Humean revisions to the best system account of lawhood

Candidate: Mike Hughes
Supervisor: Mark Jago

Submitted in partial fulfillment for PhD Philosophy, University of Nottingham 2019

Contents

Contents 3
1. Introduction 5
1.1 Thesis objective 5
1.2 Classifying theories of lawhood 10
1.3 Humean Supervenience 15
1.4 Best System Account (BSA) 25
1.5 Other problems with the BSA 31
1.5.1 The empiricist loyalty test 31
1.5.2 The non-convergence problem 33
1.5.3 The non-denumerability problem 35
1.5.4 The 'special sciences' problem 36
1.5.5 The meta-laws problem 40
1.6 Non-reductive realism 41
1.7 Humean supervenience and causation 44
1.8 Preview of upcoming chapters 51
2. Four problems with the BSA 53
2.1 Introduction 53
2.2 The zero-fit problem 53
2.3 The undermining problem 56
2.4 The problem of non-locality 60
2.5 Armstrong's problem 63
3. BSA chances and the zero-fit problem 70
3.1 Introduction 70
3.2 BSA laws and chances 71
3.3 BSA chances and the probability calculus 82
3.4 Fit as maximum likelihood 86
3.5 Fit as maximum typicality 90
3.6 $\mathbf{X^2}$ fit 93
3.7 Discussion and conclusion 98
4. BSA chances and the undermining problem 100
4.1 Introduction 100
4.2 Best-system chances 101
4.3 Undermining futures 105

4.4 The T-world proposal 114
4.5 T-worlds and the problem of undermining 124
4.6 Discussion 126
5. The BSA and quantum mechanics 144
5.1 Introduction 144
5.2 Quantum mechanics 147
5.3 Three interpretation of quantum mechanics 149
5.4 Interpretations of quantum mechanics and the BSA 153
5.5 The problem of non-locality 165
5.6 The solution to non-locality 168
6. Supervenient perfectly natural properties 178
6.1 Introduction 178
6.2 Armstrong's problem 182
6.3 Algorithmic compression theory 185
6.4 The modal entropy conjecture 190
6.5 Modal entropy and the BSA 197
6.6 Supporting the conjecture 200
6.7 Objections and conclusion 202
7. Conclusions 204
8. Bibliography 207

1. Introduction

Best system accounts of laws all have in common one central thesis: that the laws of nature are the best systematizations of empirical data. The central thesis can be traced back to Mill (1904), Ramsey (1990), and many contemporary philosophers of science and metaphysicians, most notably Lewis across a number of publications (Lewis 1973, Lewis 1980, Lewis 1986 and Lewis 1994, among others).

All of these accounts vary in terms of the degree to which they reify the laws of nature and the positions taken on issues such as scientific reduction and epistemic access. However, all suffer from problems of one sort or another.

1.1 Thesis objective

The objective of this thesis is to solve some of the outstanding problems with the best system account (BSA) developed most extensively by Lewis over a number of publications as a foundational component of his Humean supervenience program (1973, 1980, 1986, 1994, 2000).

Briefly, the BSA as formulated by Lewis states that the fundamental laws of nature of a world are the set of lawlike propositions that conjointly describe the best systematization of all the extant properties of that world, conceived of as some multi-dimensional pattern of points or point-sized objects splayed over space and time (the "Humean mosaic", to use Lewis's nomenclature). "Best" in this context is of crucial importance, and for deterministic worlds is conceived of as being measured in terms of the strength and simplicity of the set of lawlike propositions (with some measure of probabilistic fit being added in the case of indeterministic worlds).

The BSA will be described in more detail later in this chapter. Of note, it will be a variant on the canonical Lewisian version, adopting for example a stronger sense of Humeanism supervenience than that implied by Lewis (see, for example, Lewis 1994, p474).

First however, a brief synopsis of the problems associated with the BSA that will be tackled in the main body of this thesis and introduced in more detail in Chapter 2 are as follows:

1. **The zero-fit problem**: in those worlds where the best systematization involves chance-like statements (i.e. in chancy or indeterministic worlds), the best systematization should result in the actual course of history – defined as patterns of occurrence of discreet chance events - coming out as the most likely of all possible histories. But all possible histories for such chancy worlds may be infinite in number, in which case the likelihood of all worlds is the same with the chance of each possible history being zero (because 1 divided by infinity is zero). A solution to the zero-fit problem will be given in Chapter 3 of this thesis.

