcase with a cat hidden behind it, depends on supposition rather than imagery, for the image, it seems, will be the same in either case (Peacocke 1985, 19). While the imaginability of either of the latter possibilities is uncontroversial, Peacocke’s point paves the way for disputes about what it is that is being imagined – about, one might say, the content of our imaginings. (Such a dispute is what motivated Adams’s (1979) argument from almost indiscernibility, discussed in the previous chapter.) It is in such disputes that appeal to possible evidence can help further the debate.

Not in all such disputes, however. Whether appeals to possible evidence can be convincing depends on some semantic and conceptual points, in particular whether the concepts in question are functional in the same way that (as mentioned in the introduction) one might argue that the concept of time is functional, and whether one adopts the reasoning that Kripke famously uses to argue that water is essentially H₂O, and that dispositional essentialists use to argue that the laws of nature are necessary. It seems that appeals to possible evidence cannot

do much to argue against the latter positions; nonetheless, as I'll argue in section 10.4 there remains a substantive role for appeals to possible evidence to play among those who do not have an antecedent commitment to such essentialist theories

3. Imaginative resistance.

In her survey article on imagination, Gendler characterises imaginative resistance as follows:

Imaginative resistance occurs when a subject finds it difficult or problematic to engage in some sort of prompted imaginative activity. Suppose, for example, that you were confronted with a variation of *Macbeth* where “the facts of [Duncan's] murder remain as they are in fact presented in the play, but it is prescribed in this alternate fiction that this was unfortunate only for having interfered with Macbeth's sleep” (Moran 1994). If you found it difficult to imagine this, even though the author had done everything authors usually do to make such a story fictionally true, then you would be experiencing imaginative resistance (Gendler 2013, Sect. 5.2)

Gendler goes on to distinguish two ways of interpreting a claim of a philosopher to be unable to imagine something, by way of two kinds of theory: “Can't” theories, and “Won't” theories. *Won't* theories involve the would-be imaginer being unwilling to imagine something, for instance, “because doing so might lead them to look at the (actual) world in a way that they prefer to avoid” (Gendler 2013, Sect. 5.2). However, a *won't* theory with regard to scenarios like Black's spheres would bring with it an uncharitable accusation of disingenuousness. In general, then, a *can't* theory would be more appropriate to explain imaginative resistance to cases like Black's.

Simple *can't* theories often embrace some sort of *impossibility hypothesis*, suggesting that propositions that evoke imaginative resistance are impossible in the context of the stories where they appear, and that this explains why readers fail to imagine them as true in the fiction. ... Brian Weatherson (2004)

offers a more sophisticated version of a *can't* theory, suggesting that resistance puzzles arise in the face of a certain *type* of impossibility: they arise in cases where the lower-level facts of the story and the higher-level claims of the author exhibit a particular kind of incoherence. (Gendler 2013, Sect. 5.2)

A *can't* theory along the lines suggested by Weatherston (2004) can help us to understand why someone might, doubt the imaginability of Black's spheres, on principled grounds and without begging crucial questions. (Another principled, but question begging basis for doubt would be the conjunction of the claims that whatever is imaginable is possible, and that distinct indiscernibles are impossible. This entails that Black's spheres are unimaginable, but only by assuming that distinct indiscernibles are impossible, something that imaginative failure might otherwise be invoked to show). Recall the Dependency Thesis: to imagine something is to imagine an experience as of that thing. With this in hand, we can see why someone might reasonably deny that Black's spheres are imaginable, and claim that when someone feels as though they are imagining Black's spheres with ease, it really only *seems* to them that this is what they are doing. In fact, it is impossible for us to do so; for, assuming it is impossible to ever have an experience of distinct indiscernibles – because the would-be experiencer must bring some asymmetry to the scenario if they are to distinguish one object from another – the Dependency Thesis tells us that it is also impossible to imagine distinct indiscernibles: When we think we are imagining Black's scenario, we are actually imagining an experience of two discernible spheres; distinct indiscernibles cannot (qua indiscernibles) be experienced, and hence cannot be imagined.

My contention is that appealing to possible evidence can be fruitful when one is faced with imaginative resistance. When one of the participants in a debate doubts the imaginability of the scenario whose possibility is what is up for debate, this of course rules out a *direct* appeal to imagination to persuade them, and the same goes for conceivability. Such a dialectic has

the makings of dialectical deadlock, a mere clash of intuitions. This is why, as we saw in the previous chapter, it made sense for Adams to begin with the significantly different (but apparently very close) scenario, in which there are *almost* indiscernibles: that scenario is significantly different in the sense that all sides can agree on its imaginability, conceivability, and possibility.

But I have also argued that almost indiscernibility does not help; at least, in the three proposals I examined, there is no dialectically effective bridge from the possibility of discernible to the possibility of distinct indiscernibles. In the first two, the gap is bridged at the expense of making almost indiscernibility irrelevant to the argument; whereas, in the third, the argument only gets us to the conclusion that the limit of a certain series is possible, without telling us that the limit is one of distinct indiscernibles.

My suggestion that, in some debates at least, we appeal to possible evidence, has a similar rationale to Adams's argument. If the supposed unimaginability or inconceivability of something is forms part of one side's resistance to a certain possibility claim, this rules out any appeal to imagination or conceiving to justify that possibility. Where imaginative resistance is based on a *can't* theory, such as the one arising from the Dependency Thesis, appealing to the imaginability of apparently very similar scenarios does not help. To see how an appeal to possible evidence can help get around such deadlock, I'll illustrate with what is in several respects an analogous debate, and which provides some precedent for the route to modal claims I am advocating: the debate concerning whether time could pass without change.

4. Shoemaker's argument from the possibility of partial freezes.

At the centre of the following discussion is Shoemaker's (1969) argument for the possibility of time without change, and the principle that it seems to presuppose: that if we consider that there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then we should believe that *p* is possible.

Shoemaker's main target is the Aristotelian claim that "time involves change because the awareness, or realization, that an interval of time has elapsed necessarily involves the awareness of changes occurring during the interval" (Shoemaker 1969, 368). Notably, Shoemaker grants part of this claim. He says "it is plausible to suppose that as long as one is aware of the passage of time some change must be occurring, namely, at a minimum, a change in one's own cognitive state" (Shoemaker 1969, 367), and hence that "it is logically impossible ... for someone to be aware of the existence of a changeless interval during that interval itself" (Shoemaker 1969, 368). The part of Aristotle's claim that Shoemaker denies, of course, is that we should conclude from this claim, about our *awareness* of passage, that time necessarily involves change. By way of a rather ingenious thought experiment, Shoemaker aims to show that we could have empirical evidence "such that no sound argument against the possibility of such intervals can be built on a consideration of how time is measured and of how we are aware of the passage of time" (Shoemaker 1969, 368).

The next section explores what Shoemaker's thought experiment says about the reach of induction. I will be focusing on two kinds of consideration that might be brought against the possibility of time without change. The first is the empiricist worry that since experience presupposes change in mental state, one could never experience the passage of time without some change occurring during that time. The following section suggests that Shoemaker's

imagined inductions are legitimate, in spite of their jointly entailing a conclusion beyond the metaphysical possibility of direct confirmation or falsification.

The second kind of consideration has to do with the link between imaginability and possibility. Again, from an empiricist stance, one might be content to accept something beyond possible experience as being genuinely possible if one is nevertheless able to imagine that thing; but those, like Hume, who deny the possibility of time without change, might claim that they find themselves unable to imagine it. For Hume's part, he claimed that "tis impossible to conceive ... a time when there was no succession or change in any real existence" (Hume 1967, Book 1, Part 2, Sect. 4), and it seems fair to say that 'conceive' is here being used in a sense in which we might use 'imagine'. In section 7 I show how Shoemaker effectively grants, first, an imagistic guide to possibility, and, second, that no mental imagery can be conjured of time passing in the absence of change. I also show that the claim that time without change is impossible to imagine falls naturally out of a certain conception of what is to imagine something, namely the aforementioned Dependency Thesis.

According to the Dependency Thesis, whenever one imagines something, one imagines oneself from the inside having an experience as of that thing; hence it is easy to see why its advocates might think that any claim to have imagined time without change must be incoherent. Hence, by assuming that no mental imagery can be conjured of time passing in the absence of change, Shoemaker is only being fair to his natural opponent. But it also means that, if the Dependency Thesis is correct, and if Shoemaker's argument succeeds, then it succeeds in taking us beyond, not only the possibly experienced, but also beyond the possibly imagined.

I will argue that Shoemaker's argument does succeed in giving us good reason to grant the possibility of time without change. I take it that a proper assessment of the Dependency Thesis is beyond the scope of the present discussion, but I will suggest that a commitment to the Dependency Thesis is likely what made Shoemaker's argument necessary in the first place – without the Dependency Thesis, one might find time without change just as within the bounds of the imaginable as the scenario that Shoemaker asks us to imagine (at no point does he ask us simply to imagine a period of time in which nothing changes). On the one hand, Shoemaker has given us grounds to grant the possibility of time without change even if it is unimaginable; and, on the other hand, if the Dependency Thesis is false, Shoemaker's argument is superfluous: we should have no problem with directly imagining a period of time in which nothing changes.

5. Beyond potential experience, via possible inductions.

Shoemaker (1969, 370-1) asks us to imagine the following scenario. The universe is divided into three populated regions, A, B and C. The inhabitants find that each region undergoes periodic freezes, each freeze lasting one year. They can come to know this because the local freezes are staggered: by observations of clocks in unfrozen regions in the times before and after the freezes, the inhabitants find that region A freezes for one year during every third year; region B freezes for one year every fourth year; and region C freezes for one year every fifth year. Having noticed this regularity of local freezes, they realise that if each individual region continues to follow its observed pattern of freezes, then every sixtieth year will be one in which all three regions are frozen (for all and only multiples of 60 are multiples of *all three* numbers, 3, 4 and 5). Since the three regions exhaust this imagined universe, this would

amount to a *total* freeze: every sixtieth year would be one in which time passes without change.⁴⁹

It should be noted that the three universal generalisations on which Shoemaker's case depends *are* falsifiable by observation. Recall that these are: a) Region A freezes for one year every third year, b) region B freezes for one year every fourth year, and c) region C freezes for one year every fifth year. Take region A for example. All that is required to falsify a) is that some ordinal year that is a multiple of 3 should be one in which region A is observed to remain unfrozen. So, for instance, according to a), b) and c), the thirty-third year should be one in which A freezes, while B and C do not; this means that, should the thirty-third year be a counter-instance to a), the inhabitants of regions B and C would be in a position to see a) falsified. And of course, analogous remarks apply to the other two generalisations.

What is *not* confirmable/falsifiable by direct experience is the claim that is *deduced* from a), b), and c): that the sixtieth year is a year of total freeze. This is of course because observation presupposes change, and no change could occur during a total freeze (if such a freeze is possible). If there is anything wrong with a), b) and c), then, perhaps it is not so much what each of them says, but rather that their conjunction *entails* a claim which itself is neither confirmable nor falsifiable by direct observation.

However, arguably the inductive inferences made by the inhabitants are no different from more mundane ones in this respect. Consider the claim that thunder is always preceded by

⁴⁹ Newton-Smith (1980, 19-24) provides a modified version of this argument: replacing the local freezes with disappearances, he aims at the stronger conclusion that time could pass, not only in the absence of change, but in the absence of any concrete objects. Everything said here should apply equally to Newton-Smith's version.

lightning. We can easily imagine a scenario in which one arrives at this claim by way of induction: For every time one hears the sound of thunder, lightning has been observed to occur before it; having repeatedly observed this correlation, we infer that thunder is always preceded by lightning (we might also allow, as Shoemaker might in his case, that testimony plays some role in our coming to believe this). Now it sometimes happens that upon hearing the sound of thunder, I turn in hope of catching a glimpse of the lightning. When this happens I am invariably disappointed, and instead simply reminded that lightning does not usually wait around for the time it takes to turn one's head. Yet the fact that I am unable to confirm that *particular* occurrence by observation does not stop me from continuing to maintain that thunder is always preceded by lightning. On the contrary, I feel justified in positing the particular occurrence of lightning because it is entailed by the universal generalisation (in conjunction with the claim that thunder did occur) to which I have inductively inferred based on past observations. While this universal generalisation is itself confirmable/falsifiable by direct experience, it can be used to derive consequences which are not.

It may be objected that there is a difference here: the claim that lightning occurred is one about a *kind* of thing that could *in principle* be observed, whereas the claim that a total freeze occurs is not. The question is then whether this difference is important enough to draw a dividing line between those inductive inferences which are legitimate and those which are not.

The above case suggests that the use of induction to draw conclusions lying beyond what the *theorist* could confirm or falsify by direct experience is not without precedent. For in the lightning case, the fact that I *could have* confirmed/falsified the statement about the particular

lightning bolt, *had* I been looking in the right direction, makes no difference to my actual evidential situation when I make the inference. In an important sense, then, my claim is beyond what I *could* observe, and yet, arguably, this does not by itself render my inference illegitimate.

Suppose that's right. Now in order to move from the legitimacy of the inference in the lightning case to the legitimacy of the inference of the inhabitants of Shoemaker's world, we can co-opt the reasoning of scientific realists. Writing about unobservables in science, Peter Lipton makes the point well:

[T]he inferential path to unobservables is often the same as to unobserved observables. In these two sorts of case, the reasons for belief can be equally strong, so the suggestion that we infer truth in one case but not the other seems perverse. Perhaps there could be creatures whose inductive mechanisms made them constructive empiricists, but they would be different from us. (Lipton 2004, 145)

[T]he instrumentalist must claim that there is some principled distinction between inferences to the observable and to the unobservable, even if the paths are the same. On first hearing, this may sound plausible: the existence of the unobservable seems by its nature more speculative than the existence of the observable. This intuition might be strengthened by appeal to underdetermination. Only theories that traffic in unobservables have truth values underdetermined by all possible evidence. Nevertheless, the claim is misguided. The relevant distinction, if there is one, is between the observed and the unobserved, not between the observable and the unobservable. What counts for our actual epistemic situation is not ideal underdetermination by all possible evidence, but the much greater actual underdetermination by the evidence we now have. But neither the realist nor the instrumentalist is willing to abjure inferences to the truth of the unobserved, since this would make the predictive application of science impossible. To show that scientists are not entitled to infer unobservables, it would at least have to be shown why these inferences are all more precarious than inferences to the observable but unobserved, but no good reason has been given for this. (Lipton 2004, 200)

There is a great debate between scientific realists and anti-realists, and I cannot settle it here. But if we grant Lipton's points (and recall that I am proceeding from a scientific realist stance), we should also grant the legitimacy of the possible inference to the claim that time passes without change, for the path to that claim seems to be the same as the paths to claims about unobserved observables (such as unseen bolts of lightning).

It should also be noted that the hypothesis that every sixtieth year in Shoemaker's world is anomalous (i.e. that there is continuous change and no total freeze) is equally beyond direct confirmation by experience – it involves denying that there are total freezes, a denial which one cannot get from simple observation. Also, one can hardly deny the validity of the inference from a), b) and c) to the claim that there are total freezes in Shoemaker's world; so to deny that last claim, one must deny one or more of the universal generalisations. Since each of these is in all relevant respects analogous to the others, one cannot reasonably deny one without denying all. But note that if one does deny any of the generalisations, it is not because any of them has been falsified by unadulterated observation, for the regularity of the partial freezes is stipulated as part of the case.

6. Warmbrod's Criticism.

Since direct experience of a total freeze is not possible (since the would-be subject of that experience would be frozen, too), inductive inferences play a crucial role in Shoemaker's thought experiment: for, by their nature, inductive inferences involve projecting beyond what is actually experienced. However, another familiar feature of inductive inferences is that their conclusions are not logically entailed. As Shoemaker admits, "[w]hat [the inhabitants] observe is equally compatible with the generalization that freezes occur with these

frequencies *with the exception* that all three regions skip a freeze every fifty-nine years” (Shoemaker 1969, 372-3). Nevertheless, Shoemaker suggests (or at least, can be taken to suggest), the inductive data make it objectively more reasonable for the inhabitants to infer a total freeze in the sixtieth year, and hence more reasonable for us to infer the possibility of a total freeze.⁵⁰

Ken Warmbrod (2004) has disputed just that: on the contrary, he says, “Shoemaker has described a world in which it is *unreasonable* to believe in temporal vacua” (Warmbrod 2004, 270).⁵¹ This is because every observation supporting a local freeze presupposes that there is a region in which change continues during that time. Hence, “corresponding to every observation of a freeze in any region of the world, there is at the same time an observation that might be made of change in a different part of the world” (ibid). As such, the inductive data provide *at least as much* support for the generalisation that local freezes are always accompanied by change at some other region. Crucially, this would entail that the sixtieth year is not one in which a total freeze occurs. Hence, the suggestion is, it is not objectively more reasonable for the inhabitants to make the inference to a total freeze.

Warmbrod goes on to suggest that this problem with Shoemaker’s argument is general: it cannot be remedied by modifications of the case.

The underlying difficulty is implicit in the epistemology of claims about observations of freezes. The observation and timing of a freeze requires there to be a timekeeper of some sort which is unfrozen and

⁵⁰ This is in keeping with Skow’s point, which I cited in Chapter 1, section 1: “Even if it is true that each of the theories is *consistent* with my experience being as it is, it does not follow that the fact that my experience is that way fails to *favor* one of the theories over the other. This is an instance of a general epistemological truth: a body of evidence can support some hypothesis without entailing that hypothesis” (2011, 361).

⁵¹ It seems that Warmbrod understands ‘temporal vacua’ to refer to periods of time in which no changes occur, rather than periods of time in which nothing exists. I don’t think anything central to the present chapter hangs on this.

changing over the period of time in question. Hence the observational data will always support a conclusion that not everything is frozen. (Ibid, 273)

To demonstrate the generality of his objection, Warmbrod describes a modified version due to Newton-Smith (1980, 19-24). The relevant difference in Newton-Smith's case is that in the run up to every observed local freeze, the soon-to-be-frozen region displays some observable property: say, for instance, every part of that region takes on a red glow (we might also stipulate that the red glow remains but gradually fades away once the region has come back to life). Applying this modification to the observed regularities in Shoemaker's case, the fifty-ninth year is now stipulated to be one in which, at some time towards the end of that year, all regions display that same property at once. Since this red glow has been observed on every previous occasion to precede a local freeze in the region it affects, the inhabitants are supposed to infer that the sixtieth year is one in which all three regions undergo local freezes, and hence (since the three regions exhaust the universe) a year of changeless time.

Warmbrod objects as follows:

As was the case with Shoemaker's argument, Newton-Smith's depends on inductive data that support both intended and unintended generalizations. In all cases that we ever actually observe, the display of the special property signals both the disappearance [or freeze] of one object and the non-disappearance [or non-freeze] of at least one other object (since, in observed cases, there is invariably some non-disappeared changing timekeeper that observes and times the disappearance). Hence when everything displays the special property at once, we have at least as much inductive reason to infer that not everything is about to disappear as we have to infer the opposite. (Warmbrod 2004, 274; square brackets are mine)

However, that there is another way of describing what is signalled by the inductive evidence. Arguably, to simply say “every partial freeze is accompanied by a non-freeze in *something else*” is to miss out on the bigger picture, the evidence at large. It is not that the observed property signals both the freeze of one region and non-freeze of at least one other. Rather, it signals the freeze of the object(s) in the region(s) affected by the observed property, and the non-freeze of *everything else*. But since “everything else” has an empty extension when the end of the fifty-ninth year approaches (because everything has the special property towards the end of the fifty-ninth year), the occurrence of the observed property in *all* regions signals a freeze in each region, *without* signalling a non-freeze in at least one other. Hence, by this simple modification of the case, Newton-Smith has – contrary to what Warmbrod argues – described a situation in which it would be more reasonable to posit that every sixtieth year is one of total freeze than to posit that every sixtieth year is an exception to the regularity of previous freezes.

7. Beyond imagination, via imagined evidence.

Suppose we grant that Shoemaker’s appeal to induction is sufficiently similar to cases such as the lightning case to legitimise it. Shoemaker’s argument would still have to contend with the claim that total freezes are *unimaginable*.

Shoemaker’s argument seems to at least grant that our ability to imagine certain things provides evidence of their metaphysical possibility; the argument may even be thought to presuppose this: we should grant the metaphysical possibility of his scenario (in spite of its physical impossibility), because we are able to imagine it. But what about cases where we lack this capacity? Indeed, the supposition that time without change cannot be imagined is crucial if an argument like Shoemaker’s is to be of any use: if we find ourselves able to

imagine time without change, we *already* have just as much ground for entertaining that possibility as we do for entertaining the possibility of Shoemaker's imagined scenario – the whole point of which was to *provide* that ground. Hence what is needed is a way of understanding the link between imagination and possibility without construing our capacity to imagine something as a *necessary* condition on its (metaphysical) possibility.

As Gallois (1974) suggests, reflection on the reason why our capacity to imagine something is not a necessary condition for possibility can be illuminating here. It is no doubt correct to say that inability – or perhaps rather, repeated failure upon trying – to imagine something does not *entail* that it is impossible. We believe this because we know that there are things which we cannot imagine, but which we nevertheless believe to be possible; we also suspect that there are contingent limitations on our imaginative capacities more generally.

Nevertheless, “hypothesizing a necessary connection because of a general failure to imagine things of a certain sort may turn out to be the best explanation for such failure” (ibid, 62).

Call hypotheses of this type *impossibility hypotheses*. Now in some cases, hypothesising the impossibility of something to explain such failure seems too drastic: there is no need to invoke the impossibility of a chiliagon⁵²-shaped object to explain our inability to imagine one, because our inability to imagine one is easily (and more plausibly) explained by hypothesising contingent limits on our powers of imagination (ibid). The question is then whether the relevant impossibility hypothesis in the case of time without change can be similarly undercut. I suggest that it can, as follows.

Recall that Gallois proposes the following condition on visual imagining:

⁵² A thousand-sided polygon.

Whenever one can be said to image, one can be said to image a potential experience and moreover an experience which is potentially one's own. (Gallois 1974, 62)

This is a version of what came to be known in the philosophy of perception literature as the Dependency Thesis (Martin 2002), which we encountered above. We saw above that Shoemaker's argument would be of little use to anyone who claimed to be able to imagine a period of total freeze, for in that case, one would already have as much evidence for that possibility as for the possibility of the periods of partial freeze that Shoemaker asks us to imagine. (At the same time, the objection to the possibility of something based on an inability to imagine it implies that, if we did find ourselves able to imagine it, we would have reason for granting its possibility.) Hence is it reasonable to suppose that Shoemaker's argument is aimed at those who deny that periods of total freeze can be imagined; in constructing an argument that does not simply ask us to imagine a period of total freeze, Shoemaker seems to have been mindful of his target audience, those who would be most in need of convincing.

I want to suggest now that the Dependency Thesis provides a good basis for the claim that it is impossible to imagine a total freeze. For consider: what is to stop someone from imaging one of Shoemaker's regions during one of its frozen periods, and simply stipulating that that is all there is to the image, a frozen region, and nothing else? Recall that, as Peacocke (1985) notes, a single image can do double duty: when one imagines, at one time, a suitcase, and at another time at suitcase with a cat hidden behind it, the images one conjures up for each task might be precisely the same (Peacocke 1985, 19). It seems reasonable to say that what differentiates the two imaginings is that one involves a *supposition* that the other does not, namely, that there is a cat hidden from view. Similarly, we might suggest that one simply add to one's image of a frozen region the *supposition* that that region exhausts the possible

world being considered, so that what is being imagined is a period of total freeze, time without change.

But if the Dependency Thesis is correct, then any such attempt must fail. As Peacocke expresses it, “imagining always involves imagining from the inside a certain (type of) viewpoint, and someone with that viewpoint could, in the imagined world, knowledgeably judge ‘I, thus-and-so’, where the thus-and-so gives details of the viewpoint” (Peacocke 1985, 21). Since any experience presupposes the passage of time, and given that any imagining is an imagining of a potential experience, it is impossible to imagine time without change; any claim to be doing just that is based in ignorance of the feature of imagination that proponents of the Dependency Thesis claim to be constitutive of imagination (Peacocke 1985, 21).

However, while the Dependency Thesis seems to be behind the claim that it is impossible to imagine time without change, it does so in such a way as to make the impossibility hypothesis redundant, for we now have available the following argument.

- P1. Whenever one imagines something, one imagines a potential experience.
- P2. Time without change is not possibly experienced.
- C. Therefore, time without change cannot be imagined.

As we saw in the introduction, the truth of P2 is agreed by both Shoemaker’s opponents and Shoemaker himself: indeed, it is perhaps trivial, for awareness of the passage of time presupposes, at a minimum, change in one’s own mental state. As for P1, Gallois suggests, it gains plausibility from introspection.

Suppose that I am asked to image the table in the next room and describe what I am imaging... Now the list of features of the imaged table ... could equally well count as a list of features of what one perceives when one perceives a table. This brings out the connection between imaging and perceiving. In the case of visual imagery the content of an image is limited to what can be taken in at a glance. In this way visual images can be said to reproduce visual experiences that one might have (which is not to say that they are a kind of self-induced visual experience). (Gallois 1974, 63)⁵³

But now since P1 and P2 *entail* that it is impossible to imagine time without change, the search for a *best explanation* of this latter impossibility stops here. Indeed, one might go further and say that P1 and P2, taken together, render inference to the best explanation *inappropriate*. Once we see that an entailment obtains between the two premises and the conclusion (that it is impossible to imagine time without change), there is no need to resort to the form of ampliative inference to explain the inability to imagine time without change. The inability to imagine time without change is already explained by P1 and P2, both of which are consistent with Shoemaker's argument, each of which is independent of whether a total freeze is possible or not. It turns out, then, that P1 and P2, which might have been thought to provide grounds for resisting Shoemaker's argument, actually fit quite nicely with it.

I have said, following Gallois, that the impossibility of something *might* provide the best explanation for our general inability to imagine it: indeed, that this is a way of maintaining the link between our capacity to imagine something, on the one hand, and its possibility, on the other, without making the former a necessary condition on the latter. But why should hypothesising the impossibility of something always have this explanatory potential?⁵⁴ Of

⁵³ This can be fleshed out a little more in terms of Kung's (2010) account – which we will come to very shortly – of why basic qualitative content of imaginings have probative force – namely, by presenting us with a way that space could consistently be filled.

⁵⁴ Perhaps it does not: If one thinks the laws of nature are metaphysically necessary (as Shoemaker came to think later on), one might see scenarios like that of Shoemaker's (1969) as imagined *impossibilities*. But that

course, one might simply take it as a primitive fact about our imagistic capacities that we cannot sensorily imagine metaphysical impossibilities. But if P1 is correct, we have a potential explanation to hand. For we cannot experience something impossible.⁵⁵ But, by P1, something can be sensorily imagined only if it is possibly experienced. Therefore, *if* something is not possible, it cannot be imagined.⁵⁶ The upshot is that, by tying our capacity to imagine to potential experience in this way, we are able to undercut the impossibility hypothesis, while nevertheless explaining why the impossibility of something *would* explain a general failure on our part to imagine it.

However, for some theorists, this way of spelling out the explanatory potential of impossibility hypotheses will be too strong.⁵⁷ Can't we, after all, imagine some impossibilities, such as water being something other than H₂O? If so, a more sophisticated account is required to make impossibility hypotheses contenders for explaining imaginative failure.

Kung (2010) has supplied such an account. Kung distinguishes two aspects of imaginings: the *assigned* content and the basic *qualitative* content, which correspond roughly to stipulated content and imagistic content (recall Peacocke's notion of S-imagining, which helps to

stance does not let Shoemaker's thought experiment get off the ground, for the latter involves physical impossibilities).

⁵⁵ Plausible as this is, Sorenson (2002) lays some groundwork towards disputing it.

⁵⁶ Peacocke makes a similar point in his discussion of Berkeley's "Master Argument": "Since [Berkeley's envisaged] opponent is hoping to argue from imaginability to possibility, he needs a concept for which it is independently plausible to hold that what is imaginable is possible. That is plausible for imagination as characterized in this paper [i.e. as characterised by the Dependency Thesis]: the plausibility rests on the two ideas that what can be experienced is possible, and that experiences that can be imagined are possible" (Peacocke 1985, 31). We might see Shoemaker's opponent – at least, when that opponent is Hume – as rather arguing from *unimaginability* to *impossibility*. Still, it seems reasonable to think that Hume would accept that, if we *could* (perhaps *per impossible*) imagine time without change, then we would have reason for thinking it possible.

⁵⁷ For example, Yablo (1993) provides an account according to which impossibilities are imaginable.

distinguish an imagining of a suitcase from an imagining of a suitcase with a cat hidden behind it). On Kung's account, there is a sense in which we can imagine impossibilities, but generally only where the assigned content does all the work. This does not impugn the modal justificatory force of imagination in general, because on Kung's account, it is the basic qualitative content of imagining that gives it its probative force. This is because the basic qualitative content of an imaginative experience – just as that of a perceptual experience – presents a way that space *could* be filled around the perceiver (Kung 2010, 637). Since the basic qualitative content presents a way space could consistently be filled, failure to put together the qualitative content required to provide an imagining with probative force can (in principle) be explained by the inconsistency of the arrangement we are attempting to construct in our mind's eye. In this way, the explanatory potential of impossibility hypotheses is itself explained by way of a link between imagination and experience, an explanation which dovetails with the Dependency Thesis. But there remains a sense in which we can imagine impossibilities: for Kung, it is the stipulative aspect of imagining – the assigned content – that permits the imagination of impossibilities, for “stipulations and labels are virtually unconstrained, and what minimal constraints there are have no modal epistemological value” (Kung 2010, 634).⁵⁸

* * *

⁵⁸ “[T]he principal constraint on assignments is certainty. I said that so long as we find P believable, epistemically possible in the strongest sense that it is true for all we know for certain, or possibly true for all we know for certain, we will be able to imagine P via stipulation or label. Let P be some proposition whose possibility we are trying to establish via imagining. The mere fact that we find P (or possibly P) believable, and hence are capable of making the assignments required to make P true in the imagined situation, is not good evidence for P's possibility. Believability just is lack of certainty. (Let us use ‘non-certainty’ to denote lack of certainty; it avoids the unwanted connotations of ‘uncertain’.) It would be very odd if our non-certainty counted as evidence of P's possibility” (Kung 2010, 634).

We are used to the idea that induction takes us beyond the *actually* experienced. As the lightning case discussed above illustrates, there is also a sense in which we let it take us beyond what we are *able* to confirm/falsify by direct experience. What makes Shoemaker's such an interesting case is that, given the Dependency Thesis, it seems to suggest that induction can take us that far, and further – beyond the possibly (sensorily) imagined. In this section I suggested a way of making sense of this given Shoemaker's use of an imagistic guide to possibility. The premises involved came from unlikely sources: the first, P1, Gallois cited in service of Berkeley's "master argument" for idealism (to be is to be perceived); the other, the Aristotelian claim P2, is one that has been cited against the possibility of time without change.

However, as long as we are granting the Dependency Thesis, and the claim that it means that time without change cannot be sensorily imagined, the opponent of the possibility of time without change can still appeal to the Dependency Thesis to score a negative point: while there is no convincing argument from the Dependency Thesis to the claim that time without change is impossible, it does *block* any argument from the imagination of time without change to its possibility. In the remainder of the chapter I will defend the view that appeal to possible evidence can provide positive reason to believe the possibility of something, even if that thing is beyond the limits of sensory imagination.

4.

Part II: Possible evidence as a guide to possibility.

8. An implied modal epistemology.

There are two ways one might understand the aim of Shoemaker's argument, one positive and the other negative. Le Poidevin (2010) favours the negative: Shoemaker argues only indirectly for the possibility of time without change, "by showing what is *wrong* with a particular argument against that possibility" (Le Poidevin 2010, 129; emphasis added), namely, the Aristotelian argument that "time involves change because the *awareness, or realization, that an interval of time has elapsed necessarily involves the awareness of changes occurring during the interval*" (Shoemaker 1969, 368; my italics). By showing that it is logically possible for there to be some *evidence* that periods of changeless time have occurred or will occur, Shoemaker has effectively rebutted the premise that I have just italicised.

However, I would like to explore – and ultimately endorse – the other reading: by considering a possible scenario in which the most reasonable conclusion would be that there are periods of changeless time, we *actually* have positive reason for thinking that time without change is possible. This is evidently how Warmbrod reads Shoemaker: "there is a possible world in which inductive evidence makes it reasonable to believe that there are periods of changeless time. So temporal vacua are logically or conceptually possible" (Warmbrod 2004, 269). And again, where he says:

The thrust of Shoemaker's thought-experiment is thus that the observable facts of the imaginary world make it objectively more reasonable for inhabitants to believe in temporal vacua than to disbelieve in

them. But if Shoemaker is correct about this much, then surely the observable facts of the imaginary world also make it objectively reasonable for *us* to believe that there is a possible world where temporal vacua occur. (Warmbrod 2004, 270).

As we saw above, Warmbrod goes on to suggest that the best response to Shoemaker's argument is to show that the observable facts of the imaginary world *do not* make it objectively more reasonable for the inhabitants to believe in temporal vacua than to disbelieve in them. I suggested above that Newton-Smith's modification of Shoemaker's argument gets around Warmbrod's criticism; still, it is interesting that Warmbrod seems to accept that, *had* Shoemaker shown that it would be objectively more reasonable for the inhabitants to believe in temporal vacua, then he would have shown that it is objectively more reasonable for *us* to believe in the *possibility* of temporal vacua.

This principle seems worthy of further consideration. For, on the one hand, the validity of the principle is far from obvious; but, on the other hand, if we abstract from Shoemaker's particular argument, it would seem to represent a more general modal epistemological principle, which I'll refer to as the Possible Evidence Principle (PEP):

PEP: If there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then we should believe that *p* is possibly true.

What is meant by my 'objectively' in the principle? It is objectively reasonable to conclude that *p* when that conclusion is drawn once all relevant information *besides p itself* has been rationally considered (failure to exclude *p* itself would of course render the principle trivial). To make this clearer, two further phrases need unpacking: "all relevant information" and "rationally considered".

In practice, it will be up to us to decide whether we have included, in the world we are considering, all relevant information to the question in hand. It is our privilege as imaginers to determine whether anything has been left out. The inhabitants of Shoemaker's world may not be able to tell whether the three regions they can observe exhaust the possible world they occupy – whether, say, there is a hidden fourth region occupied by changing things – but we can settle this by stipulation, without begging any questions. We cannot of course simply stipulate that in the *sixtieth* year, a hidden fourth region occupied by changing things does not pop into being. That would beg the question by ruling out by stipulation an alternative account of the inductive data. But it would also beg the question to merely stipulate that such a region *does* pop into being every sixtieth year, for this would be to effectively stipulate that periods of changeless time do not occur in the world under consideration. That such a fourth region comes into being can legitimately be offered as part of an alternative account of the inductive data, but such an account would seem to have little going for it besides the fact that it would allow someone who believes that periods of changeless time are impossible to hold onto that belief.

When using the principle to argue against certain theorists, there is a natural restriction that must be placed on the our stipulations about the world under consideration: we can stipulate any information about the world in which the evidence is given, as long as our stipulation can be granted by us *and our opponents* and does not beg any crucial questions. When using the principle for oneself, to help decide whether one should grant that possibly p, the restriction to respect the opponent's premises is lifted, and one need only restrict our stipulations to what is possible by one's own lights.

I have also said that the information at the world we are considering must be rationally considered. The point here is to ensure that PEP doesn't mandate inferences to possibility claims based on what might be concluded by possible agents whose capacity to reason has somehow been compromised – perhaps by an evil demon who ensures that they have a “feeling of rightness or correctness” (Plantinga 1993, 104) when presented with false theories, or perhaps simply by being under the influence of rationality-altering drugs.

The following sections are intended to help illustrate how this principle can be put to use, first by a comparison to what is in many respects an analogous modal epistemology, due to Hanrahan (2007, 2009) (Section 9); then by responding to some potential objections to the principle (Section 10); then by providing some further examples of the principle at work (Section 11).

9. Comparing Hanrahan's Modal Explanationism.

The idea that possible evidence is a guide to possibility also emerges from Rebecca Hanrahan's “Imagination and Possibility” (2007). However, there are many respects in which our views differ. Hanrahan's is a valuable discussion, but in my view it undersells the reach of possible evidence. To see this, some exposition is needed.

Central to Hanrahan's modal epistemology is the idea that “imagination is to the possible as perception is to the actual” (Hanrahan 2007, 127). A similar view can be found in Yablo (1993), where he says that “conceivability involves the appearance of possibility” (7), and is a natural view for proponents of the Dependency Thesis. As Hanrahan notes, “Peacocke seems to be alluding to this analogy when he claims that our inference from what is imaginable to what is possible is based on two claims. “What can be experienced is possible

and [...] experiences that can be imagined are possible”” (Hanrahan 2007, 127, footnote 9; Peacocke 1985, 31). (I suggested effectively the same claims earlier as helping to explain why impossibility hypotheses can in principle explain failures of imagination: to say that what can be experienced is possible is to say that one cannot experience an impossibility; combined with the claim that to imagine something sensorially is to imagine an experience of that thing, we get the claim that if something is impossible, it is therefore sensorially unimaginable.)

But given that non-actualised possibilities do not impress themselves upon our senses, a metaphysical/causal interpretation of this analogy is bound to fail; actualities can cause a perception in us, but non-actualised possibilities cannot. As such, considering the conditions under which a perception is veridical will shed little to no light on the conditions under which an act of imagining depicts a possibility. With this in mind, Hanrahan proposes, “we would do better to consider the conditions under which *we are justified* in accepting a perception as veridical. We need, in other words, to give an epistemological rather than a metaphysical interpretation of this analogy” (Hanrahan 2007, 128; emphasis added).