2. **The undermining problem**: In worlds where the best system results in fundamental laws that contain (non-trivial) probability functions – that is, in chancy worlds – the probability functions are parameterized by the systematization of the chancy events in that world. But, by the nature of the world being chancy, the relative frequencies of these chancy events could have been very different. So, in those worlds where the relative frequencies did turn out very different, there may be probability functions appearing in fundamental laws with different parameters. Thus, it is to define some worlds that are nomologically possible in that they share the same deterministic and indeterministic laws but have different relative frequencies of some chance event. Yet, those same worlds would have – according to their own best systematization – different indeterministic laws and would therefore not be nomologically possible ("nomologically possible" here being

defined as not differing with respect to the fundamental laws). Thus, a situation arises where a lawful pattern of chance events can contradict the laws that made them lawful in the first place.

3. **Quantum mechanics and the problem of non-locality**: when quantum mechanics is considered as at least a possible candidate for the fundamental laws of the actual world, and given the Humean ban on any necessary connections built into the Humean mosaic of the actual world, a problem arises when certain phenomena in quantum mechanics suggest necessary connections may lawfully exist between two non-local points in space. This phenomenon, known as entanglement, has been demonstrated empirically, and therefore is thought to represent a known counterexample to the Humean ban on any necessary connections built into the Humean mosaic. This problem is not fatal to the BSA, for the best systematization process may proceed given any set of properties – whether non-local or not – that are built into the Humean mosaic. Yet, such non-local properties would imply that the BSA could no longer be formulated consistent with Humean metaphysics. The overarching aim of this thesis is to demonstrate how the BSA can be formulated using the (qualitatively) parsimonious metaphysics of Lewis's Humean supervenience, and thus the problem of non-locality is one whose solution is desirable. A Humean response to the problem will be given in Chapter 5.

4. **Armstrong's problem (also known as the immanent comparisons problem)**: the definition of "best" must presume some universal language or set of predicates common to all possible worlds in order to objectively measure the simplicity of a set of lawlike propositions and therefore determine whether the set is the best when traded off against the resulting strength. But no such universal language or set of predicates seems to be available to a reductive realist who must reject the

extravagant ontologies of the primitive realists and such primitive notions as natural kinds. A solution to this problem is critical for the BSA to be considered a realist theory of lawhood, as it relies on such objective measures of simplicity and strength and their trade-off. The mathematical framework already exists for measuring both simplicity and strength, but the predicates comprising the language over which the measures are made does not. A Humean method for identifying those predicates, or rather the second-order property that the properties designated by such predicates must possess, will be the purpose of Chapter 6 of this thesis.

A more in-depth rendering of the BSA and the problems associated with it, as well as how it fits in with an overarching classification scheme of other metaphysical theories of lawhood will be the main purpose of the first two chapters of this thesis.

In the following section this overarching classification scheme of philosophical theories of lawhood will be introduced with a brief discussion on the motivation and ontological commitment of each variant. Then, in the final section, the overarching metaphysical thesis of Humean supervenience, as developed primarily by Lewis, will be introduced, thus providing the philosophical context of the BSA and the main objectives of this thesis.

An overview of Humean supervenience and the BSA will be given in sections 1.3 and 1.4 below. In section 1.5 some of the other problems associated with the BSA will be introduced by way of providing a more comprehensive evaluation of the theory that goes beyond the restricted set of problems that will be the main focus of this thesis. No firm response will be offered to many of these problems, which lies beyond the scope of this thesis, but a brief synopsis of responses that have been made in the published literature will be noted wherever appropriate.

Sections 1.3 through 1.5 may be considered as laying the metaphysical groundwork for a Humean theory of lawhood and its subsequent development and evaluation in the form of the BSA. This overarching Humean metaphysics will then be contrasted with the non-Humean non-reductive realist approaches to lawhood in section 1.6.

Laws of nature in Lewis's Humean metaphysics formed the basis upon which other nomological notions were constructed. Namely, a possible world was in the privileged set of nomologically possible worlds only if they shared the same laws of nature (Lewis, 1986):

> More often than not, modality is restricted quantification; and restricted from the standpoint of a given world, perhaps ours, by means of so-called 'accessibility' relations. Thus it is nomologically necessary, though not unrestrictedly necessary, that friction produces heat: at every world that obeys the laws of our world, friction produces heat. It is contingent which world is ours; hence what are the laws of our world; hence which worlds are nomologically 'accessible' from ours; hence what is true throughout these worlds, i.e. what is nomologically necessary. (Lewis 1986, p7)

This then allows for notions of physical possibility and necessity and related counterfactuals to be defined over that set of possible worlds. Likewise, a Humean notion of causation can similarly be defined. Unlike nomological modality however and the BSA, Lewis's Humean theory of causation suffered fatal problems that he failed to fully overcome (see, for example, Bigaj, 2012; Collins, 2000). Although the generally recognized failure of Lewis's – and many others' – Humean theory of causation has little direct impact on the success or failure of the BSA, the Humean theories of causation often presume the validity of Humean theories of lawhood. To this end, a discussion of Humean theories of causation and how they may relate to the BSA will be given in 1.7.

Finally, before a closer examination of the four problems in chapter 2, in the final section of this chapter an overview of the other chapters comprising this thesis will be given.

1.2 Classifying theories of lawhood

It will be assumed throughout this thesis that metaphysical theories of lawhood are theories that aim to explain what laws of nature are, in any possible world in which they exist. The approach will be exclusively philosophical, not empirical, and thus no recourse will be made to evidence from the fields of the history, sociology or anthropology of science, for example. These fields of endeavor aim to describe the particular historical development of laws of nature in worlds like the actual world, or how human societies or individual scientists conceive of laws of nature. In contrast, the objective of this thesis is the further development of a metaphysical theory of lawhood, one that is true in all possible worlds, irrespective of whether there are humans or other sentient beings inhabiting those worlds to discover the locally prevailing laws or to study them.

Metaphysical theories of lawhood may be grouped per multiple criteria. Two criteria against which most theories may be classified are whether they propose laws of nature to be real, mind-independent, entities or not, and whether they may be reduced to something more ontologically fundamental or not. The first of these criteria addresses the realism issue; the second the reductionism issue.

The result of applying these two criteria to the extant literature should be four categories of theory type: realist reductionist, realist non-reductionist, anti-realist reductionist and anti-realist non-reductionist. However, anti-realist theories which deny the real, mind-independent, nature of laws appear to be universally committed to a reductionist stance towards lawhood, explaining away laws in terms of a cognitive agent's system for interpreting the world, for example. It may be argued that a position which can be called expressivism regarding laws of nature could be formulated where the attitudes towards the patterns observed in the Humean mosaic are somehow irreducible, thus representing a counterexample to the reduction of anti-realist laws. Perhaps so, but this would appear to

require a very particular theory of mind. In any case, for the purposes of this thesis at least, anti-realist theories will be construed as universally eliminativist in character.

However, the converse does not hold that realist theories are all committed to laws being non- reducible to something more ontologically fundamental. That is, not all realist theories hold laws to be ontologically primitive.

Among the most well-known and well developed eliminativist theories that deny the mind-independent existence of laws is van Fraassen's constructive empiricism (van Fraassen, 1980). Constructive empiricism states that scientific laws are not literally true – in the sense that they are real, mind-independent entities – but that they may nonetheless be empirically adequate. 'Empirically adequate' is intended by van Fraassen to avoid the metaphysical baggage implied by "true", for it is measured in terms of the degree of concordance or fit between our empirical observations of the world and the predictions of the theory. This degree of concordance, or fit, is all that is required to evaluate a scientific law.

The notion of empirical adequacy is closely related to that of verisimilitude or truthlikeness, after Popper (1976), whereby – in at least some cases of scientific theorizing – our theories approach, perhaps asymptotically, a perfect fit with the observable data. The key difference is that truthlikeness presumes that there is such a thing as true scientific laws for our own scientific theories to approximate; empirical adequacy makes no such assumption.

Although constructive empiricism does not claim scientific laws can ever be true in the same sense that – for example – an observation of a cat on a mat can be true, it accepts that there are regularities out there in the world that can be described and classified and reasoned about. It must accept this, for a theory is only empirically adequate when it

successfully 'latches on to' these regularities. This relationship of 'latching on to' lacks precision of course, but the same may be said of relationships invoked in realist theories, 'supervenes' being the most pertinent example for this thesis.

The relationship of 'latching on to', along with constructive empiricism's acceptance that the world has order and regularities even if it does not have laws that explain them, is often cited as a response to the no miracles argument, after Putnam (1975). The no miracles argument starts with the premise that science is successful in predicting observable phenomena in the world. It is then assumed as a second premise that the best explanation for this success is that science is saying something about the real world. The final premise – often not discussed in formulations of the argument – is that inference to the best explanation is a reliable inference. The conclusion is then that science says something about the real world and that theories such as constructive empiricism that deny this must therefore be false.