Hanrahan’s solution can be referred to as *modal explanationism*. This involves a method of conceiving that combines two others: sensorial imagining and storytelling. Sensorial imagining involves calling forth an image of some sort, and here ‘image’ is intended to cover imaginings of the non-visual as well as visual kind. “To imagine that p is to call forth a mental image of a state of affairs in which p is true” (Hanrahan 2007, 128). Storytelling can function, like Peacocke’s supposition-imagining, to make distinctions that images alone do not – to distinguish, for example, an image of a suitcase from an image of a suitcase with a cat hidden behind it. But Hanrahan also intends storytelling as a distinct method of

conceiving: “To conceive that p via a story is to describe or narrate the conditions associated with p’s truth” (Hanrahan 2007, 129).

Just as, according to explanationism, inference to the best explanation can justify claims about the actual world by looking at actual evidence, Hanrahan’s modal explanationism has it that we can justify claims about possible worlds by thinking about possible evidence, evidence conceived of through the combination of sensorial imagining and storytelling mentioned above. In particular, the kind of evidence relevant to Hanrahan’s project consists in non-factive mental states; just as one can have a bear represented to one’s mind without there actually being a bear there, when one conceives of a possible perception *as of* there being a bear in one’s backyard, we have not necessarily conceived of a possible world in which there *is* a bear in one’s backyard. Possible percepts can be misleading, just as actual percepts can⁵⁹. This is crucial, if the aim is to justify the claim that imagination is a guide to possibility, for it would beg the question to assume that whenever we imagine a possible percept, the world imagined is one in which that percept is veridical (Hanrahan 2007, 136, 142).

Modal claims are then justified by comparing the merits of competing potential explanations of possible percepts. Hanrahan asks us to consider her epistemic twin, Rebecca, who has the same mental life as Hanrahan up to the point where Hanrahan intentionally conjures up an image of a bear in her backyard; Rebecca, on the other hand, has an image of a bear in her backyard come to her unbidden.⁶⁰ Hanrahan herself has recently heard that there was a bear

⁵⁹ Note that I am using ‘percept’ here in a non-factive sense.

⁶⁰ The reader may be wondering about the choice of example – after all, who would deny the possibility of a bear? Hanrahan responds: “Even if I concede that I have knowledge of this possibility, I do not yet know how I came to have this knowledge. I do not know how my belief in this possibility is justified. Given the lengths to which epistemologists have gone to show that we are justified in believing other claims we all take for granted, it would be absurd to insist that our modal beliefs, or any subset of them, are so secure as to need no

sighting on someone's porch, and she notes that her neighbours keep many bird feeders stocked up with feed, and that this can be a draw for bears.

Hanrahan argues that her epistemic twin is justified in taking her images as of a bear to be veridical, and that this in turn justifies Hanrahan in believing that (B) *A bear is in Hanrahan's backyard*, is possibly true. It is argued that the mental state of Hanrahan's epistemic twin, and the account that the latter gives of that state, satisfy three requirements, sufficient to justify the epistemic twin in taking her mental representations as of a bear to be veridical:

1. The account of that state as a veridical perception must possess all of the epistemic virtues to a greater degree than any of the other accounts we could have given of that state and it must also increase or at least preserve the epistemic virtues of our best explanation over the long run.
2. The account of that state as a veridical perception must, of course, include or imply that the state is a product of the normal workings of our senses which accurately reflects the way the world is.
3. The state must be sensorial in nature and it must usually (*though not always*) be highly forceful, vivacious, and determinate. (Hanrahan 2007, 130)

Hanrahan justifies the satisfaction of requirements 1 and 2 as follows:

Now because Rebecca possesses nearly all the beliefs I possess, all the information I have recently gathered about bears is available to her. Thus, I know that the story she can tell about these images that best preserves the epistemic virtues of her best explanation is one that deems them to be veridical perceptions. In this story, this bear had in the past found a mound of birdseed to dine on in Rebecca's neighbor's backyard. And this bear (being a creature of habit) returned to these bird feeders (via

justification" (Hanrahan 2007, 127). My own hopes for the reach of possible evidence are thus more ambitious: thinking about possible evidence can potentially justify controversial possibility claims, that we may not have already assented to, and that may be beyond justification by imagination.

Rebecca's yard) to see if there was more that very morning as Rebecca was looking out her back window. Thus, requirements 1 and 2 above have been satisfied. (Hanrahan 2007, 137-138)

As for requirement 3, its satisfaction depends on how vivid Hanrahan's own imagery was, for, being epistemic twins, Hanrahan's and Rebecca's mental states are the same on this score. If it is granted that Hanrahan's sensorial imagination of a bear in her backyard is highly forceful, vivacious and determinate, then it must also be granted that her (possible) epistemic twin's mental representation of a bear also possesses these qualities. But given that some people's imaginations can manifest these qualities to a high degree, we should granted the former point, and hence that 3 is satisfied by her epistemic twin's mental state (Hanrahan 2007, 138-139).

The next step in the argument is to show that, because Hanrahan's epistemic twin is justified in taking her mental state to be veridical, and thus in believing that there is a bear in her backyard, Hanrahan is also actually justified in the claim that there could be a bear in her backyard. Hanrahan does this by noting, first, that Rebecca is justified in taking her image to be a veridical one, and, but for the fact that Rebecca's image came to her unbidden, they are epistemic twins; as such, it seems only reasonable that Hanrahan would draw the same conclusion as Rebecca were the former in the latter's shoes. Secondly, Hanrahan notes that we have been given no reason to think that Rebecca was intoxicated, or her vision obscured or impaired, or that there is a holographic representation in her garden; and since Rebecca's imagery comes to her unbidden, it makes more sense to conclude that her mental state is veridical, than it does to take it to be a product of her imagination, as Hanrahan's was (Hanrahan 2007, 140-141).

There are several benefits to Hanrahan's modal explanationism. One is that it can (indeed, is primarily intended to) provide justification for everyday modal claims (hence the mundanity of the bear example). As she points out, modal epistemologists have a tendency to focus their attention on controversies of metaphysics and philosophy of mind; but when it comes to justification of more mundane modal beliefs, Hanrahan's view is arguably on a firmer footing than, say, Chalmers's (2002) view, which relies on an idealised notion of conceivability, or Bealer's (2002) intuitionism, which is arguably more alien, and its justificatory connection more mysterious, than Hanrahan's explanationism. As Armstrong said, "To infer to the best explanation is part of what it is to be rational. If that is not rational, what is?" (Armstrong 1983, 59; quoted in Hanrahan 2007, 133, footnote 23).

Modal explanationism is also intended to avoid two problems: *the problem of incompleteness* and *the problem of adjudication*. The former befalls what Hanrahan refers to as *set-constructionism*, according to which possible worlds are maximally consistent sets of propositions⁶¹. These sets are infinitely large, so it is (practically) impossible to justify modal beliefs by constructing worlds that make them true. For less strict standards of justification we might settle for constructing a small subset of those infinite propositions; but doing so seems to run a high risk that one of the propositions of the larger set contradicts one from the smaller set. By contrast, modal explanationism is supposed to allow for incompleteness; no one demands that our explanations of actual goings on consist of complete descriptions of the actual world in order for abductive arguments to justify beliefs about the actual world, and so it would be similarly unreasonable to demand such completeness in best explanations of possible phenomena (Hanrahan 2007, 144).

⁶¹ Hanrahan notes that Chalmers's view combines set-constructionism with modal intuitionism (Hanrahan 2007, footnote 40).

The second problem, of adjudication, arises when competing theorists fail to agree on whether something is imaginable or conceivable. A prominent case is Chalmers's "brute intuition" that philosophical zombies are possible – that there is a possible world just like ours except that our counterparts are without consciousness (Chalmers 1996, 96). One may deny that intuition, or, alternatively but to similar effect, one may have an intuition of something else that contradicts it. In an example of the latter, Frankish (2007) mounts an argument against the conceivability of zombies, from an intuition that conflicts with it: the conceivability of "anti-zombies". Anti-zombies are beings very much like zombies would be (if they were possible), except that the purely physical facts about them necessitate their consciousness. But there is no room in modal space for both zombies and anti-zombies, for, if the purely physical features of anti-zombies necessitate their consciousness, they therefore make any would-be zombies conscious too. Frankish suggests that the result is a stand-off, and hence "that considerations of conceivability have little role to play in debates about the nature of consciousness" (Frankish 2007, 666).

Hanrahan argues that her brand of modal explanationism does not suffer from this problem. Now of course there may be problems of adjudication on an explanationist framework; for example, there may be disagreement about which is the simplest explanation. As Jiri Benovsky comments on Shoemaker's argument for the possibility of time without change (which is in part an argument from the simplicity of one account over another):

According to Shoemaker, it would then be simpler, and consequently more reasonable, to avoid believing in such exceptions [to the periodic local freezes described in Shoemaker's thought experiment]. But here, issues about simplicity arise: is it more complicated to believe in a more complicated regularity law or is it more complicated to believe in the possibility of time without

change? One kind of simplicity is mathematical simplicity (equations predicting the behaviour of local freezes would need to be more complicated (but, actually, not very complicated)), while the other kind of simplicity is a conceptual one: to conceive of and understand the notion of time without change.

(Benovsky 2012, 765)

This also illustrates that something similar can be said of coherence. Internal coherence is of course a virtuous property for a theory to have, and can be appealed to in order to justify one theory over another. But one might also judge a theory by how it coheres with other beliefs that one already has. Someone who holds that time is change will therefore be more resistant to the idea that a theory that invokes extended periods of time without change is more virtuous than one that does not (perhaps making the point in terms of ‘conceptual simplicity’).

Nonetheless, Hanrahan’s claim is that adjudication is *possible* on modal explanationism, and not possible on modal intuitionism. While in *some* cases disagreement may persist, on which is the most fruitful, coherent, and simple theory, there will at least be some cases where the abductive success of a theory is enough to overcome some original misgivings about the concepts appealed to in that theory, or misgivings based on the extent to which that theory coheres with previously held beliefs.⁶² That is, it is at least possible for opposing parties to evaluate each other’s theory based on abductive criteria; the same cannot be said (or at least,

⁶² Indeed the success of the theory may make it reasonable to change some of those original beliefs, as Locke appears to do in the fourth edition of his *Enquiry Concerning Human Understanding* concerning the concept of action at a distance. As Kochiras notes in his *Stanford Encyclopedia* entry on Locke’s philosophy of science: “Locke is initially sympathetic to the [orthodox corpuscularian] proviso [of contact action], writing in the first three editions of his *Essay*, “How bodies operate one upon another...is manifestly by impulse and nothing else. It being impossible to conceive that body should operate on *what it does not touch*.” (E II.viii.11, editions 1–3) Yet for the fourth edition, he replaces that claim about how bodies do operate with one about how we can conceive of them operating: “How Bodies produce *ideas* in us is manifestly *by impulse*, [this being] the only way which we can conceive Bodies [to] operate.” (E II.viii.11, edition 4). He also omits a clause appearing in II.viii.12 of previous editions that denied unmediated action at a distance” (Kochiras 2014, section 3.2.2). Locke’s apparent change of mind is down to the success of Newton’s theory of gravity, which implies the possibility of unmediated action at a distance.

it is not clear that the same can be said) for apparently opposing modal intuitions (Hanrahan 2007, 145).

This apparent advantage of Hanrahan's account is due to her adoption of the imagination-perception analogy, the idea that the imagination is to the possible as perception is to the actual.

But what if the problem between the two [disputants] is located in the fact that one can call forth images associated with the proposition in question, while the other cannot? Obviously, two people cannot share and evaluate each other's imagined experiences. But if we allow that there is some relationship between what we imagine and what we perceive, one person can provide some guidance to another. Consider here how you might get someone who did not know what a chimera was to imagine one. You would ask them to imagine a lion, a goat, and a serpent, then you would tell them how to combine the three. In this way, some kind of conversation can be had about what is or could be imagined. Thus, adjudication is at least possible on my modal epistemology; it is not in modal intuitionism. (Hanrahan 2007, 145)

It is less clear, however, that Hanrahan's account enjoys much of an advantage over intuitionism in the way described here. What is being described here is effectively recombination: to show someone what a chimera is, and that it is imaginable, we can get them to recombine things more readily imaginable. Presumably something similar can be done for modal intuitionism: where there is agreement on the intuitability of some things, they can be recombined to produce a further idea, provided that the recombination does not result in a contradiction. Moreover, one may have principled reasons for claiming that something is in fact unimaginable, perhaps despite initial appearances. As I have suggested above, this may be the case if both a) to imagine something is to imagine from the inside an experience of that thing, and b) for whatever reason the thing to be imagined is impossible to

experience (as in the case of perfect symmetry and time without change, if the Dependency Thesis is granted). As we have seen in the case of distinct indiscernibles, for example, one might have a principled reason for blocking the move from the imaginability of distinct spheres qualitatively identical but for some small asymmetry, to the imaginability of distinct spheres qualitatively identical (without the asymmetry). In such cases, it will not be a simple matter to show someone how to imagine something, for they will naturally claim that what results is merely apparent imagining of the possibility in question.

What I am proposing with PEP is a way of arguing in favour of certain possibilities whether or not they *can* be imagined. For this purpose, some features of Hanrahan's account are crucial. One is the role of ampliative inference. The justificatory criteria Hanrahan appeals to (1, 2 and 3 above) are purely epistemic, and hence do not necessitate the *truth* of one explanation over another (Hanrahan 2007, 131). Given that her aim is to *justify* the move from the imagination of something (a bear in her backyard, in her example) to its possibility, appeal to imaginability per se is not enough. That's where explanationism comes in: the possibility of a real bear best explaining the possible appearance of a bear is invoked to justify the move from the imaginability to the possibility of a bear in the backyard.

For me, too, imaginability is not enough, but for a different (albeit still dialectical) reason, and explanationism is brought in not to justify the move from imaginability to possibility, but rather to justify certain possibility claims even where their *un*imaginability has been granted. As we have seen, there are some cases where imaginative resistance is firmly rooted in a claim about what it is to imagine something, the Dependency Thesis, which is a natural view to have when the analogy between imagination and perception is endorsed (imagination is to possibility as perception is to actuality). Where this happens, direct appeal to the imagination

of the possibility in question will get us nowhere, for that is in part where the disagreement lies; and where this imaginative resistance is principled, it will not be a simple matter to demonstrate imaginability in the same way that one can demonstrate the imaginability of a chimera by getting someone to imagine first a lion, a goat and a serpent (this brings out the limitation of recombination mentioned at the start of this chapter).

But by appealing to possible evidence, or possible reasoning given a certain possible scenario, we can avoid begging the question of imaginability. The move I am offering is similar in spirit to the move made by Adams in his argument for the possibility of distinct indiscernibles: in that dialectical setting, it was hard to see how imaginability of two qualitatively indiscernible objects (i.e. Black's spheres) could lead to that possibility without begging some questions; hence Adams appealed to a possibility that both sides could agree on, and proceeded from there. Similarly, in order to make progress in a dialectical setting where the imaginability of something is part of what is disagreed on (and, again, as Hanrahan noted, "two people cannot share and evaluate each other's imagined experiences" (2007, 145)), we can appeal to possible (perhaps imagined) evidence, and possible reasoning given that evidence. If we can think of a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then we should conclude that *p* is possible. As in Hanrahan's account, such a move means proceeding from the epistemic to the metaphysical, from what would be most reasonable to conclude given certain possible facts, to what is possible. The risk of begging the question is therefore dramatically reduced, for reasonable belief does not entail the truth of what is believed.

Instead, the inference from possible evidence to possibility is ampliative, which is suggestive of a further analogy to build on the analogy between imagination and perception: just as, for a

scientific realist, ampliative inferences can warrant belief in unobservables, possible ampliative inferences can warrant belief in sensorially-unimaginables. Insofar as these possible inferences are abductive, my account shares a further feature with Hanrahan's: its viability depends on the viability of explanationism, the view that inference to the best explanation is a source of justification (Hanrahan 2007, 130). Like Hanrahan, I'll pass over the task of defending explanationism, too great an undertaking to face here; the same goes for scientific realism, in particular the idea that inference to the best explanation can warrant belief in unobservables. On this last point, it is perhaps as well to note that scientific realism lies at the foundation of my broader project, not just of justifying belief in the possibility of sensorial-unimaginables, but also of taking contingent evidence in general, and science more particularly, to be a guide to metaphysics. As Hawley has argued, the view that science can be a guide to metaphysics is best served by scientific realism (Hawley 2006, sections 4 and 5).

There are also some conditions that must be imposed on the epistemic situations involved in the inference from possible evidence to possibility. One is that were we in these possible situations, we would not be less informed on the subject matter than we already are; another is that it is not part of the possible situations in question that the capacity to reason is diminished, e.g. by an interfering demon. Similarly, given that Hanrahan seeks to infer possibility based on what her epistemic twin infers, it makes sense that part of her argument hinges on the claim that her epistemic twin is a reliable witness: that the twin reasoned well and that she "was not impaired either physically, mentally, or intellectually at the time in question" (Hanrahan 2007, 139).

A further point of analogy between Hanrahan and my own route to possibility claims is that we both assume some possibilities in order to infer others. With this in mind, Hanrahan considers the following objection:

I want to justify my belief that B is possibly true. In order to provide myself with this justification, I have made an assumption that I have tried to defend. But in this defence, I have employed some of my beliefs about what is possible. For example, my above argument depends on the claim that the images we produce via our imagination could have come to us via an alternative pathway. So, I am seemingly trapped in a circle. (Hanrahan 2007, 141)

A similar objection might be put to me: in any argument from possible evidence to a possibility claim, the possibility of at least some of that evidence must be assumed. For example, in Shoemaker's case of evidence for the possibility of time without change, we must assume that his world of partial freezes (apart from whether it is a world where every sixtieth year is one of total freeze) is itself possible. However, the objection poses more of a problem for Hanrahan than for me, for it is not my aim to provide a general route to possibility claims to rival all others. Rather, I grant that there are other routes to possibility claims. It is not my aim to justify the link from imagination to possibility, whereas this is Hanrahan's aim.

Hanrahan responds that the circle is not a vicious one.

When we seek to establish that a proposition is possibly true, we need not start from scratch, throwing out all of our modal beliefs. We only need to avoid assuming the possibility of the proposition whose modal status is in question. This I have done. (Hanrahan 2007, 141-142)

The trouble is that Hanrahan goes on to tentatively endorse a “limited skepticism about possibility”, based on the limitations of her own modal epistemology, namely that “[b]y design, it is radically ill-equipped to deal with possible worlds that are radically different from this [i.e. the actual] world” (Hanrahan 2007, 143). Hanrahan’s epistemology is intended as a general alternative to Chalmers’s and Bealer’s, and the suggestion seems to be that we should be sceptical of modal claims that cannot be justified by Hanrahan’s method. But it seems that Hanrahan herself appeals to such modal claims – for example, the claim that the images we produce via our imagination could have come to us via an alternative pathway. One might think this is justified by everyday experience; but then why not say the same of other mundane modal claims like the one that Hanrahan justifies using her method (that it is possible for there to be a bear in her backyard)? The kinds of possibilities that Hanrahan’s epistemology can justify are mundane, the possible inferences she refers to based on “variations in our immediate surroundings” (Hanrahan 2007, 143). Thus the modal claims she takes for granted are on a par, as far as the need for justification is concerned, with those she seeks to justify.