The response of the constructive empiricist to the no miracles argument in terms of the 'latching on to' relation seems to work: it proposes a relationship (or at least a vague idea of a relationship) that connects the real world with our scientific theories and laws, and explains why the latter should be able to describe our world so successfully. Van Fraassen (1980) proposes models as the tool we use to latch on to the world, such models being human constructions that depict reality.

> To present a theory is to specify a family of structures, its models; and secondly, to specify certain parts of those models (the empirical substructures) as candidates for the direct representation of observable phenomena (van Fraassen, 1980; p.64).

The proposition that human constructions – scientific models – connect our observations to reality, and explains why the empirically adequate laws we derive from those models should not be miraculous, goes some way to responding to the no miracles argument. But

there is another much more difficult objection from the realist against anti-realist theories such as van Fraassen's constructive empiricism. The objection is closely related to the no miracles argument, but appeals instead to the explanation of the world in terms of the laws, rather than vice-versa. That is, anti-realist theories cannot explain why the world is the way it is, since laws must be conceived of as impotent to determine the course of nature and the observations we make upon it. In which case, it becomes miraculous not that our laws and scientific theories should be so accurate, but that we observe such regularity and order in nature, as opposed to randomness and disorder.

Accepting this version of the No Miracles argument forces one to adopt a governing theory of lawhood. Such a theory states that laws of nature are not merely descriptions of the underlying observations or explanations of them, but somehow determine them and restrict what is nomologically possible and what not. Clearly, this is a perspective inconsistent not only with anti-realist theories such as constructive empiricism, but also for the Humean who does not adopt modal realism. Humean supervenience with modal realism and the plenitude of concrete possible worlds can explain the way the world is, simply because every possible way for a world to be exists. Some of these worlds are jumbles of random events, others not.

Another way to respond to arguments in support of governing accounts of lawhhod without appeal to modal realism is with an unappealing dilemma arising from the question What is the explanation for the laws themselves being the way they are? That is, if the properties of primitive stuff is 'explained' by appeal to primitive laws of nature, this leaves the primitive laws of nature themselves in need of explanation. The response to this must be to either stipulate that it is of their nature to be primitive and in no need of explanation, or that they are in need of explanation by some other yet higher order objects, resulting in an infinite regress of explanation. Accepting that it is their nature to be primitive and in no

need of explanation, the governing theories of lawhood fail to achieve the same level of ontological parsimony as theories – such as the version of the BSA to be developed in this thesis – which have the same explanatory power but without the addition of primitive laws of nature. That is, unlike the non-reductive realist who has to posit the existence of both laws and the things that they govern, and also unlike the eliminativist anti-realist who simply denies the mind-independent existence of laws, the BSA theorist (who avoids positing primitive perfectly natural properties a la Lewis) need posit only the configuration of primitive stuff as it appears in the Humean mosaic. Thus, qualitatively speaking, such a theorist obtains both explanatory equivalence with the non-reductive realist, while having a more parsimonious ontology.

The version of the BSA to be developed in this thesis also avoids the horn of the dilemma ending in an infinite regress by accepting that primitive stuff (in the language of Lewis 'points or points-sized objects') can have properties independently of other things. This, in fact, will be shown in Chapter 6 of this thesis, where a concept of supervenient perfectly natural properties will be introduced whose characters are determined by the configuration of primitive stuff throughout all of modal space. This is in contrast to Lewis who actually had to introduce the notion of primitive perfectly natural properties into his ontology.

The non-reductive realist who adopts a governing account of lawhood may accept the dilemma above, but proceed on the basis of 'whatever works!'. Such a line of reasoning may be to note that the argument above in favor of the BSA makes appeal to ontological parsimony on the basis that 'explanation' means 'reduction' in some sense, and thus that reductive theories are more parsimonious than non-reductive theories. Yet, the non-reductive realist may accept their more extravagant ontology on the basis that no reductive theory of lawhood works. 'Works' here may be expounded upon by citing numerous

problems with reductive theories of lawhood, some of which can be more easily dealt with than others. It is beyond the scope of this thesis to explore all of these problems and consider the non-reductive position in greater detail. Nonetheless, this thesis will go some way towards defending the reductive realist position by showing how at least one reductive realist theory – namely, the BSA account – can be made to work by proposing solutions to four of the problems that have been raised against it, where these solutions all stay firmly within the realm of Humean metaphysics.

The version of the best system account of lawhood that will be considered canonical throughout this thesis is that developed by Lewis (1980, 1986, 1994) as one of the key theories in his overarching program of Humean supervenience.

The BSA as developed by Lewis has some advantages over both anti-realist and non-reductionist theories as shown above, but also some problems, many of which are distinct to it alone. The resolution of four of these above-mentioned problems will be the subject matter of this thesis, from chapters 3 onwards.

1.3 Humean Supervenience

Humean supervenience can be summarized as the doctrine that modality is not part of the fundamental ontology of the world, but supervenient upon it. The form of supervenience appealed to here is global supervenience. Therefore, modality (and by implication all modal notions) are supervenient upon the entire world, considered as an s-dimensional object, where s denotes the signature of the world ('signature' being the number of dimensions and which of those are spatial: 3 or more spatial dimensions plus 1 time dimension in the actual world, for example). Call the class of all properties that are part of the fundamental ontology of the world F type properties, and all modal properties S type properties. Paraphrasing Lewis, global supervenience is then the thesis that there can be no variation in the pattern of instantiation of S type properties without variation in the pattern of

instantiation of F type properties. Humean supervenience is, according to Lewis (1980, 1986, 1994), the thesis that this relationship holds between F type properties and S type properties in the actual world and 'worlds like ours' (Lewis 1994, p474).

With the addition of a more comprehensive logic of possible worlds, it is possible for the Humean to develop more sophisticated supervenient modal notions. Most significantly for the purposes of this thesis, modality comes in many forms, which – for a possible worlds theorist – may be categorized in terms of the relative cardinalities of two mutually exclusive sets of possible worlds.

At one extreme are the modal notions of logical, mathematical or analytic possibility and necessity, where all possible worlds are in the set of logically, mathematically or analytically possible worlds, and all logically, mathematically or analytically impossible worlds either do not exist as per Lewis (1986), or if they do exist in some ersatz sense, are in the set of impossible worlds. To say that, for example, 2+2=4 is mathematically necessarily true, or that 2+2=5 is mathematically impossible, can be recast as saying that 2+2=4 is true in the set of all logically, mathematically or analytically possible worlds; that 2+2=5 is true in precisely zero possible worlds, and so on for logically necessary statements or analytically true statements such as "All bachelors are unmarried".

At the other extreme of modality lies the more restricted notion of technical possibility. Something is technically possible only if the current levels of technology permit it. For example, it is technically possible to build a 50-storey building, but it is currently technically impossible to build a 500-story building. As such, notions of technical modality are inherently relative to a species' level of technological development and are of more concern to the engineer than the philosopher. Nonetheless, the calculus of possible worlds can be used to define the notion schematically: Some engineering feat is technically possible only if it there are possible worlds that share the same laws of nature and level of

technological development as the actual world but where the feat of engineering has been accomplished.

Between these two extremes lies the notion of nomological or physical modality (the two terms will be considered interchangeable throughout this thesis). Something is physically or nomologically possible if it is true in at least one of those worlds that share the same laws as the actual world. In large part, defining this set of possible worlds, and thereby nomological modality, without circularity or brute stipulation is one of the crowning achievements of Lewis's version of Humean supervenience and the BSA. A more in-depth discussion of how this is achieved will be provided in the succeeding section subsequent chapters of this thesis. For the moment, some crucial notions of Humean supervenience will be introduced.

The notion of nomological modality appears to have been the original target of the Humean's 'no necessary connections' doctrine first argued for by Hume himself (Hume 1955, p77). More specifically, the doctrine restricts modality as an ontological primitive in that it says nothing about supervenient notions of modality. Whether the Humean should also object to so-called metaphysically necessary connections – such as those that make mathematical statements true – is an issue that can be safely disregarded for the purpose of this thesis, since the concern is exclusively with nomological modality.