I have claimed that consideration of possible evidence can provide a route to possibility claims where appeals to imagination or conceivability are inappropriate, e.g. where imaginability or conceivability of a possibility is part of what is disputed. I, too, rely on the analogy: the imagination is to the possible as perception is to the actual. But taking that seriously encourages us to endorse the Dependency Thesis, which poses difficulties for certain claims to have imagined certain things. This is not to say that only the imaginable/conceivable are possible (just as, from a scientific realist perspective, the unobservability of something does not entail its non-existence); it just means that appeals to the imagination are rendered ineffective in such discussions. I don’t present my route to

possibility claims as a *general* alternative, to supplant all other routes to possibility claims. But in disputes where imaginative resistance is a factor, and dialectical settings where the problem of adjudication arises, appeal to possible evidence can make some headway where appeals to imagination or modal intuition of the possibility in question are inappropriate. If I make use of imagination to justify the possibility of some evidential situation, I do not assume what is to be shown, for I do not aim to justify the claim that imagination is a guide to possibility; thus I can, without circularity, appeal to imagination or modal intuition to justify modal claims regarding possible evidence, in order to consider what conclusions it would be reasonable to infer. In terms of the kind of dialectic I have been concerned with in the present chapter, this makes sense, for those who take imaginative resistance to be a barrier to justification of possibility have effectively granted imagination as a guide; that is, the idea that imaginative resistance is a barrier to the justification of a modal claim presupposes that imaginability *would have* provided some justification.

10. The Possible Evidence Principle, and Potential Objections.

10.1. Figuring things out about possible evidence; figuring out what is possible. One might think that any argument that proceeds as Shoemaker's does – inferring possibly-p from an appraisal of merely possible evidence in favour of p – makes use of an objectionable conception of possible worlds, a conception to which Kripke famously objected: "Possible worlds' are *stipulated*, not *discovered* by powerful telescopes" (Kripke 1981, 45); "Generally, things aren't 'found out' about a counterfactual situation, they are stipulated" (Kripke 1981, 49). But in Shoemaker's case, it seems that we are supposed to be figuring something out: we are figuring out what would be most reasonable to conclude given the evidence possessed by Shoemaker's imaginary inhabitants; and in doing so – according to the principle presently under consideration – we are *finding out* what possible worlds there are.

Kripke's concerns in the quoted statements above seem to be primarily about reference and identity, but his characterisation of possible worlds – how we come about them – is reasonably popular, and many may see it as underscoring a certain oddness about this way of proceeding. Indeed, this oddness is what makes the principle an interesting one: if it can be defended, it would seem to represent a novel way of theorising about what is metaphysically possible. It may be that there is not much to be said to defend it from Kripke-style insistence on possible-worlds-as-stipulated – it is simply a feature of the principle I am now considering that, sometimes, possible worlds are found out, not simply stipulated. A couple of remarks are worth making in this connection, however.

First, it is often the case in metaphysics that stipulation is inappropriate, in particular when doing so involves asserting what is to be shown (unless of course the stipulation is provisional, for example as part of a *reductio*); or when it involves assuming the falsity of a key claim motivating the opponent's view (as when the possibility of distinct indiscernibles is argued for by stipulating distinct locations of duplicate objects; for the relationist about space, distinctness of objects grounds distinctness of regions, not vice versa (Hacking 1975, 251-252)). As I noted above, Shoemaker's is indeed such a case: it would have been dialectically inappropriate for Shoemaker to ask his opponents to consider a possible world in which time passes without change; similarly, Shoemaker could only address those who claim that it is impossible to *imagine* such a scenario – as advocates of the Dependency Thesis will – by doing something *other* than asking that we imagine time passing in the absence of change. In such cases, appealing to possible evidence seems to be appropriate.

Secondly, there is an interesting grey area here, due to the fact that inductive inferences play a crucial role in Shoemaker's thought experiment. Suppose that Shoemaker is right that it is objectively more reasonable for the inhabitants to conclude that the sixtieth year is one of total freeze, and moreover some principle needed to underscore our moving from this to the claim that *time without change is possible* is correct: by considering what should be concluded given a certain possible evidential situation, we can draw conclusions about what is possible. Everything that Shoemaker asks us to imagine – everything that is stipulated about his scenario – is compatible with the imagined world being one in which time *never* passes without change. As such, it seems a little strange to say that in considering what would be most reasonable for the inhabitants to conclude, we are *figuring out* what is true at *that particular* possible world. This is strange because there are countless worlds conforming to Shoemaker's stipulations (varying in trivial respects – say, the number of inhabitants in the world, the kinds of clothes they wear, etc.), and they cannot all be worlds in which time passes without change – we have to allow that there is some such possible world in which a total freeze never occurs, even though everything happens there as Shoemaker describes. And yet, since all such worlds are the same in respects relevant to the question of whether it is reasonable to conclude that time passes without change, it seems natural to think that the conclusion drawn about one should be drawn about all. I will return to this issue in the next chapter.

10.2. Possible evidence, 'water' and H₂O.

There is a widespread view that water is essentially, and therefore necessarily, H₂O. We may find ourselves able to imagine a world where some watery stuff, comprising elements XYZ (= something other than H₂O), is referred to using the term 'water', is clear, drinkable, usable for washing, etc.; but it is impossible for water to be anything other than H₂O. That's

because, on this view, 'water= H_2O ' is an a posteriori necessity, albeit arrived at by a combination of a priori and a posteriori means. We start with an a priori claim: 'water' rigidly designates whatever plays the water-role in the actual world. Then our empirical investigation reveals that H_2O plays that role in the actual world. In other words, water is essentially H_2O .

This view is widely held among philosophers (though not universally – see, e.g., Sidelle 2002), and may be thought to present a problem for the claim that possible evidence is a guide to possibility. I am defending the view that, if there is a possible situation in which it would be objectively reasonable to conclude that p, then p is possibly true. Isn't there a possible situation in which it would be objectively reasonable to conclude that water=XYZ (not H_2O)? If so, my principle may seem to go against the widespread view that water is necessarily H_2O .

There is, however, a straightforward reply to this worry (a reply which echoes Chalmers's response to the claim that we can conceive of impossibilities, as given in Hanrahan 2009, footnote 4). In the argument for the claim that water is necessarily H_2O , we should note that 'water' can be understood in two ways: in its primary or in its secondary intension. The primary intension of a term provides the way in which its referent is fixed at the world at which it is used. In the present case, we can consider the primary intension of 'water' to be 'that which plays the watery role' (Chalmers 1996, 57). In the actual world, that which plays the watery role is H_2O , but there are plausibly worlds in which that role is played by XYZ, something other than H_2O . Understood in its primary intension, then, 'water is XYZ' is possibly true. The secondary intension of water, on the other hand, is what the term rigidly

designates: in the case of the actual world, taken in its secondary intension, 'water' rigidly designates H₂O, and so refers to H₂O whichever world we are considering.

Now consider my principle with 'water is XYZ' plugged in. If there is a possible world in which the observable facts make it objectively reasonable to conclude that water is XYZ, then we should conclude that 'water is XYZ' is possibly true. We now have to ask, is 'water' being taken here in its primary or its secondary intension? Suppose it is the primary intension. In that case my principle does tell us that we should conclude that water is possibly XYZ, and hence not necessarily H₂O. However, this is unproblematic, as it is only when given in its secondary intension that 'water' is supposed to refer to a kind of stuff that is necessarily H₂O. Taking 'water' in its secondary intension, however, the antecedent of the conditional is false – there is no world in which the observable facts make it objectively reasonable to conclude that water (=H₂O) is XYZ (≠H₂O); one could only reasonably conclude as much if one were lacking some crucial empirical knowledge of what water is.

10.3. An objection from possible justification by testimony.

We have been focussing in this chapter on justification of possibility claims by way of a consideration of possible evidence, but in previous chapters I also suggested that consideration of what would be reasonable to conclude in various worlds, can also lead us to the contingency of certain metaphysical theses where the evidence pertaining to a metaphysical theory varies across modal space. That is, contingency falls naturally out of my modal principle when there are worlds in which it would be objectively reasonable to conclude that p, and worlds in which it would be objectively reasonable to conclude that q, where p and q are incompatible or rival claims.

This leads us to another potential problem for the view, based in the idea that testimony is a source of evidence. Take, for example, Goldbach's conjecture, the claim that every even integer greater than 2 is the sum of two primes. We can imagine being justified in believing its truth, and we can imagine being justified in believing its falsity, by imagining a scenario in which a very reliable mathematician claims to have proven its truth, and imagining a similar scenario in which a very reliable mathematician claims to have proven its falsity.⁶³ Or, to make the case stronger, replace the reliable mathematician with the entire community of mathematicians. Why can't my principle be used to justify the contingency of Goldbach's conjecture in this way? This is a worry about over-generating possibilities, as Goldbach's conjecture is presumably necessarily true, if true at all.

Here, again, the phrase '*objectively* reasonable' is doing a lot of work in my principle. If it is correct that Goldbach's conjecture is a matter of necessity, then only one set of possible mathematicians will be correct. As such, one set of mathematicians must have made an incorrect step in their supposed proof, in principle discoverable by all, and hence it is not *objectively* reasonable to affirm their conclusion. It may be that, as far as the possible mathematicians have seen, there is no such incorrect step; and as far as any layperson is concerned, it is only reasonable to believe what an overwhelming number of respected mathematicians are saying (after all, I have never proven Fermat's last theorem, yet I believe it to be true because I have seen it reported that it has been proven). But this only leaves us with evidence from the limited point of view, either of the mathematicians who have failed to

⁶³ Thanks to Aaron Cotnoir for suggesting a case along these lines. A similar case is given in Yablo (1993): "Thus I can imagine some leading number theorist announcing an error in Euclid's proof from which it emerges that there is a largest prime number after all; the error takes years of training to understand but the authorities are convinced, and I, naturally, defer to their superior knowledge. Although my imagined self is convinced, my actual self is not; I find a largest prime unimaginable and so I suppose the imagined authorities to be mistaken" (Yablo 1993, footnote 46).

notice the error in their proof attempt, or of the layperson who has not seen the proof for themselves (and would not be able to evaluate it, even if they had).

To put this another way, in the Goldbach's conjecture case, presumably it is either true or false that the possible mathematician in question has managed to prove it. If she has, then the antecedent of my principle is satisfied; if not, it is not, testimony only makes it reasonable to conclude that she has proven Goldbach's conjecture from the standpoint of someone who has not correctly gone through the steps of the proof. We are justified, by my principle, in saying that Goldbach's conjecture is possibly true only in the former case, for only then would it be objectively reasonable to conclude that Goldbach's conjecture is true. If the mathematician has not managed to prove it, but claims to have done so, then it is at best subjectively reasonable to conclude that Goldbach's conjecture is true on the basis of her testimony. But whether she has or not is not something we can simply stipulate either way, for proof is a success term, and hence it is only possible for the mathematician to prove it if it is indeed provable; we can't decide whether it is provable without deciding whether it is possibly true. Any case where the possibility in question must be decided first in order to apply my principle will of course be one for which my principle is useless.

For Hanrahan, the problem remains, since she only requires warrant from the perspective of her epistemic twin. Presumably Hanrahan does not know whether Goldbach's conjecture is true, and so neither does her epistemic twin. But *from the subjective point of view of a layperson*, it does seem reasonable to believe that a mathematical conjecture is correct if the whole community of mathematicians is agreed on its proof – that is why I believe that Fermat's conjecture is correct. Hence Hanrahan's epistemic twin is justified in believing Fermat's conjecture; but Hanrahan also has an epistemic twin who justified in believing that

Fermat's conjecture has been refuted. So if possible subjective justification of an epistemic twin, who is generally reliable for believing p , provides justification for one to believe possibly- p , then Hanrahan's modal epistemology would seem to justify both that Goldbach's conjecture is possibly correct, and that it is possibly false. Hence it seems that any principle that warrants the inference of a possibility from possible evidence must somehow exclude merely subjective possible justification.

This response, however, raises a further rather pressing issue. I have assumed that Goldbach's conjecture is a matter of necessity. Based on that assumption, I claimed that only one set of mathematicians in our possible scenarios can be correct in their claim to have proven or disproven it, and hence that it can only be objectively reasonable to conclude that it is true, or to conclude that it is false, but not both. We now have to consider whether this same line of reasoning can be adopted by any necessitarian, to defend the necessity of some matter of metaphysics in spite of variation across modal space in the evidence that apparently pertains to that matter.

10.4. Objection that the principle lacks probative force against necessitarians about laws.

While far from a universal opinion, it is often supposed that worlds with differing laws of nature are possible. This is especially natural for a scientific realist, who thinks that inference to the best explanation is a valid mode of inference. Presumably (on this view) there was a time when inference to the best explanation warranted the belief that the actual world is Newtonian. It turned out, however, that some Newtonian predictions were inaccurate or false. For example, one prediction of Newtonian physics is that the speed of light, as it makes its way through the ether, will vary depending on the movement of the observer: faster, if the observer is traveling towards it, slower, if the observer is traveling away from it.

However, the measurements of Michelson and Morley showed this not to be the case: whether the Earth is traveling towards or away from the Sun, the Sun's rays the speed of light was found to be the same (Dainton 2010, 316).

But what if the Michelson-Morley experiments had different results, corroborating rather than refuting the Newtonian prediction? More to the point, what if *all* Newtonian predictions withstood rigorous testing? Plausibly, the possibility of this happening lies behind the thought that a Newtonian world is possible – since there is a possible world in which it would be reasonable to believe that the world is Newtonian, we should think that a Newtonian world is possible. In effect, my principle may be behind the claim that the laws of nature are contingent.

However, it is difficult to see how my principle can be used to argue for the contingency of the laws of nature against dispositional essentialists such as Alexander Bird (2007).

Borrowing from Kripke, Bird deploys a re-description strategy to explain away the apparent evidence from imagination that other laws are possible (Bird 2007, Chapter 8). In short, the idea is that, when we think we are imagining a world in which, say, an airplane travels faster than the speed of light, we are actually imagining a world in which an airplane travels faster than something else, call it *tlight*. More generally, when we think we are imagining a world in which some other, non-actual laws of nature are correct, we are in fact imagining alien worlds, not accessible from the actual world. There may therefore be no use in appealing to possible evidence to refute the dispositional essentialist claim that the laws of nature are necessary, as the dispositional essentialist will employ an error theory, re-describing what is being imagined much as Kripke would have us re-describe apparent imaginings of water being XYZ as actually being imaginings of a qualitatively indistinguishable scenario.

This does not however render a principle sanctioning appeals to contingency of evidence useless for the purpose of establishing contingency in some matter of metaphysics. Appeals to contingency of empirical evidence to support the contingency of metaphysical theses whose plausibility is affected by that evidence are of course useless to those who do not accept the former contingency. But this does not mean that arguments appealing to contingent evidence can only preach to the converted: there may be theorists who are not necessitarians about the laws of nature, who are nonetheless necessitarians about many matters of metaphysics. By showing that a principle lying behind their commitment to the contingency of the laws of nature also supports the contingency of metaphysical theses whose plausibility varies across modal space, as the evidence varies across worlds, a substantive argument has been made. In short, an argument that proceeds from imagination of a world that violates the laws of nature might not move a necessitarian, but can still have some force for the rest of us, for such an argument might show that there is more contingency surrounding the kinds of things we are interested in than we initially thought⁶⁴.

How about Shoemaker's case? Can an argument like Shoemaker's convince someone who has an antecedent commitment to the necessary co-occurrence of time and change? Perhaps such theorists might take a similar line to the one I took in the previous subsection. Based on the assumption that time without change is impossible, they might claim that it *cannot* be *objectively* reasonable to conclude that time passes without change, just as it cannot be objectively reasonable to conclude that Goldbach's conjecture is true if in fact it is possible to disprove it.

⁶⁴ I am thinking here of metaphysical issues as characterised by Hawley (2006), so as not to preclude by an act of definition that science can be a guide to the nature of time, space, properties, causation, persistence, numbers, composition of objects and possible worlds (Hawley 2006, Section 1). See my introduction.

Whether one takes this line will depend on the strength of one's commitment to the impossibility of time without change. We should not imagine Shoemaker's inhabitants as thinking through what to conclude given their abductive evidence isolated from any consideration of the strengths and weaknesses of the relevant metaphysical theories; rather, we can stipulate that Shoemaker's inhabitants are well-versed in the reasons one might have for thinking that time and change necessarily coincide (they have read, for example, Le Poidevin 2003, Chapter 2, and chased up the references). When thinking about what the observable facts make reasonable to conclude, we, too, should take into account the reasons one might have for denying the possibility of time without change. The question is whether these reasons are strong enough to overturn the simplicity of Shoemaker's preferred account of the inductive data. One can argue that they *are* strong enough, of course; but, as I noted with regard to Warmbrod's claim that a different interpretation of the inductive data is preferable, such an argument seems to presuppose that, *had* the observable facts been such as to make it objectively reasonable to conclude that time passes without change at Shoemaker's world, then we should conclude that time without change is possible. To deny the antecedent of some instance of PEP is not to deny the principle itself. Perhaps one will not grant that Shoemaker's argument is in fact a successful argument from possible evidence to a possibility claim. In that case, some further examples of PEP at work may serve to illustrate its uses.

11. Further examples of the principle at work.

11.1. Universe growing in size.

Structurally, this argument is almost perfectly analogous to Shoemaker's argument for the possibility of time without change. There is a world comprised of three inhabited regions

which exhaust this world: A, B and C. Region A grows in size by 1 unit every three years; region B by 1 unit every four years; region C by 1 unit every 5 years. Given the observed regularity of these changes in size, the inhabitants of this world infer the same pattern will continue into the future, and hence that the sixtieth year will be one in which the whole world grows by 1 unit (for all and only multiples of sixty are multiples of all three numbers, 3, 4 and 5 (Oppy 2004, section 1.2)⁶⁵. As with time without change, the global growth cannot be directly observed, for since all regions would grow together by an equal amount, there is nothing in the world that could be held up to measure the growth. That's not yet to say that the global growth is non-relational, for, on the assumption that a global growth takes place, the whole world will be a greater size *that it was before*; hence the inference to a global growth does not directly contradict relationism about size. It does however mean that there could be changes in size that are not observable.^{66,67} It thus seems that relationism about size and the view that there can be no concrete differences that are not in principle observable (a view which provides some motivation for relationism) come apart.

11.2. *Earth and the Aleph.*

⁶⁵ Oppy's (2004) argument proceeds from local doublings to a global doubling, but the numbers don't add up in such a way that would make for a sixtieth year where all three regions grow by the same amount. For consider that, since region A is meant to double every three years, by the sixtieth year it would have doubled 20 times; but region B would have only doubled 15 times in that same period, and region C 12 times. Having the regions all grow by one unit each time, instead of doubling, remedies this by ensuring that each would grow by the same amount as the other in the sixtieth year (even if the different regularities of growths in the different regions is such as to render them all different sizes as time progresses).

⁶⁶ This assumes Euclidean space (Oppy 2004, footnote 4; Dainton 2010, Chapter 14.1).