In addition to clarifying the various notions of modality under Humean supervenience, clarification is also needed of what exactly constitutes the F type properties upon which such modal notions supervene. The most that Lewis says on this is that they are perfectly natural properties (PNPs), which according to Lewis are to be understood as properties instantiated by 'points, or point-sized occupants of points' (Lewis 1994, p474). This insistence that PNPs be instantiated only by properties of points or point-sized objects seems unnecessarily restrictive, as the properties of a fundamental theory could be of

finite-sized objects. For example, consider gunky or uninteresting possible worlds in Schaffer's (2003) sense, where there are no such things as the fundamentally smallest components of the physical world, only an infinite descent into infinitely divisible matter where each successive division looks much like the preceding one. In such worlds, the best systematization would have to be in terms of finite-sized objects such as the smallest finite regions of such gunky material that display different properties than anything smaller. Moreover, given the failure of additivity in such infinitely divisible worlds some mathematical formulations of laws as traditionally conceived could fail to work (Arntzenius, 2008).

In any case, a revision of Lewis's original formulation of PNPs will be given in chapter 6. For the moment, it suffices to accept Lewis's formulation and say that PNPs are instantiated by properties of the points or point-sized objects in the s-dimensional manifold of each possible world. Lewis's term for this s-dimensional manifold is the *Humean mosaic*. Although the analogy is misleading insofar that mosaics are typically 2-dimensional objects, whereas the actual world is at least 4-dimensional, the term has become established and so will be retained throughout this thesis.

The final clarification is of the term 'worlds like ours'. In this context, Lewis used the term merely to demonstrate the theoretical viability of Humean supervenience as a metaphysical doctrine. That is, because it did not, according to Lewis, imply any logical contradiction, it is viable as a metaphysical thesis (Lewis, 1994). But by the common understanding of 'metaphysical', a metaphysical thesis must be one that is true in all possible worlds, not just some subset that can be described as being 'worlds like ours'. One aim of this thesis is to demonstrate that the BSA is an account of lawhood in all possible worlds, and therefore constitutes a metaphysical theory of lawhood. Therefore, Humean supervenience will henceforth be assumed to be a stronger, metaphysical thesis, dubbed

here as *strong* Humean supervenience. Strong Humean supervenience is a modal strengthening of Lewisian Humean supervenience, henceforth referred to here as *weak* Humean supervenience. Weak Humean supervenience may be defined as follows:

- Weak Humean supervenience: Fundamental laws of nature, dispositions, and chances supervene on the distribution of fundamental intrinsic local properties in that world (past, present, and future).

Strong Humean supervenience is then the additional thesis that:

- Strong Humean supervenience: Weak Humean supervenience is true in all possible worlds.

This strengthening of Weak Humean supervenience implies that all F type properties are local and intrinsic in all possible worlds. That is, they are all denoted by unitary predicates that make no reference, implicit or otherwise, to any points or point-sized objects. This is another way of restating the founding doctrine of Humeanism: that there are no necessary connections in nature. This is because such a relationship between two or more fundamental points or point-sized objects comprising the Humean mosaic would be a necessary connection in the following sense: the properties of one point or point-sized object would necessarily depend upon those of another point or point-sized object.

Strong Humean supervenience is best considered a generalization of weak Humean supervenience. It states that not only in *worlds like ours* are there no necessary connections because all (fundamental) properties are local and intrinsic, but that this is true in all possible worlds. Alternatively, it states that all worlds are worlds like ours. This generalization is important for any metaphysical theory of lawhood, since it is reasonable to think that metaphysical truth should be true in all possible worlds.

Apart from strong Humean supervenience being a modal generalization of Lewisian Humean supervenience, the above definition also makes explicit that the thesis applies to

fundamental properties and particulars. Although Lewis himself often neglected to make explicit that his version of Humean supervenience applied to fundamental particulars when formulating his theses of recombination, for example, his reference to points and point-sized objects makes this restriction implicit. Hume himself however, seems to have considered his requirement for no necessary connections to be general, applicable to the human level of macroscopic objects as well as fundamental particles. Hume was, of course, writing before the time of modern microscopy and modern physics.

It may be asked why the Humean so doggedly insists on all fundamental properties being intrinsic. Other views – such as structuralism, most recently proposed by Ladyman and Ross (Ladyman, 2007) – that relax this requirement are available, after all. The reason can be found in the parsimony of the Humean approach over that of a structuralist one: a structuralist who allows fundamental necessary connections into their ontology needlessly stipulates the existence of such necessary connections. This is needless because, according to a Humean, they can always be analyzed into simpler, intrinsic properties. Therefore, they should not be considered fundamental and need not be admitted into a parsimonious ontology. This, after all, was the great insight of Hume: that any statement invoking a necessary connection such as “Billiard ball A hit billiard ball B and *caused* it to move” is empirically equivalent to one that invokes no such necessary connection, such as “Billiard ball A moved from point X to point Y, and then billiard ball B moved from point Y to point Z”, and so on. For the Humean, the prohibition on necessary connections is, ultimately, the application of Occam’s Razor to fundamental ontology.

A similar line of reasoning applies to the prohibition on non-local properties, which apply to two or more spatiotemporally separated (i.e. non-local) points or point-sized objects. Such non-local properties are extensionally equivalent to a non-unitary predicate, where

multiple points or point-sized objects are required for instantiation of the corresponding non-local property.

Despite this, it has often been claimed that this prohibition on non-intrinsic, and by implication non-local, predicates violates established observations in quantum physics. Such observations are predicted by quantum theory and under many interpretations require that the properties of one particle depend on the properties of another spatiotemporally distinct particle. This is thought to imply precisely such necessary connections between the quantum particles that are prohibited by a Humean metaphysics (for example, Loewer, 1997; Karakostas, 2009; Butterfield, 2011). This topic will be explored in more detail in chapter 5, where it will be argued that such necessary connections do not violate the Humean 'no necessary connections' principle since they remain supervenient.

Other notions often associated with Humean supervenience are the principle of recombination, similarity relations between worlds, and modal space. Regarding recombination, it is a consequence of there being no necessary connections in the fundamental ontology of any world, that all possible worlds can, according to Humean supervenience, be considered as re-arrangements of the Humean mosaic, or parts thereof, of any other world according to some principle of recombination. Every possible re-arrangement of the Humean mosaic of any possible world is then a possible world itself, resulting in a set of all possible worlds.

Lewis (1986) uses a Humean principle of recombination and the set of all possible worlds it generates as the basis of a theory of modality. This will be explored in much greater detail in the next section for the case of nomological necessity. He also used it as the basis of counterfactual reasoning by imposing on this set of all possible worlds an ordering relation, based on the degree of similarity to the actual world, or some other index world. To avoid

the situation where some world may be more similar to the actual world than the actual world itself, the notion of strong centering is also often stipulated as an axiom, which states that no world is more similar to the actual world than the actual world itself.

The resulting Stalnaker-Lewis theory of counterfactuals is tangential to Lewis's theory of modality and the BSA and will not be considered further, but the ordering relation imposed on the set of all possible worlds is particularly relevant for the refinements to the BSA that will be proposed in subsequent chapters. Coupling the set of all possible worlds with the ordering relation we get a space of possible worlds. This space of all possible worlds shall be referred to as modal space, $\mathcal{M}$, throughout this thesis. $\mathcal{M}$ provides not only the basis of a more formal definition of modal notions, but also a way to more easily conceptualize the different degrees of modality mentioned above. Specifically, each can be considered a subspace of worlds within a more encompassing region of modal space (see Figure 1.1 below). For example, all possible worlds that share the same laws of nature and nomological characteristics as the actual world – a characteristic that will henceforth be referred to as isonomic – may then be defined as a subspace within $\mathcal{M}$. A proposal for how this will be defined using the BSA to establish the nomological character of each world, in a manner that avoids some of the problems with Lewis's approach, will be the main objective of this thesis, developed in chapters three onwards.

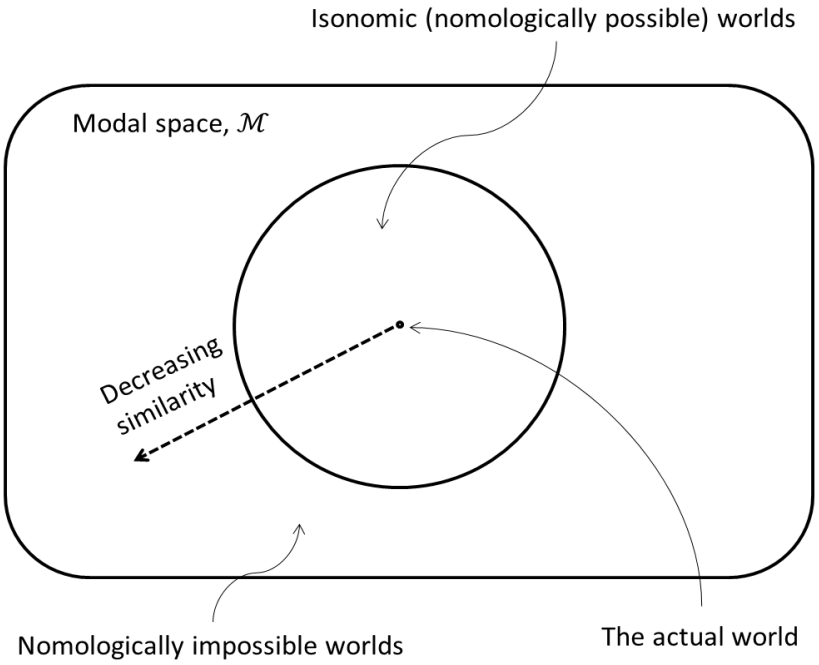


Figure 1.1: Depiction of the actual world located in modal space with strong centering and an ordering similarity relation between worlds.