⁶⁷ In my exposition I referred to the local *doubling* in size of the regions, and this seems to suggest that they are *growing* with respect to some others. But a region can become double the size of another in virtue of the *shrinking* of the other. Indeed, for a relationist about size, these are different descriptions of the same occurrence. Thus the relationist can describe the local doublings in size in two ways: for example, take the twelfth year, in which regions A and B both double in size relative to region C. Observing this, the inhabitants might say that A and B have grown to double the size of region C; but they might also say that region C has shrunk to half the size of A and B. For the relationist, the difference is only at the level of description. Therefore, the relationist must separate the concept of doubling from that of growing, and that of halving from that of shrinking.

Josh Parsons offers an argument from possible evidence to suggest that it is possible for an object to be a proper part of itself, based on a variation of J.L. Borges's story "The Aleph" (2000; originally published in 1949). In Borges's story, an object, the Aleph, in Beatriz Viterbo's house, is described as containing the whole universe; which is to say, it is a part of itself.

Suppose that, prior to the events of Borges' story, Beatriz Viterbo discovered the Aleph for herself. Looking at her cellar staircase through a powerful electron microscope, she finds a tiny apparent replica of the Earth among the microscopic constituents of the 19th step. Looking closer, she sees an even tinier apparent replica of herself in her house, looking into a tiny electron microscope. Zooming out, she sees that the tiny Earth is surrounded by a Solar System, and indeed by a tiny apparent replica of the entire known universe, all forming a microscopic part of the 19th step. Looking into the sky with a powerful radio telescope, she sees that what she thought was the universe is in fact a part of a gigantic intergalactic replica of her cellar, as seen from the 19th step. Being a parsimonious reasoner, she concludes that the apparent replicas are one and the same thing – that the Earth she found in her staircase is the Earth she lives on; that the gigantic step in the sky is the one in her house. (Parsons 2013, 1)

Parsons ultimately uses this story as a springboard for a mereology according to which the proper part relation is not transitive. I'm not concerned here with his further arguments for that claim; but it is worth noting that Parsons is explicitly (Parsons 2013, footnote 6) making a Shoemaker-style argument. Parsons goes on to say:

I am not just asking you to consult your intuitions to decide that a work of literature is internally consistent. The story also describes a way in which a clearly consistent set of evidence could favour a scenario in which there is a mereological circularity. But scenarios cannot be favoured by evidence unless they are consistent. (Parsons 2013, 2-3)

These last remarks are acknowledgment that appeal to possible evidence can be more powerful than appeal to intuitions, particularly given that different theorists may hold contrary intuitions. An argument along these lines is also the best bet for a conclusion such as the one drawn by Viterbo; conjuring up the imagery given in the story is not enough to determine whether what is being imagined is a mereological circularity, or a Russian doll universe containing parts that, while qualitatively exactly similar, are distinct. That is, as we saw above, such imagery can underdetermine the world imagined, giving rise to the problem of adjudication. Where intuitions among different theorists differ, Viterbo's appeal to parsimony can be brought in to help adjudicate.

11.3. The Turing test. Turing (1950) famously proposed a test for the presence of mind, thought or intelligence in machines.

I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 , to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning. ... I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.
(Turing 1950, 442)

Putting aside the correctness of Turing's predictions, and the demandingness of the test he proposes, he is effectively presupposing something like my modal principle: given that a computer could convince us – the average interrogator – that it is human more often than not, we should conclude that it is possible for machines to think. While Turing is perhaps more

concerned about what *will* be the case, than he is about what *could* be the case, there seems to be a more general principle underlying his test: if there is a possible world in which the observable facts make it objectively reasonable to believe/conclude that p, then we should grant that p is possible.

12. Concluding remarks.

Appeal to possible evidence can help make progress on a debate concerning a certain possibility claim where differing claims about what can be imagined would have left the dialogue at a stalemate. Where one side in the debate has a principled reason for why a certain scenario is beyond justification by imagination (as I suggested would be the case for time without change if the Dependency Thesis is correct), then imagination of course cannot be appealed to in order to justify the possibility of that scenario. By illustrating how PEP might be appealed to, my aim has been to show that possible evidence can *in principle* provide one with justification for a claim that lies beyond possible experience and beyond sensorial imagination. In this sense, PEP can be seen as expanding a scientific realist stance to the modal realm, with potentially surprising results.

I hope that the cases provided also show how one might go about justifying a possibility claim using PEP. Of course, it is open to anyone to disagree on what the observable facts in a given world make reasonable to conclude (perhaps by disputing one theory's claim to be the most virtuous explanation of the possible phenomena, just as one might reject an actual inference to the (purported) best explanation, by pointing to the virtues of an alternative explanation). This was Warmbrod's (2004) tactic, to rebut Shoemaker's argument for the possibility of time without change. But, as I suggested in my discussion of Warmbrod's

critique, such a tactic seems to grant the general capacity of possible evidence to provide evidence of possibility.

However, the reader will have noticed that I have only provided a route to possibility claims.

In the next chapter, I will explore the capacity for a consideration of possible evidence to lead to necessity claims.

5. Abduction, Agglomeration, A Priority, and Necessity.

1. Introduction.

This chapter looks at how the principle PEP discussed in the previous chapter creates what Faraci (2013) has called a “modal echo of the lottery paradox”. That principle effectively says that, if there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then we should conclude that *p* is possible. In keeping with this principle, we should say that, for every world in which it would be objectively reasonable to conclude that *p*, we should say that *p* obtains in that world. But if that’s true, then it is natural to think that *p* obtains in every such world. Applying similar reasoning to a priori justification, we may end up concluding that all a priori claims should be necessity claims – at least to the extent that what is a priori reasonable to conclude at one world is equally reasonable at any world. I’ll suggest, however, that we should follow a prominent solution to the lottery paradox, and reject the agglomeration step that this argument from a priority to necessity would require. I’ll end the chapter with a discussion of Biggs’s (2011) case for the claim that modality is centrally abductive, in which case both possibility and necessity claims can be justified by abduction.

2. Modal echoes of the lottery paradox.

Let’s continue to suppose that it is objectively reasonable for any inhabitant of a world conforming to Shoemaker’s stipulations to believe that every sixtieth year is one in which time passes without change. In that case, by the principle we have been considering, PEP, we should believe that it is possible for time to pass without change. There are very many possible worlds conforming to Shoemaker’s conditions: some in which all the inhabitants

wear only blue socks, some in which the socks are red, and so on for all kinds of irrelevant but world-distinguishing features. Call these “Shoemaker worlds”.

Now take some particular such world, w . It is objectively reasonable – let us suppose – for the inhabitants of w to believe that every sixtieth year in w is one in which time passes without change. But what should *we* believe about w ? PEP tells us that, if there is a possible world in which the observable facts make it objectively reasonable to believe p , then *we* should believe that p is possible. It is hard to see why possible evidence gives us reason to believe that p is possible unless we have reason to believe that p is true at w – unless, that is, our believing that it is objectively reasonable for the inhabitants of w to believe that p gives *us* reason to believe that p is true *at* w .

However, if our beliefs about w ought to reflect what we think the inhabitants ought to believe about it – that w is a world in which time passes without change – we soon find ourselves in a paradox, one curiously reminiscent of Kyburg’s (1961) lottery paradox. Think of a fair lottery in which it is reasonable to believe that one ticket is a winner (perhaps this is reasonable simply because the lottery is carried out like a raffle, in such a way that all tickets are sold, and hence that the winning ticket, chosen at random, must correspond to one of the sold tickets). Suppose also that there is a threshold for reasonable belief: if one considers that the probability of a certain event is above that threshold (say, .70), then it is reasonable to believe that that event has occurred or will occur (hence our assumption that it is reasonable to believe that some ticket is a winner amounts to assuming that the event of there being a winner has a probability greater than that threshold). Now if there is a sufficient number of tickets sold, the probability of each being a loser will fall well above the threshold for reasonable belief. Thus it will be reasonable to believe of each ticket that it is a loser. But by

conjoining (agglomerating) reasonable to beliefs about each ticket (that it is a loser), we end up with a claim about all tickets: every ticket is a loser. This contradicts our first assumption, however, that it is reasonable to believe that one ticket is a winner. That's the lottery paradox, in a nutshell.

If we allow inferences to possibility claims in line with PEP, as discussed in the previous chapter, something similar can be said to arise in the modal realm. While Shoemaker-worlds may differ from each other in any number of respects, each has its membership in that class by conforming to Shoemaker's stipulations, which pertain to what evidence obtains. Hence, when it comes to what they should believe about the existence of changeless intervals, what holds for one such world holds for all: in every such world, the inhabitants should believe that time passes without change. But the inhabitants of w know no more nor less of the relevant facts about their world than we do; as such, one might think that we should likewise believe that w is a world in which time passes without change – and likewise, for each of the other worlds in the class. But, if we should say of *each* of these worlds that time passes without change in that world, then, surely, we should say that they are *all* worlds in which time passes without change.

But that seems like something we *shouldn't* believe. After all, as we noted in the previous chapter, the evidence available to the inhabitants of Shoemaker's scenario is compatible with every sixtieth year being an exception to the perceived regularity of local freezes. The inhabitants may have strong reason to believe that time passes without change, but their evidence does not *entail* this; surely, there are plenty of worlds conforming to Shoemaker's stipulations in the relevant respects, but in which time never passes in the absence of change.

We have, then, the following claims:

- a*) At every Shoemaker-world, it is objectively reasonable for the inhabitants to conclude that time passes without change.
- b*) For every Shoemaker-world, if it is objectively reasonable for the inhabitants to believe that time passes without change at that world, then we should believe that time passes without change at that world.
- c*) At every Shoemaker-world, the evidence available to the inhabitants is compatible with a total lack of timeless intervals.

From a*) and b*), it seems natural to infer that all Shoemaker-worlds are such that there are timeless intervals. But from c*), it seems natural to infer that some Shoemaker-worlds are totally lacking timeless intervals. (Though additional premises are needed to make these inferences valid.)

There are hints here of a more general issue, and indeed something very similar has recently been noticed to arise out of a certain metaethical discussion. Faraci (2013) shows how a similar problem arises out of Brown's (2013) discussion of moral error theory, in particular from what Brown says about the *argument from relativity*, as originally offered by John Mackie (1977). The details of the argument from relativity needn't concern us here; what is important is that it is an abductive argument for error theory: error theory provides the *best explanation* for the diversity in moral opinions and practices that we find in our world, in different places and across time. Brown (2013) argues that, if that is correct, then the argument from relativity is not merely an argument in favour of error theory for the actual world, but also for *all* relevantly similar worlds.

In itself, this doesn't seem particularly problematic. At least, for one who would argue for error theory at the actual world, the stronger modal conclusion may be welcomed – so much the worse for moral realism. The problem Faraci identifies, however, has more to do with the general form of Brown's argument. Faraci's discussion focuses on the case of abduction, but he notes that the issue he raises with the form of Brown's argument can be generalised to apply to all forms of non-deductive reasoning. Here is the argument form at its most general:

1. From set of features S of W^A , one has license to conclude (through non-deductive reasoning) that C in W^A .
2. For any world W^* , if W^* has features S , one is warranted in concluding that C in W^* .
3. Therefore, if one is warranted in concluding that C in W^A on the basis of S , one is warranted in concluding that C is true in any world with features S . (Faraci 2013, 753, footnote 5).

In Brown's case, S is any set of features that amount to divergence in moral opinions and practices, W^A can be considered to be the actual world (though any world relevantly similar to the actual world would do the job), and C is the truth of error theory. In order to establish his conclusion that all worlds with similar divergence in moral opinions and practices are worlds in which error theory is true, Brown needs to be able to infer the following from (3.) (where W is the set of worlds relevantly similar to W^A):

4. Therefore, [if one is warranted in concluding that C in W^A on the basis of S], one may conclude that C is necessarily true across set W_+ , which includes W^A and all worlds in set W .⁶⁸ (Faraci 2013, 753).

⁶⁸ My modification in square brackets is to make (4) cohere with the argument form as generalised to apply to all forms of non-deductive reasoning.

What's wrong with this argument form? Faraci approaches the problem by first showing there are cases in which the argument form would lead us from true premises to an obviously false conclusion.

Consider a recently estranged couple. It may well be that the best explanation for my seeing them walking hand-in-hand is that they have reconciled. And it may well be that, having recognized this fact, I have license to believe that they have reconciled. But surely it would be extreme for me to conclude that they have reconciled in all possible worlds like this one. Indeed, this further conclusion seems not merely unjustified, but obviously false—after all, one of the defining features of abduction as a form of non-deductive reasoning is that its conclusions are not necessitated by its premises, that the couple's reconciliation is the best possible explanation of what I saw but not the only one.^[69] (Faraci 2013, 752)

What is interesting about the argument form, however, is that, in spite of the possibility of such invalid instances of it, it is hard to say exactly where it goes wrong.

Faraci considers two suggestions that would find fault in the form of reasoning that underlies the last quoted passage by imposing limits on the scope of abduction. (Since the problem Faraci raises can be generalised to all forms of non-deductive inference, analogues of these suggestions would need to be given for the other forms of non-deductive arguments.) The first is to say that abduction can only legitimately lead us to conclusions *at* the actual world. This is quickly dismissed, however, since it would mean that actual abductive conclusions, such as that the recently estranged couple mentioned above have reconciled, would not have been warranted, had the situation differed in some irrelevant respect (e.g. had I been wearing a differently coloured hat) (Faraci 2013, 753). As such, to claim that abductive inferences are

⁶⁹ This may need to be modified or qualified if it is to fit with Biggs's (2011) claim that metaphysical necessities can be established by abduction (see section 6 below); for, if abduction did justify a necessity claim for being the best explanation of some explanandum, presumably this would be to claim that the best explanation is the only possible one.

only legitimate at the actual world is to render them absurdly sensitive to irrelevant contingencies; in the case of the recently estranged couple, I should be able to say that, had I been wearing a differently coloured hat, my conclusion ought to be the same (that they have reconciled).

The second suggestion is to say that abduction can only legitimately lead us to conclusions *about* the actual world. (Again, bear in mind that, since the problem can be generalised to all forms of non-deductive argument, a corresponding restriction would be required for each of these forms, or else a general one that covers all, such as: no non-deductive form of argument can reach conclusions about other possible worlds.) However, this is an odd move, particularly once we have granted that abduction can reach conclusions *at* other possible worlds, for it is hard to see how we can decide on what the inhabitants of a possible world *ought* to conclude without at the same time lending our own support to the truth-at-that-world of the conclusion that we recommend for them (unless of course we have independent evidence – perhaps a stipulation about the world being imagined – that tells *us* that, while *they* should conclude that *p*, *p* is false at their world). According to this restriction on the use of induction: while I can make assertions about what agents at other possible worlds *should* conclude via abduction, I cannot rightly make any assertions about whether they are correct – unless, that is, the latter assertions of mine are not based in abduction. I can make a claim about what someone in another possible world should conclude, based on what I know of their evidence, but if I am to say anything about whether they are *correct*, I must base my claim on something else, stipulation perhaps. This will not help, however, in those cases – like Shoemaker's – where the aim is to decide on whether something is possible or not, where stipulation would beg the question. Now it is of course part of the present suggestion – that abduction should not be used to draw inferences about other worlds – that abduction cannot

help us in those cases either. However, citing Biggs (2011), Faraci also notes that “[i]t has even been argued that our epistemic access to modal facts is *centrally* abductive” (Faraci 2013, 754).

None of this is decisive against the second proposed restriction on the proper use of abduction (the restriction to drawing conclusions at, but not about other worlds). But it does suggest that more is needed to justify the restriction, beyond the fact that it would relieve us of a paradoxical situation.

Perhaps the proposed restriction might be supported by a general modal epistemology suggested by Kripke’s remarks (mentioned in the previous chapter), according to which we cannot *find out* about possible worlds.⁷⁰ However, at least as just stated, this seems false. For one thing, it is not clear why we should take *stipulation* and *finding out* to be opposed – a degree of stipulation seems to be compatible with finding out something about the world in question. On a view according to which we *can* find out about possible worlds, (alleged) instances of this are bound to involve some stipulation, as is the case with our consideration of Shoemaker’s world(s): first we stipulate that regular freezes are observed by the inhabitants, and then we ask what they should conclude about their world given their observations. For another thing, it is hard to deny that some features of a possible world, or set of possible worlds, can be found out: since knowledge is not closed under (unknown) entailment, we might stipulate some features of a possible world without yet realising the consequences of those stipulations; then, when we come to reflect on the things we have stipulated to be the case at that possible world, we may realise that logic compels us to

⁷⁰ “Possible worlds’ are *stipulated*, not *discovered* by powerful telescopes” (Kripke 1981, 45); “Generally, things aren’t ‘found out’ about a counterfactual situation, they are stipulated” (Kripke 1981, 49).

recognise certain other things to be the case as that world, i.e. we can *find out* about the world in question. Once we have stipulated some features of a possible world, the presence of these features will usually entail the presence of others; but we may of course initially fail to recognise these entailments. Perhaps the proposed modal epistemology might be modified then, saying instead that we can only find out about a possible world via a combination of stipulation and *deduction*. However, this would be to directly rule on the question at issue – that of whether we can use abduction, or any other non-deductive argument, to learn about possible worlds – with a ‘no’ verdict; whereas the point of invoking the Kripke-inspired modal epistemology was to provide some justification for that verdict, beyond the fact that it would relieve us of a certain paradoxical situation.

In the lottery paradox, a tension arises from the fact that, on the one hand, the probability of any particular ticket being a winner seems to give reason enough to believe, of any particular ticket, that it won’t win; while, on the other hand, by agglomerating these beliefs, one would end up with the conclusion that *no* ticket will win. This conclusion seems dubious; after all, lotteries are often won, in spite of the extremely low probability facing each individual player, and, besides, we can easily describe the case so that we know that there definitely *is* a winner (perhaps a million or so scratch cards are handed out at random, one of which is the winning ticket).

Given the similarity between the lottery paradox and the problem that arose in relation to non-deductive arguments (as illustrated via Shoemaker’s and Brown’s arguments), one might expect a solution to either one to provide a solution to the other. The solution that Faraci invokes is that it is the agglomeration of beliefs that we should let go of.

The only plausible solution, I submit, is to admit that just as we may believe that each ticket will lose but not that all will, we may believe that error theory is true in each world like ours but not that it is true in all of them. This result is strange, and it is counterintuitive, but without it, Brown's argument could plausibly serve as a premise in a *reductio* of abduction itself, for abductive reasoning – indeed, I believe, non-deductive reasoning of any kind – would lead to contradiction. (Faraci 2013, 755).

Having cast doubt on the reasonableness of placing restrictions on the proper use of abduction (indeed, of restrictions on the proper use of non-deductive reasoning in general), the rejection of belief agglomeration seems to be the only reasonable alternative.

In the next section, I'll consider whether the general structure of the above dialectic can be mapped onto another. First, I'll show that similar considerations that prevented us from restricting abduction to drawing conclusions about the actual world can also shore up the inference from a priority to necessity in the face of an argument made by Ross Cameron (2007). Then, I'll consider whether the move to reject belief agglomeration in response to the problem raised above can bolster Cameron's case against making the inference from a priority to necessity.

3. A priority, necessity, and agglomeration.

Inferences along the lines of PEP can warrant possibility claims, but it is hard to see how a consideration of possible evidence can warrant necessity claims. In particular, if we reject agglomeration, as above, even restricted necessities cannot be justified by a consideration of possible evidence (e.g. we cannot infer that B theory is true at every world in which special relativity is true, even if every world in which special relativity is true is one in which B theory is favoured by the empirical evidence).

However, PEP is restricted to what *observable* facts make it reasonable to conclude, and one might think that metaphysics is largely, or wholly, a subject for a priori reasoning, and as such, one whose questions call for necessity claims as answers.

As discussed in my Introduction, in his 2007 paper “The Contingency of Composition”, Ross Cameron rejects the step from the a priori of an answer to the special composition question to its necessity. This rejection is not based on the purported examples of contingent a priori truths – Cameron says that examples such as Kripke’s ‘Jack the Ripper committed the East End murders (if anyone did)’ (Kripke 1980, 79-80) “turn on a linguistic trick” (Cameron 2007, 107)⁷¹. Instead, Cameron blocks the inference from a priority to necessity on different grounds. The argument from a priority to necessity is based on the claim that the way the world is plays no part in my coming to know a priori propositions: since the way the world is is irrelevant in my coming to know a proposition a priori, no difference across possible worlds could make an a priori known proposition false. Cameron rejects this claim.

Knowledge is factive, so I can only come to know a proposition on the basis of a priori justification *if the world cooperates to make that proposition true*. There is a sense in which the world is irrelevant to a priori propositions, but there is also a sense in which the world is relevant. It is irrelevant in that my justification for an a priori proposition does not come from experience – I do not have to look into the world to find justification for believing it; it is relevant in that, if the world does not cooperate, I may end up with an a priori justified belief in a falsehood. And since justified belief is not factive, while my a priori justification for p may give me sufficient reason to believe that p is true, it does not give me reason to believe that p *must* be true, in the sense of being metaphysically necessary. Hence, according to

⁷¹ See my Introduction, Section 1.2.

Cameron, the argument from a priority to necessity fails because it depends on the false claim that that the way the world is is irrelevant to my coming to know a proposition a priori.

However, we can see why someone might nevertheless maintain a link between a proposition's being justified a priori and its necessity, even while granting that justification does not entail truth. If we truly have sufficient a priori warrant for believing a proposition, then our warrant for believing that proposition is independent from the way the world is. But in that case, surely we should expect that that same degree of justification is available in any world – any world in which conscious reasoning is possible, at least. If our justification is actually sufficient to warrant belief, and *if there is no way the world could be that would diminish that justification*, then the same a priori considerations would suffice for drawing the same conclusion at any world.

So it seems that, where a conclusion is drawn purely on the basis of priori reasoning (assuming that there is such a thing as *pure* a priori reasoning/justification), we should say that, *at any world in which thinking beings are concerned to find an answer to the relevant question, the same conclusion should be drawn*. It is hard to see how we could avoid recommending the same conclusion to (all) inhabitants at (any) other worlds unless we admit either that our own justification for that conclusion is not purely a priori, or that our own justification is not sufficient to warrant belief in that conclusion in the first place.

We must now consider, as Faraci did, whether we can move from what ought to be concluded *at any other possible worlds*, to claiming that we should draw the same conclusion *about any other possible world*. In the case Faraci was concerned with, it seemed reasonable to respond in the affirmative; the question is whether those same considerations – revolving around

abduction in particular, but also applicable to non-deductive reasoning more generally – apply to the present case, where the justification is supposed to be wholly a priori.

Faraci offered basically two considerations against restricting abduction so that it cannot rightly be employed to reach conclusions about other worlds. One was that it would be odd to say: the inhabitants of *w* should conclude that *p* is true at *w*, but I should not (Faraci 2013, 754). That is, what we think the inhabitants of a possible world ought to conclude, on the one hand, and what we think we ought to conclude about that world, on the other, seem to stand or fall together – unless we have some independent evidence concerning the proposition in question which they do not, or vice versa. The other was that there are metaphysicians who have supposed that abduction can be employed to reach conclusions about possible worlds – “such as in arguments for modal realism” (Faraci 2013, 754); while Biggs (2011) has argued that abduction is central (Faraci 2013, 754).

Similar considerations seem to apply no less to the case of a priori justification. It is hard to see how we can accept a conclusion about the actual world on purely a priori grounds unless we are willing to accept that that same conclusion is equally reasonable at every other possible world: *every* world is relevantly similar in terms of the justificatory force of those a priori considerations, for their being a priori suggests that their justificatory force does not depend on the way the world is⁷². Meanwhile, a priori reasoning seems to be in at least as good standing among philosophers as abduction when it comes to justifying modal claims.

⁷² I say “suggests” rather than “entails” because it is open to externalists about justification to say that the justificatory force – in terms of what *should* be believed – of a priori considerations is somewhat diminished in worlds in which the proposition in question is false.

This much *might* help strengthen the inference from a priority to necessity (though perhaps only if a priori justification is considered to immune to defeat by experience); as we'll shortly see, whether it helps will depend on what we say about agglomerating beliefs.

Against inferring necessity from a priority, Cameron pointed out that, while it may be true that a priori propositions can be known on the basis of justification that holds independently of the way the world is, a priori reasoning is fallible: "When we say that a proposition is *a priori*, this means that it can be known on the basis of justification which is not empirical. But justification does not entail truth" (Cameron 2007, 108). But if we believe that our a priori reasons for believing a proposition would warrant our believing that proposition at any possible world, we might also think it natural to say that those reasons warrant our believing that that proposition holds at *every* possible world; that is, we might think that, from being warranted in believing that a proposition is true at each possible world, we are warranted in believing that it is true at *all* possible worlds – i.e. that the proposition is necessary.

4. Should we reject agglomeration of a priori justified beliefs?

As with the cases mentioned above (Shoemaker's argument, Brown's extension of Mackie's argument, the lottery paradox) the move to the necessity claim crucially involves agglomerating (what we believe would be) warranted beliefs. But should we reject agglomeration in the case of a priori warranted beliefs?

One might think that there is a difference here, such that we may agglomerate a priori warranted beliefs (even if we reject agglomeration in the other cases, as above), and hence arrive at their necessity. In the case of the lottery paradox, agglomerating the relevant beliefs

leads to a universal claim which we can stipulate, as part of the case⁷³, to be false. But there is no analogue of this in the above discussion about a priori warranted beliefs – we cannot stipulate that the universal claim (e.g. that every world is a world where van Inwagen’s is the correct answer to the special composition question) is false, for that is precisely what is at issue: whether we should believe that a priori warranted beliefs are necessarily true, i.e. true at every world.

The question is then whether the rejection of belief agglomeration, a rejection which we saw endorsed by Faraci (in response to the lottery paradox and the paradox arising from Brown’s extension of Mackie’s argument), is suitably general such that it gives us reason not to agglomerate (what we believe would be, and therefore *are*) warranted beliefs about each possible world where the warrant is wholly a priori; that is, whether our rejecting agglomeration in the above lottery-style cases gives us reason to reject it in the case of a priori warrant across worlds.

One might think that consideration of the “preface paradox” – a paradox very similar to the lottery one – also provides some support for the rejection of agglomeration in the case of a priori warranted belief. In the preface paradox, the author of a book rationally believes each of the claims she makes in that book; but, regarding herself as fallible, she also rationally believes that the conjunction of those claims is false (or at least, doesn’t believe the conjunction). But if it is a principle of rationality that $B(p)\&B(q) \rightarrow B(p\&q)$, the author

⁷³ Recall that we can construct the case so that the winning number has already been picked out – perhaps in the manner of a raffle of which we know that the tickets have sold out. Before the number of the winning ticket has been revealed, we might consider each ticket and ask: Is it a winner? Provided that the chance of each being a winner is sufficiently small, so that we are warranted in believing of *each* that it is a loser, agglomerating the beliefs about each will result in a universal claim – “Every ticket is a loser” – that conflicts with what we know: that one ticket is (or will be) a winner.

cannot rationally believe each claim without rationally believing their conjunction. In his survey article “Epistemic Paradoxes”, Sorensen notes,

The preface paradox does not rely on a probabilistic acceptance rule [e.g. a rule that “permits belief in any proposition that has a probability of at least .99” (Sorensen 2014, Section 3)]. The preface belief is generated in a qualitative fashion. The author is merely reflecting on his similarity to other authors who are fallible, his own past failing that he subsequently discovered, his imperfection in fact checking, and so on. (Sorensen 2014, Section 4)

Similarly, the metaphysician reasoning about (for example) the special composition question might reflect on her similarity to other metaphysicians who are fallible, perhaps recalling Kant on the modal status of certain claims of geometry.

But the allusion to Kant on geometry is only appropriate if we grant that a priori claims can be, at least in principle, empirically defeated: in that case, we cannot exclude the possibility that some world in which a belief would be a priori warranted is nonetheless a world in which that belief is false. But what if the best characterisation of a priori justification is one according to which a priori justification is empirically infeasible? Such a view is defended in Warenski (2009), for whom a priori justification is essentially empirically infeasible but nonetheless fallible. As Tahko (2008, 59) notes, the fallibility of a priori justification is acknowledged by many modern accounts of the a priori (e.g. Bonjour’s 1998). However, we should also bear in mind that knowledge based on fallible justification need not be knowledge of a contingent truth. Baron Reed (2002) provides a quite general account of fallibilism that “permits fallible knowledge of both contingent truths and necessary truths” (152). If a proposition being a priori justified amounts to it being empirically infeasible and yet

fallible, one cannot base one's rejection of the a priority-necessity link simply on the fallibility of a priori justification.

The strength of the case against inferring necessity from a priority may thus depend on what it is we are talking about when we talk about a priori justification. But if pressed to weigh in on the issue without having settled the essential features of a priority, I would perhaps put the issue on this: It would be ad hoc, and not very informative, if Kyburg's rejection of agglomeration were restricted to those cases where we already know that the universal claim is false. Kyburg's rejection of agglomeration is to apply to uncertain inference, and a priori inference is arguably uncertain – at least where those inferences are not analytic.

5. Contingentism, abduction, and agglomeration.

The problem Faraci raised for allowing abduction to tell us what is true at other worlds should arise wherever contingentism is assumed. Suppose that we come to see that our evidence (whatever that may be) supports mereological universalism at the actual world, and that our evidence does not entail that conclusion, while believing that mereological universalism is only contingently true. Then, of any world relevantly similar to the actual world, we should conclude that mereological universalism is true. But given the assumptions just made, we should also think that there is a possible world relevantly similar to our own, but at which mereological universalism is false. But that would mean that we cannot make the inference, from each world relevantly similar to the actual world being a world in which mereological universalism is true, to the claim that mereological universalism is true at *all* worlds relevantly similar to the actual world. In line with Faraci's discussion, it seems that contingentists again have reason not to agglomerate beliefs, in this case about the metaphysical facts – inferred to via abduction – of other possible worlds.

Consider, for example, contingentism about composition. Under what conditions can we say of some things that they compose an object? For the contingentist, the correct answer to the special composition question varies across modal space: perhaps there are worlds in which nihilism is correct – i.e. the xs never compose an object – as well as worlds in which universalism is correct – i.e. the xs always compose an object. It seems to be a natural feature of contingentism about composition that two worlds could be exact qualitative duplicates, while differing in the facts about composition. Thus most contingentists, it seems fair to say, will deny

(C) Worlds that are intrinsic duplicates are compositional duplicates. (Rasmussen 2014, 537)⁷⁴

Rasmussen goes on to note,

Consider that without the principle, there cannot be any deeper *explanation* of why composition occurs when it does. Compositional facts would be inexplicable. Moreover, if (C) is false, then it seems it would be possible, in principle, for objects to ‘flicker in and out’ of existence without *anything* changing position. So, Peter van Inwagen may be presently correct: there are no tables or chairs. But a second later, Hud Hudson may become right: there are tables, chairs and towernoses – regardless of whether anything changed position or orientation!” (Rasmussen 2014, 537).

Why should we worry about the mere possibility of objects flickering in and out of existence without anything changing position? It seems that Rasmussen is raising two worries here, one about the potential for explanation, and another about scepticism. The explanatory worry

⁷⁴ “The thought here is that if you fix the facts about the existence and arrangement of the basic ingredients of a world, then you thereby fix the world’s compositional facts. Put differently, if two worlds *differ* with respect to what objects compose what, then they differ in some other, non-compositional respect” (Rasmussen 2014, 537).

is that without (C), we cannot explain why composition occurs when it does. One thought that may be behind the worry is that composition can only be explained by a claim that holds necessarily – that the facts that explain why composition occurs when it does must *entail* that composition occurs. Whether or not that is true is an important question if we are interested in working out the details of contingentism. For if some facts can explain composition (or its absence) only when they entail it (or its absence), then surely the contingentist must concede that composition is brute: in a world in which van Inwagen is correct, we cannot explain why composition occurs when it does, except by saying: because van Inwagen is correct at this world. That amounts to an explanation of the form: P, because P.

However, even if we grant that a hypothesis needs to necessitate in order to explain, it is not clear why this must be metaphysical necessitation. Cameron (2007) considers the possibility that there may be mereological laws, such that the supervenience of composition facts on intrinsic facts can be sustained without being metaphysically necessary.

[E]ven if it is only a contingent fact that there are laws of mereology, this still affords some benefits to the contingency theorist. Firstly, we can sustain counterfactuals of the form ‘if there were another simple there would be another complex object’ without accepting necessitarianism. And secondly, it lets us define a notion of ‘mereological necessity’: a restricted modality, such that something is mereologically necessary iff it is true in every world in which the actual laws of mereology hold. This may prove useful because it lets us hold that the intra-world supervenience claim, while metaphysically contingent, is mereologically necessary, and this might go some way to defusing the intuitions against the possibility of the mixed-worlds. Sure, there is a sense in which they are impossible: but they are only mereologically impossible, not metaphysically impossible. (Cameron 2007, 114)

More work may well be needed to give us reason to think that there are such laws. However, insofar the concern is that we ensure that there is a deeper explanation for why composition

occurs when it does, less than metaphysically necessary supervenience of the composition facts on intrinsic facts is the more conservative hypothesis for meeting that concern.

This also goes some way to answering the other worry, about the threat of scepticism: That, if we allow that facts about composition *can* vary independently of base-level, non-compositional facts, then we must concede that we have no way of telling whether Hudson, van Inwagen, or any other rival theorist is correct about the actual world, for the world would be qualitatively indiscernible either way. Arguably, this sceptical worry is analogous to more general knowledge scepticism, yet few would take the latter to give us reason to think that brain-in-vat scenarios are impossible. The mere possibility of such epistemically troubling worlds should not force us to delimit modal space – it is enough that we look to simplicity as a guide for what goes on in the actual world. The notion of mereological laws allows us to accept the possibility of the mixed worlds that Rasmussen describes, while nonetheless maintaining that ours is not a world in which objects are flitting in and out of existence regardless of changes in intrinsic features.

It seems that both those who affirm (C) and those who do not can agree that there is something unpalatable about the suggestion that the actual world is one in which objects are flickering in and out of existence in the way that Rasmussen described. Those who do *not* affirm (C), however, may nonetheless reject the actuality of that scenario by way of principles which, by themselves, do not entail that the scenario is *impossible*. Instead, they might claim, first, that we should posit existence facts – in particular, facts about composition and identity – that would best explain the qualitative arrangement of the world as we know it; and, secondly, that, on balance, the best theory will be one according to which existence facts do not vary in the way that Rasmussen worries they *might* if (C) were false.

This way of thinking might yet lead us to a metaphysical necessity claim: for we might consider that any considerations of theoretical virtue that make it most reasonable to deny that the flickering scenario is actual would be just as forceful in *any* possible world. Hence, if we allow ourselves to reason abductively about other possible worlds, we would end up concluding that, of any possible world we might consider, the flickering scenario does not obtain. And *if we allow ourselves to agglomerate* these potential beliefs about each possible world, this will amount to saying that the flickering scenario is *impossible* – thus potentially undermining the contingentist’s project (insofar as contingentism commits one to the possibility of mixed worlds).

Take Josh Parsons’s (2013) case for contingent mereological universalism. Parsons endorses the General Sum Principle:

(GSP). There is a mereological sum of the Xs iff there are the Xs. (Parsons 2013, 332).

But he denies that it is a conceptual truth; rather, for Parsons, it is empirical and contingent. First, Parsons argues that “composite objects play a crucial role in the best explanations of my experience” (Parsons 2013, 332). Nihilists will of course deny this, arguing instead that such explanations need only cite particles arranged in certain ways; but Parsons urges that in many cases, the nihilist’s attempt to explain and predict referring only to particles arranged x-wise would lead to an “intractable many-body problem” (Parsons 2013, 333). From there – having accepted that we have good reason to allow the existence of those objects that feature in our scientific and proto-scientific explanations – Parsons argues that we should be universalists.

For there is no limit to what kinds of things compose other things that is both adequate and acceptable. To be adequate, a limit on composition must not rule out any objects that are cited in scientific explanation; to be acceptable, it must not be unacceptably anthropocentric. For example, the thesis that there are all and only the composite objects cited in actual scientific explanation is the adequate composition thesis *par excellence*, but it is hopelessly anthropocentric. What a grand stroke of luck that the universe provides us with just those objects that we actually cite in our explanations. (Parsons 2013, 333).

Parsons's argument is partly an abductive one, partly an indispensability one, and partly a "no miracles" argument: many of our best explanations cite composite objects, so we have reason to posit them. But it would be something of a miracle if the only composite objects were the ones that *actually* figured in such explanations; if we want a theory that is free from miraculous implications, we should endorse mereological universalism.

As we have seen, Parsons believes mereological universalism to be *contingently* true. But if we find Parsons's argument sufficiently compelling to warrant belief in mereological universalism at the actual world, wouldn't we find it equally compelling in any other possible world? What could possibly reduce the force of the argument?

Perhaps something could: perhaps, Parsons says, there might be "some non-anthropocentrically specifiable feature that the composite objects cited in our explanations had – if, for example, the only objects that ever played any role in our explanations were perfectly continuous – that would be a reason for thinking that those were all the composite objects" (Parsons 2013, 333). This seems to be another instance where PEP, or something like it, is presupposed, for Parsons is effectively claiming that, given that there is a world at

which it is objectively reasonable to conclude that the only composite objects are perfectly continuous, then we should say that a theory of restricted composition is possibly true.

However, it is not clear that the feature that Parsons offers here as an example would avoid the miracles worry that he takes to favour (contingent) universalism. It seems fair to say the only way that *only* perfectly continuous objects would be cited in our explanations is if our theoretical concerns were contingently limited. And even if the only objects that ever played any role in our explanations were perfectly continuous, or medium-sized, or malleable, or whatever, the objection Parsons raised for imposing a limit on what kinds of things compose other things would still apply: “What a grand stroke of luck that the universe provides us with just those objects that we *actually cite* in our explanations” (Parsons 2013, 333; emphasis added); in the same spirit, if the only objects ever cited in our explanations were of some particular kind, we might equally say: What a grand stroke of luck that the universe provides us with just those composite objects that our *actual explanations* cite. The shift in emphasis and wording is slight, but does make for a different, though related, point: in any world in which the only composite objects cited in (proto-)scientific explanations are of a particular kind, it would be just as unlikely that the explanations we actually give are the only ones worth giving.

Suppose that’s right: Parsons’s argument for mereological universalism at the actual world would be equally compelling at any other world. Then, if we may reason abductively about other worlds to find out what is the case at those worlds, then it seems we should say that, for any world we might consider, mereological universalism is the case at that world. If we agglomerate these beliefs, about which theory of composition is the correct one at each

world, then we end up with the conclusion that, at *every* world, mereological universalism is the case, i.e. mereological universalism is *necessarily* true.

But again, the question arises: Should we agglomerate these beliefs about each world? The line of argument just given is again similar to that which results in the Lottery Paradox, which may be taken to show that we should not agglomerate beliefs. In the lottery case, we are asked to consider a lottery in which the probability of any particular ticket winning is small enough that one is justified in believing, of any particular ticket, that it is not a winner. This is perhaps true of many of the lotteries we are used to. But then if we agglomerate these (potential)⁷⁵ beliefs about *each* ticket, we end up concluding that *no* ticket is a winner. That, however, is counterintuitive; at least for many of the lotteries we are used to, one can expect – or at least, should not rule out – that some ticket is a winner. In any case, there is no reason why we may not build it into the case that some ticket *is* a winner, and that this is known to the ticket holders – it is just that the chance of it being any particular ticket is vanishingly small.

One can therefore remain a contingentist, in spite of maintaining the actual truth of a thesis on the basis of an argument whose force is not diminished across possible worlds: one would simply need to deny that one should agglomerate one's (potential)⁷⁶ beliefs about each possible world. If the modal case and the lottery case are relevantly similar, this denial finds independent support from consideration of the lottery paradox.

⁷⁵ It is not usually required that one actually consider each ticket. It should be enough that any ticket that one *might* consider has a sufficiently low chance of being a winner to warrant the belief that it is not a winner.

⁷⁶ One might object that one does not, and moreover could not, hold beliefs about each individual world; but, as with the lottery case, it should be enough that, for *any* world that one *might* consider, mereological universalism is warranted on the basis of the same argument that Parsons makes for its actual truth.

In the previous chapter we saw how an inference to the best explanation, based on possible evidence, can justify a possibility claim, and I attempted to articulate and defend a principle, PEP, underlying such inferences. Putting a priori inferences to one side (PEP is conveniently restricted to inferences where observable facts play a crucial role), the principle cannot, however, justify necessity/impossibility claims. Even if *every* world is one in which the observable facts make it objectively reasonable to conclude that p, we cannot, on that basis alone, conclude that p is necessarily true; likewise for not-p: from the claim that it would be objectively reasonable to conclude that not-p at each world, we cannot on that basis alone infer that that p is impossible. To do so would be to agglomerate beliefs about each world into a belief about all worlds, which, as we have just seen, is a form of inference that leads to paradox.

This has consequences for certain otherwise tempting inferences. I'll give three examples. First, as we saw in Chapter 1, Prosser (2007) argued that, given that it is nomologically impossible for us to experience the passage of time, and that our experience as of the passage of time is the main reason we have for thinking time passes, we should conclude that time does not pass. One of Prosser's premises states that the mental supervenes on the physical with nomological necessity. Strengthening that premise, to say that mental-physical supervenience holds with metaphysical necessity, the argument becomes one for the metaphysical impossibility of experiencing the passage of time. In that case, every world is relevantly similar: if the impossibility of experiencing passage gives us reason to think that time does not pass at this world, it gives us reason to think, of any possible world we might consider, that time doesn't pass at that world. Should we then say that the passage of time is

impossible? If we reject agglomeration, we cannot infer from ‘At each world, time doesn’t pass’ to ‘Every world is such that time doesn’t pass’.

A second example: suppose Warmbrod was right when he said:

The observation and timing of a freeze requires there to be a timekeeper of some sort which is unfrozen and changing over the period of time in question. Hence the observational data will always support a conclusion that not everything is frozen. (Warmbrod 2004, 273)

In that case, in each world that conforms to Shoemaker’s stipulations, the conclusion to be drawn is that there are no periods of time without change. Should we not then say that *no* Shoemaker-world is one in which time passes without change? Not if we reject the agglomeration of claims about each world.

A third and final example. Della Rocca (2005) discusses the question of whether it is possible for two duplicates to permanently coincide, that is, to exist at the same place throughout their existence. Now it seems fair to say that Occam’s razor will always tell us to deny that there are permanently co-located duplicate objects, and that this applies to any possible world we might consider. Should we not then say that there is no world in which permanently co-located duplicates reside? No – at least, not simply on the basis of what would be reasonable to conclude at every possible world.

But abductive inferences may yet lead us to necessity/impossibility claims without having to agglomerate beliefs about each world. The following section illustrates Biggs’s claim that

modality is centrally abductive, and notes some ways in which our understanding of abductive inference needs to be tweaked to accommodate that claim.

6. Abduction and Modality: Choosing a necessity claim from a range of potential explanations.

As we have seen, Faraci cites Biggs's (2011) paper, "Abduction and Modality", as making trouble for the attempt to restrict the power of abduction to justifying claims *at* worlds by those who inhabit them, but not *about other* worlds. But Faraci's characterisation of abduction cannot accommodate Biggs's central claim, that abduction can justify claims of metaphysical necessity. As we saw, Faraci said: "one of the defining features of abduction as a form of non-deductive reasoning is that its conclusions are not necessitated by its premises, that the couple's reconciliation is the best possible explanation of what I saw but not the only one" (Faraci 2013, 752). While Faraci is surely right that there is more than one possible explanation in the case he mentions – where the thing to be explained is the appearance of a recently estranged couple walking hand-in-hand – this will not be so in all cases, particularly not those that Biggs is concerned with; for, as we will see shortly, in the example that Biggs offers, each of the candidate (i.e. 'potential') explanations rules out the very possibility of each of the others.

Biggs argues that the epistemology of modality is centrally abductive. Several points are offered in favour of this view. One is that it aids a popular stance on the mind-body problem, a posteriori physicalism, which is comprised of two theses:

- (i) Physical properties necessitate all properties.
- (ii) Physical truths do not a priori entail some phenomenal truth. (Biggs 2011, 285)

According to Biggs, standard modal epistemologies cannot accommodate a posteriori physicalism, since standard modal epistemologies “hold that failure of a priori entailment entails failure of necessitation” (Biggs 2011, 285). Hence, given any standard modal epistemology, (i) entails that (ii) is false, which is to say that standard modal epistemologies cannot countenance both of the theses that are essential to a posteriori physicalism.

Another point offered in favour of the claim that modality is centrally abductive is that it provides an attractive response to Peacocke’s (1999) integration challenge. If there are real, mind-independent necessities/possibilities, it would be a striking – to the point of being puzzling – fact if investigating our own minds revealed them to us. Yet this is what standard modal epistemologies involve, when they claim that failure of a priori entailment entails failure of necessitation (Biggs 2011, 287-288). By contrast, Biggs says, “an abduction-centered modal epistemology for scientific entities holds that no amount of reasoning, however idealized, entails the truth of modal beliefs about scientific entities. The shift is an improvement, because one expects truths about mind-independent properties of scientific entities ^[77] to outstrip ideal justification” (Biggs 2011, 291).

Biggs then goes on to show that abductive principles can favour certain explanans over others where the explanans in question are either modal claims, or else claims that entail modal claims (beyond those less metaphysically interesting possibility claims that are entailed by

⁷⁷ “Broadly scientific entities include inter alia electrons, trees, and pains. Broadly scientific entities exclude inter alia numbers and logical axioms. This domain-restriction is standard in various discussions in the metaphysics of science. Discussions of physicalism, causation, and laws of nature, for example, often take broadly scientific entities as their explananda. For ease of presentation, ‘broadly’ will be presumed” (Biggs 2011, 284).

claims about what is actual). In the example Biggs gives, the explanandum is the co-occurrence of water and H₂O, and three potential explanations are offered:

- (I) H₂O is identical with water.
- (M) H₂O merely metaphysically necessitates water – where ‘merely’ indicates that the necessitation holds in only one direction.
- (N) H₂O merely nomologically necessitates water – where ‘merely’ indicates that the necessitation is not metaphysical. (Biggs 2011, 297 and 299)

As Biggs notes, these three potential explanations are mutually exclusive. If H₂O is identical with water, as according to (I), then water and H₂O metaphysically necessitate each other (assuming, that is, that property identities are metaphysically necessary). Thus (I) rules out (M) because (I) has it that the necessitation goes in both directions, while (M) has it that the necessitation holds in only one direction; and (I) rules out (N) because (I) implies that the necessitation is metaphysical, not merely nomological. Similarly, (M) and (N) rule each other out because they contradict each other on the strength of the necessitation – metaphysical according to the former, and nomological according to the latter.

That the competing candidate explanations are mutually exclusive is typical of an abductive assessment (Lipton 2004, 57). Less typical, however, is the strength with which Biggs’s potential explanations exclude each other, for if any of (I), (M), or (N) is true, then the other two candidates are each ruled out *as possibilities*.

At first sight, this may seem to be at odds with Lipton’s characterisation of inference to the best explanation. Operating on an assumption of inferential and explanatory realism for much of his book on inference to the best explanation, Lipton (2004) says that “for something

to be an actual explanation, it must be (at least approximately) true” (57). Hence, since competing explanations are typically mutually exclusive, “Inference to the Best Explanation cannot be understood as inference to the best of the *actual* explanations”; instead, Lipton suggests, we ought to distinguish *actual* from *potential* explanation (Lipton 2004, 57). Now a natural reading of ‘potential’ is one according to which it entails possibility. On this interpretation, “the potential explanations of some phenomena are those that do explain them in a possible world where our observations hold” (Lipton 2004, 59). However, this cannot be the sense in which Biggs’s candidates are potential explanations, for once we infer to the best explanation, we will look back on the other candidates and deny that they were possible explanations after all; that’s because, as we have seen, in assenting to the *best* explanation among the three candidates that Biggs considers – whichever that turns out to be – we would be committed to the *impossibility* of any of its competitors obtaining. Thus, if we are to allow abductive inferences to necessity claims, choosing from a list of mutually exclusive candidate explanations, we need some way of cashing out ‘*potential* explanation’ that doesn’t presuppose the possibility of every candidate explanation. (Incidentally, such a reading of ‘potential’ is also required – albeit for a different reason – if we are to allow abductive inferences to possibility claims, along the lines discussed in the previous chapter. This is because, if an inference to the best explanation is to justify a certain possibility claim, it would beg the question to assume the possibility of all candidate explanations.)

Biggs cashes out the sense in which his candidates, (I), (M) and (N), are potential explanations by appealing to conceivability.

Conceivings play a central role in generating that list: a theory is a potential explanation if that theory is conceivably true. Accordingly, one knows that (M) and (N) are potential explanations for the invariable co-occurrence of water and H₂O if one can conceive of scenarios that verify them. Conceivings, therefore, play a significant role in modal inquiry without playing a direct role in selection of the best explanation. (Biggs 2011, 318)

While this involves commitment to the idea that metaphysical impossibilities are conceivable, perhaps Biggs would not take this to be a great cost – it is already a commitment of those a posteriori physicalists who grant the conceivability, but not the possibility, of zombies (microphysical duplicates of us who lack consciousness); and after all, as we have seen, Biggs intends his abductive modal epistemology to suit such theorists.

But perhaps it is worse than that: it is not simply that *some* impossibilities are conceivable; rather, among the conceived theories that make up our list of candidate explanations, conceived *possibilities* are the *exception*. In Biggs's example, at most one of the candidate explanations is possible, and the other two are impossible. This threatens to make conceivability worse than merely fallible when it comes to establishing possibility claims; if all but one of our candidate explanations are impossible, then conceivability may be positively *unreliable* as a guide to possibility.

To remedy this issue, we can draw on Rosen's (2006) notion of correct conceivability, which we met in the Introduction. Recall that a proposition is correctly conceivable just in case it harbours no contradiction "given a full specification of the natures of the items it concerns" (Rosen 2006, 24). With this notion in hand, we can say that there is a sense in which all three potential explanations, (I), (M) and (N), are conceivable – but they are not all *correctly* conceivable. In this way, Biggs's abductive modal epistemology can be seen as dovetailing

with Rosen's method of correct conceivability. Correct conceivability can only help us learn about modality when we have sufficient insight into the natures of the entities we are interested in. But it is not clear that a "full specification" of the nature of water can be given without saying something about its modal profile; nor is it clear that empirical investigation alone can reveal modal properties. As Miller has noted, "Some find it puzzling how discovering something about the microstructure of water reveals that its essence is H₂O, since it does not appear that scientists are in the business of discovering essences" (Miller 2010, 967). While drawing on the empirical data, an inference to the best explanation can at once be seen as teaching us which of the potential explanations, (I), (M), or (N), is correctly conceivable, while at the same time filling in the modal profile of water.

6. Final Thoughts.

The first three chapters of the present work discussed several cases in which contingent evidence could reasonably be brought to bear on certain issues of speculative metaphysics. Some of the examples given were of actual contingent evidence being marshalled in favour of a certain metaphysical thesis. Others were suggestions of merely possible evidence. Once we admit the former as evidence of actual truth, it is hard to see why we cannot reasonably admit the latter as evidence of *possible* truth, for we can always ask of contingent evidence: what if it did not obtain? What bearing *would* its non-obtaining have on the case for the metaphysical thesis whose truth it *actually* supports?

For many metaphysical debates, particularly those concerning space, time, and identity, empirical science has much to contribute. But often the scientific debates that would have an impact on theory choice are themselves far from settled. For example, I briefly mentioned in the Introduction that relativistic considerations have been thought to weigh in on the debate over whether spacetime points are substantival or somehow reducible to relations among objects; and in Chapter 3 that a certain interpretation of quantum phenomena has been suggested to refute the principle of identity of indiscernibles (PII). What if it turns out that general relativity and quantum theory cannot be integrated? Allowing merely possible evidence to be a guide to possibility means that, even if some further data emerges such that the scientific theory lending weight to a certain metaphysical claim is superseded, the mere possible truth of such a scientific theory can be taken to bear on the possible truth of the relevant metaphysical claims. So unless the account of quantum phenomena upon which French (1988) bases his rejection of PII is *necessarily* false, there may be an equally strong argument against the necessity of PII based on merely possible evidence.

I mentioned in Chapters 4 and 5 Kripke's remarks about possible worlds: that they are not foreign countries that we can look at with a powerful telescope; that possible worlds are not found out, but stipulated (Kripke 1980, 45). To some extent, Kripke was right about possible worlds and the nature of our access to what goes on in them. Which particular possible world we are referring to is in large part a matter of stipulation. We cannot stipulate everything about a possible world of course, but we hope that what we have left out of our characterisation of a possible world under discussion has no bearing on the case at hand.

There comes a point when we take it that we have specified enough about a possible world to allow us to draw conclusions about it. When, in Chapter 4, I talked about 'Shoemaker's world' of partial freezes, there is an incredible amount that I left unspecified (Is football the most popular sport at that world? Do they eat animals there? How many toes do the inhabitants have?) – none of which, I hope, was relevant. It would take a Godly intellect to determinately imagine such worlds in every detail, and presumably no worldly philosopher is possessed of such capabilities. But that doesn't mean that the worlds we imagine are somehow metaphysically vague entities: we can't determinately imagine such worlds, but we nonetheless imagine such worlds as being determinate (Yablo 1993, 29). Our inability to specify all the facts about a possible world does not mean that we cannot learn about that world based on what we have specified. The same goes for an actual inference to the best explanation. When we make the inference, we take something of a leap of faith, assuming that we are not lacking in some crucial information that would favour a theory other than the one we are inferring to.

I have taken it throughout that imagination can, in principle, justify possibility claims. If consideration of merely possible evidence is to provide us with justification of further

possibility claims, we cannot start from scratch: we need some initial possibility claims to work with. So I am happy to allow imagination to play the role of telescope. But this telescope has a limited range. Whatever the correct account of imagination is, there are at the very least dialectical constraints on the extent to which imagination can be appealed to in order to justify possibility claims. Where a dispute about what is imaginable is at the heart of a dispute about what is possible, there is a threat of dialectical deadlock if sensory imagination and mere stipulation is all we have to go by.

With this in mind, I looked to Shoemaker's (1969) argument for a way forward. That argument proceeded from a consideration of merely possible evidence to the claim that time could pass without change (without simply asking us to imagine a world where nothing changes for a while). Abstracting from the particulars of Shoemaker's case (and drawing on Warmbrod's (2004) critique), I arrived at a general principle (PEP): If there is a possible world in which the observable facts make it objectively reasonable to conclude that *p*, then it is objectively reasonable for *us* to conclude that *p* is possible. The idea that it is legitimate to argue for possibility claims in this way is based on the idea that actual and possible justification are, to some extent, on a par. If we can draw on actual contingent evidence to favour conclusions not directly confirmable by experience, or to justify belief in entities not directly observable, why can we not draw on merely possible evidence to favour conclusions that are, for whatever reason, beyond justification by sensory imagination? If imagination is to the possible as perception is to the actual, imagination can furnish the raw data with which to extend our knowledge of possibility.

PEP is a tool for those who wish to argue for a possibility, and is geared towards expanding the sphere of possibility claims which we can claim to have justified. There are two respects

in which PEP is somewhat biased towards those wishing to expand this circle. One is that, by arguing against an argument like Shoemaker's (1969) for the possibility of time without change, one has not thereby argued for the impossibility of time without change; at most, one has thereby attempted to rebut an argument for it. If an argument satisfying PEP succeeds, it establishes a possibility claim; if it fails, nothing is gained, but nor is anything lost.

Secondly, as we saw in Chapter 5, unless the proposition satisfying PEP is itself a necessity claim, PEP cannot be used to justify even restricted necessities – even if the observable facts at *every* world make it objectively reasonable to conclude that *p*, we cannot conclude that *p* is necessary. Miller's (2010) suggestion (which we met in the Introduction), that “a synthetic thesis is necessary just in case it is preferable to its competitor theses in all worlds, and contingent where it is preferable to its competitors in only some worlds” (Miller 2010, 970) has some initial plausibility; but while I have effectively endorsed the latter claim by endorsing PEP, we saw in Chapter 5 that the former claim is based on a form of reasoning that can generate a modal analogue of the lottery paradox, and hence should be rejected.

I closed Chapter 5 with a brief discussion of Biggs's abductive modal epistemology. As we saw, Biggs has it that an inference to the best explanation can justify necessity claims. This is not achieved by way of agglomerating claims about each and every world, but simply by having necessities in one's shortlist of candidate explanations. It is perhaps tempting to look to necessities for one's explanations because necessities are explanatorily powerful. And since stronger necessities explain weaker ones, the stronger the necessity, the greater the explanatory power. One may also be tempted to argue for the necessity of one's preferred theory for less noble reasons. If one can convincingly argue that one's theory is metaphysically necessary, then one has knocked all competitors out of the park.

What becomes of contingent evidence when one argues for the necessity of one's metaphysical thesis? If such evidence is initially thought to count against one's preferred theory, then it must be somehow discounted. One way to do this is to explain it away – to provide an explanation of the evidence apparently in conflict with one's theory, such that that explanation and one's theory are compatible. But one had better make sure that one can explain away the problematic evidence at *all* worlds, or a conflict will remain. If consideration of merely possible evidence can in principle justify possibility claims, it is a tall order to establish metaphysical necessities.

Bibliography.

- Adams, Robert M. 1979. "Primitive Thisness and Primitive Identity." *The Journal of Philosophy*, lxxvi, 1 (January 1979): 5–26.
- Anscombe, G. E. M. 1959. *An Introduction to Wittgenstein's Tractatus*. London: Harper and Row.
- Armstrong, D.M. 1983. *What Is a Law of Nature?* Cambridge: Cambridge University Press.
- Armstrong, David. 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Balashov, Yuri. 2005. "Times of Our Lives: Negotiating the Presence of Experience". *American Philosophical Quarterly*, Vol. 42, No. 4: 295-309.
- Balashov, Yuri. 2010. *Persistence and Spacetime*. Oxford: Oxford University Press.
- Baldwin, Thomas. 1996. "There Might Be Nothing". *Analysis* 56: 231-238.
- Bealer, George. 2002. "Modal Epistemology and the Rationalist Renaissance", in *Conceivability and Possibility*, ed. T. Gendler and J. Hawthorne (Oxford: Clarendon Press, 2002): 71–126.
- Bennett, M. R, and P.M.S Hacker. 2003. *Philosophical Foundations of Neuroscience*. Oxford: Wiley-Blackwell.
- Benovsky, Jiri. 2012. "The Causal Efficiency of the Passage of Time". *Philosophia* 40:763–769.
- Benovsky, Jiri. 2015. "From Experience to Metaphysics: On Experience-based Intuitions and their Role in Metaphysics". *Nous* 49: 4 684–697.
- Biggs, Stephen. 2011. "Abduction and modality". *Philosophy and Phenomenological Research*, 83(2): 283–326.
- Bird, Alexander. 2007. *Nature's metaphysics: Laws and properties*. Oxford: Oxford University Press.
- Black, Max. 1952. "The Identity of Indiscernibles." *Mind*, lxi: 153–64.
- Blackburn, Simon. 1986. "Morals and Modals", in G. MacDonald and C. Wright, ed., *Fact, Science, and Morality*. Oxford: Oxford University Press.
- Bohn, Einar. 2009. "An Argument Against the Necessity of Unrestricted Composition". *Analysis* 69 (1): 27-31.
- Bohn, Einar. 2014. "Unrestricted Composition as Identity", in D. Baxter and A. Cotnoir (eds) *Composition as Identity*, 143–65. Oxford: OUP.
- Borges, Jorge Luis. 2000. *The Aleph: Including the Prose Fictions from The Maker*. London: Penguin.
- Bourne, Craig. 2006. *A Future for Presentism*. Oxford: Clarendon.
- Bricker, Phillip. 1991. "Plenitude of Possible Structures". *The Journal of Philosophy*, Vol. 88, No. 11, Eighty-Eighth Annual Meeting American Philosophical Association, Eastern Division (Nov., 1991): 607-619.
- Brown, James Robert. 2004. "Why Thought Experiments Transcend Experience", in *Contemporary Debates in Philosophy of Science*, ed Christopher Hitchcock (Malden, Mass: Blackwell): 23-43.

- Brown, Philip. 2013. "The Possibility of Morality". *Philosophical Studies* 163 (3): 627-636.
- Butterfield, Jeremy. 2006. "The Rotating Discs Argument Defeated". *British Journal for the Philosophy of Science* 57 (1):1-45.
- Byrne, Alex. 2007. "Possibility and imagination". *Philosophical Perspectives*, 21: 125–44.
- Callendar, Craig. 2000. "Shedding Light on Time". *Philosophy of Science (Proceedings)*, 67: 589–99.
- Calosi, Claudio. 2016. "Composition is Identity and Mereological Nihilism". *The Philosophical Quarterly* Vol. 66, No. 263: 219-235.
- Calosi, Claudio. Unpublished. "Persistence and Change in Minkowski Spacetime". philsci-archive.pitt.edu/8629/1/Persistence_and_Change_in_Minkowski_Spacetime.pdf . Accessed 10/07/2016.
- Cameron, Ross. 2007. "The Contingency of Composition". *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, Vol. 136, No. 1, Selected Papers from the 2006 Bellingham Summer Philosophy Conference (Oct., 2007): 99-121.
- Cameron, Ross. 2012. "Composition as Identity Doesn't Settle the Special Composition Question". *Philosophy and Phenomenological Research* 84 (3): 531-554.
- Chalmers, David. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Chalmers, David. 2002. "Does Conceivability Entail Possibility?" in *Conceivability and Possibility*, ed. T. Gendler and J. Hawthorne (Oxford: Clarendon Press, 2002): 145–200.
- Craig, W. L. 2000. *The Tensed Theory of Time: A Critical Examination*. Dordrecht: Kluwer Academic Publishers.
- Cross, Charles B. 2009. "Causal Independence, the Identity of Indiscernibles, and the Essentiality of Origins." *Journal of Philosophy*, 106 (5):277-291.
- Dainton, Barry. 2006. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.
- Dainton, Barry. 2010. *Time and Space: Second Edition*. Durham: Acumen.
- Davies, Paul. 1995. *About Time: Einstein's unfinished revolution*. Penguin: London.
- Della Rocca, Michael. 2005. "Two Spheres, Twenty Spheres, and the Identity of Indiscernibles". *Pacific Philosophical Quarterly* 86 (4):480–492.
- Earman, J. 1989. *World Enough and Spacetime*. Cambridge, MA.: MIT Press.
- Earman, J. and J. D. Norton. 1987. "What Price Spacetime Substantivalism? The Hole Story". *British Journal for the Philosophy of Science* 38: 515–25.
- Effingham, Nikk. 2013. "Harmoniously Investigating Concrete Structures". *Thought* 2 (3): 190-195.
- Efrid, David & Stoneham, Tom. 2005. "The Subtraction Argument for Metaphysical Nihilism". *The Journal of Philosophy* 102: 303-325.
- Evnine, Simon. 2008. "Modal Epistemology: Our Knowledge of Necessity and Possibility". *Philosophy Compass*, 3/4 (2008): 664–684.

- Faraci, David. 2013. "Brown on Mackie: Echoes of the Lottery Paradox". *Philosophia* 41: 751-755.
- Forrest, Peter. 1995. "Is Space-time Discrete or Continuous? – An Empirical Question". *Synthese* 103: 327-354.
- Forrest, Peter. 2012. "The Identity of Indiscernibles". *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2012/entries/identity-indiscernible/>>. Accessed 19/09/2014.
- Frankish, K. 2007. "The Anti-Zombie Argument". *Philosophical Quarterly*, 57: 650–666.
- French, Stephen. 1989. "Identity and Individuality in Classical and Quantum Physics". *Australasian Journal of Philosophy*, 67: 432–446.
- French, S.: 1998, "On the Withering Away of Physical Objects", in Castellani (ed.), *Interpreting Bodies*, Princeton University Press: Princeton.
- Gallois, Andre. 1974. "Berkeley's Master Argument". *The Philosophical Review*, Vol. 83, No. 1: 55-69.
- Gendler, Tamar Szabo, and John Hawthorne. 2002. "Introduction", in *Conceivability and Possibility*, ed. T. Gendler and J. Hawthorne (Oxford: Clarendon Press, 2002): 1-70.
- Hacking, Ian. 1975. "The Identity of Indiscernibles". *The Journal of Philosophy*, Vol. 72, No. 9 (May 8, 1975): 249-256.
- Hales, Steven D., and Johnson, Timothy A. 2003. "Endurantism, Perdurantism, and Special Relativity", in *The Philosophical Quarterly*, Vol. 53, No. 213: 524-539.
- Hanrahan, Rebecca. 2007. "Imagination and Possibility". *Philosophical Forum* 38 (2):125–146.
- Hanrahan, Rebecca. 2009. "Consciousness and Modal Empiricism". *Philosophia*, 37(2): 281-306.
- Hart, W. D. 1988. *Engines of the Soul*. Cambridge: Cambridge University Press.
- Hawley, Katherine. 2001. *How Things Persist*. Oxford: OUP.
- Hawley, Katherine. 2006. "Science as a Guide to Metaphysics?" *Synthese* 149: 451–470.
- Hawley, Katherine. 2006. "Weak Discernibility". *Analysis* 66.4, October 2006: 300–303
- Hawley, Katherine. 2009. "Identity and Indiscernibility". *Mind* 118 (469): 101 - 119.
- Hawthorne, John. 2001. "Causal Structuralism", *Philosophical Perspectives* 15: 361–378.
- Hume, David. 1967. *A Treatise of Human Nature*. Oxford: Clarendon Press.
- Ingram, David. Ms. "Thisness and Material Objects".
- Kment, Boris. 2006. "Counterfactuals and the Analysis of Necessity." *Philosophical Perspectives* 20: 237–302.
- Kochiras, Hylarie. 2014. "Locke's Philosophy of Science", *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/sum2014/entries/locke-philosophy-science/>>.
- Kripke, Saul. 1981. *Naming and Necessity*. Blackwell.

- Kyburg, Henry. 1961. *Probability and the logic of rational belief*. Middletown: Wesleyan University Press.
- Le Poidevin, Robin & D. H. Mellor. 1987. "Time, Change, and the 'Indexical Fallacy'". *Mind* 96 (384): 534-538.
- Le Poidevin, Robin. 1990. "Relationism and Temporal Topology: Physics or Metaphysics?" *Philosophical Quarterly*, Vol. 40, No. 160: 419-432.
- Le Poidevin, Robin. 2003. *Travels in Four Dimensions: The Enigmas of Space and Time*. Oxford: Oxford University Press.
- Le Poidevin, Robin. 2007. *The Images of Time: An Essay on Temporal Representation*. Oxford: Oxford University Press.
- Le Poidevin, Robin. 2010. "Time without Change (in three steps)". *American Philosophical Quarterly* 47 (2): 171-180.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, David. 1991. *Parts of Classes*, Oxford: Basil Blackwell.
- Lipton, Peter. 2004. *Inference to the Best Explanation*. Routledge: London.
- Locke, J., 1975, *An Essay Concerning Human Understanding*, ed. P. H. Nidditch, Oxford: Oxford University Press.
- Lowe, E. J. 1998. *The possibility of metaphysics: Substance, identity, and time*. Clarendon Press.
- Lowe, E.J. 1987. "The Indexical Fallacy in McTaggart's Proof of the Unreality of Time". *Mind* 96 (381): 62-70.
- Mackie, J. L. 1977. *Ethics: Inventing right and wrong*. Harmondsworth: Penguin.
- Markosian, Ned. 1998. "Simples". *Australasian Journal of Philosophy*, Vol. 76, No. 2: 213-228.
- Martin, M.G.F. 2002. "The Transparency of Experience". *Mind & Language*, Vol. 17 No. 4 September 2002: 376-425.
- McDaniel, Kris. 2010. "Composition as Identity Does Not Entail Universalism". *Erkenntnis* 73 (1): 97-100.
- McTaggart, J.M.E. 1908. "The Unreality of Time". *Mind* 17: 457-74.
- Mellor, D. H. 1974. "Special Relativity and Present Truth". *Analysis*, 34: 74-8.
- Mellor, D. H. 1998. *Real Time II*. London: Routledge.
- Merricks, Trenton. 1994. "Endurance and Indiscernibility". *Journal of Philosophy* 91 (4):165-184.
- Merricks, Trenton. 2001. *Objects and Persons*. Oxford: Oxford University Press.
- Miller, Kristie. 2009. "Defending Contingentism in Metaphysics". *Dialectica*, Volume 63, Issue 1: 23-49.
- Miller, Kristie. 2010. "Three Routes to Contingentism". *Philosophy Compass* 5/11: 965-977.
- Moran, R. 1994. "The Expression of Feeling in Imagination," *Philosophical Review*, 103(1): 75-106.
- Nerlich, Graham. 1994. *The Shape of Space: Second Edition*. Cambridge: Cambridge University Press.
- Newton-Smith, W. H. 1980. *The Structure of Time*. London: Routledge.
- Nolan, Daniel. 2004. "Classes, Worlds and Hypergunk". *The Monist* 87(3): 3-21.

- Noordhof, Paul. 2002. "Imagining Objects and Imagining Experiences". *Mind & Language*, Vol. 17 No. 4 September 2002: 426–455.
- Oppy, Graham. 2004. "Can we Describe Possible Circumstances in which we would Have *Most* Reason to Believe that Time is Two-Dimensional?" *Ratio (new series)* XVII 1 March 2004: 67-83.
- Parsons, Josh. 2007. "Theories of Location". *Oxford Studies in Metaphysics: Volume 3*, 201-232, Dean Zimmerman (Ed.). Oxford: OUP.
- Parsons, Josh. 2013. "Conceptual Conservatism and Contingent Composition". *Inquiry* 56 (4): 327-339.
- Parsons, Josh. 2013. "The Earth and the Aleph". Manuscript. Available at <http://www.joshparsons.net/draft/aleph/aleph.pdf> . Accessed 12/05/15.
- Parsons, Josh. 2015. "A Phenomenological Argument for Stage Theory". *Analysis* 75 (2): 237-242.
- Paul, Laurie. 2010. "Temporal Experience". *Journal of Philosophy* 107 (7):333-359.
- Peacocke, Christopher. 1985. "Imagination, Experience and Possibility: A Berkeleyian View Defended". In J. Foster and H. Robinson (eds.), *Essays on Berkeley*. Oxford: Oxford University Press: 19–35.
- Peacocke, Christopher. 1999. *Being Known*. Oxford: Clarendon Press.
- Penrose, Roger. 2004. *The Road to Reality*. London: Jonathan Cape.
- Plantinga, Alvin. 1993. *Warrant and Proper Function*. Oxford: Oxford University Press.
- Priest, Graham. 2006. *In Contradiction: Second Edition*. Oxford: Oxford University Press.
- Prosser, Simon. 2007. "Could we Experience the Passage of Time?" *Ratio* 20 (1): 75-90.
- Prosser, Simon. 2012. "Why does time seem to pass?" *Philosophy and Phenomenological Research*, 85 (2012): 92-116.
- Putnam, Hilary. 1967. "Time and Physical Geometry". *Journal of Philosophy*, 64: 240–7.
- Putnam, Hilary. 1975. *Mind, Language and Reality, Philosophical Papers ii*. Cambridge: Cambridge University Press.
- Rasmussen, Joshua. 2014. "Continuity as a Guide to Possibility". *Australasian Journal of Philosophy*, Vol. 92, No. 3: 525–538.
- Rea, Michael C. 1998. "Temporal Parts Unmotivated", in *Philosophical Review*, 107: 225-260.
- Reed, Baron. "How to Thing About Fallibilism". *Philosophical Studies* 107 (2): 143-157.
- Rodriguez-Pereyra, Gonzalo. 2004. "The Bundle Theory Is Compatible with Distinct but Indiscernible Particulars". *Analysis*, 64: 72-81.
- Rosen, Gideon. 2006. "The Limits of Contingency". *Identity and Modality*. Ed. F. McBride. Oxford: Oxford University Press: 13–38.
- Salmon, Nathan. 1989. "The Logic of What Might Have Been". *Philosophical Review*, 98: 3-34.
- Saunders, Simon. 2002. "How Relativity Contradicts Presentism", in C. Callender (ed.), *Time, Reality and Experience* (Cambridge: Cambridge University Press): 277–

92.

- Saunders, Simon. 2006. "Are quantum particles objects?" *Analysis* 66: 52–63.
- Schaffer, Jonathan. 2005. "Quiddistic Knowledge". *Philosophical Studies* 123 (1-2): 1-32.
- Schuster, M. M. 1986. "Is the Flow of Time Subjective?" *The Review of Metaphysics*, 39: 695–714.
- Shimony, A. 1993. "The Transient Now", in his *Search for a Naturalistic World View, Volume 2* (Cambridge: Cambridge University Press): 271–90.
- Sidelle, Alan. 2002. "On the Metaphysical Contingency of Laws of Nature," in *Conceivability and Possibility*, T. Gendler and J. Hawthorne (eds.), Oxford: Clarendon, pp. 309–336.
- Sider, Ted. 2001. *Four-Dimensionalism*. Oxford: OUP.
- Sider, Ted. 2013. "Against Parthood". In Karen Bennett and Dean W. Zimmerman, eds., *Oxford Studies in Metaphysics*, volume 8 (Oxford: OUP, 2013): 237–93.
- Sklar, Lawrence. 1974. *Space, Time, and Spacetime*. London: University of California Press.
- Skow, Bradford. 2009. "Relativity and the Moving Spotlight". *The Journal of Philosophy*, CVI, 12 (December 2009): 666-78.
- Skow, Bradford. 2011. "Experience and the Passage of Time", in *Philosophical Perspectives*, 25, *Metaphysics*: 360-387.
- Sorensen, Roy. 1992. *Thought Experiments*. Oxford: Oxford University Press.
- Sorensen, Roy. 2014. "Epistemic Paradoxes". *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2014/entries/epistemic-paradoxes/>.
- Stanford, Kyle. 2001. "Refusing the Devil's Bargain: What Kind of Underdetermination Should We Take Seriously?" *Philosophy of Science*, Vol. 68, No. 3, Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 2001): S1-S12.
- Tahko, Tuomas. 2008. "A New Definition of A Priori Knowledge: In Search of a Modal Basis". *Metaphysica* 9 (2):57-68.
- Traynor, Michael. 2013. "Actual Time and Possible Change: A Problem for Modal Arguments for Temporal Parts". *Thought* 2 (3): 180-189.
- Traynor, Michael. 2014. "Phenomenal Experience and the Metaphysics of Experience". *Proceedings of the Aristotelian Society* 114 (3pt3): 381-388.
- Turing, Alan. 1950. "Computing Machinery and Intelligence," *Mind*, 59 (236): 433–60.
- Van Fraassen, Bas. 1985. "Statistical Behaviour of Indistinguishable Particles: Problems of Interpretation", in P. Mittelstaedt and E. W. Stachow (eds.), *Recent Developments in Quantum Logic*, Mannheim: B.I. Wissenschaft, pp. 161–187.
- Van Fraassen, Bas. 1991. *Quantum Mechanics: An Empiricist View*, Oxford: Oxford University Press.
- Warenski, Lisa. 2009. "Naturalism, Fallibilism, and the A Priori". *Philosophical Studies*, 142(3): 403-426.
- Warmbrod, Ken. 2004. "Temporal Vacua". *Philosophical Quarterly* 54 (215): 266–286.
- Weingard, Robert. 1972. "'Relativity' and the Reality of Past and Future Events". *British Journal for the Philosophy of Science*, 23: 119–21.

- Williams, Bernard. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.
- Wittgenstein, Ludwig. 1958/1965. *The Blue and Brown Books*. New York: Harper Torchbacks.
- Yablo, Stephen. 1993. “Is Conceivability a Guide to Possibility?” *Philosophy and Phenomenological Research* 53.1 (1993): 1–42.
- Yablo, Stephen. 2000. “Textbook Kripkeanism & the open texture of concepts”. *Pacific Philosophical Quarterly* 81: 98–122.