Lewis (1986) famously adopted one other thesis regarding modal space, which will be here referred to as concrete modal realism (Lewis 1986, p81). That is, every point in $\mathcal{M}$ corresponds to some possible world, and vice-versa, and every possible world has a mind-independent existence in the same manner as does the actual world. Call the first of these requirements the completeness condition, and the second the concreteness condition. The concreteness condition can be equivalently stated as: every possible world is as concrete as the actual world.

Concrete modal realism is often considered an extreme position that is unmotivated. In fact, it logically follows from a principle of recombination that all possible worlds will be as equally concrete as the actual world if they are themselves formed from re-arrangements of the same stuff as the actual world, which is assumed to be concrete. 'the same stuff' is not here meant to denote the token of stuff – whatever that may ultimately discovered to

be – that constitutes the actual world, but rather the same type of physical stuff in the broadest sense of 'physical'.

Lewis (1986) also discusses some arguments for the position, but it is not the purpose in this thesis to defend the position or discuss it in any detail. Instead, wherever possible, reference will be made simply to "modal realism", allowing for equivocation on the issue of concreteness. In one instance however – specifically, in chapter 4 – it will be necessary to deal with the issue head on and presume concrete modal realism as part of the theory, at which point this will be made explicit.

In brief, throughout this thesis, Humean supervenience will be adopted with the Lewisian notions of modal space and of modality left largely intact, along with concrete modal realism. Humean supervenience thus understood comprises three inter-related theses, as follows:

- **No necessary connections**: There are no necessary connections between fundamental properties. That is, all particulars are fully specified by their intrinsic fundamental properties – properties that cannot differ between duplicates – and all worlds are built up from such particulars according to some principle of recombination consistent with this requirement.
- **Concrete modal realism**: Via some principle of recombination, any logically consistent proposition involving fundamental properties is true of at least one point in $\mathcal{M}$, and each point in $\mathcal{M}$ is called a possible world, and all possible worlds are as concrete – whatever the term 'concrete' means is unimportant – as the actual world.
- **Nomic necessity**: For any world, the set of fundamental laws of that world, L, define some subspace of $\mathcal{M}$ as the set of all nomologically possible worlds that share the same laws L, and there is not some other set of laws L^* which defines an overlapping subspace of $\mathcal{M}$.

It will not be the purpose of this thesis to defend any of these claims, but to presuppose them as the metaphysical setting within which the BSA can be made to work as a viable

realist theory of lawhood. The BSA may then be added as a fourth claim to complete an overarching metaphysical scheme for Humean supervenience, introduced as follows, and expounded on in the next section:

- Best-system account of fundamental laws (BSA):
  - I. The set L of fundamental laws of any world collectively represent the best systematization of the Humean mosaic of that world, and
  - II. the entire history of the world, considered as an s-dimensional object, comprises the Humean mosaic and
  - III. the best systematization is measured in terms of the strength, simplicity and (in the case of indeterministic worlds) fit, in comparison to alternatives.

1.4 Best System Account (BSA)

Of all explanatory philosophical theories of lawhood the best system account (BSA) is the frontrunner in terms of being able to provide an intuitively acceptable account of fundamental laws without stipulating such laws as ontological primitives. The theory was conceived originally by Mill (1904) and Ramsey (1990) and later developed further by Lewis (1980, 1986, 1994), and for that reason is sometimes also referred to as the MRL theory of lawhood.

The generic form of the BSA states that laws are the axioms or theorems of a best systematization of observable phenomena. The idea is that of a competition between different sets of propositions, each set being a candidate for the best system. The winner (or winners in the case of a tie) of the competition is that set which is best. The propositions of the winning set, along with all propositions logically derivable from them, are then the laws of nature.

The meaning of 'best' in the case of deterministic possible worlds Lewis defines as the optimum trade-off between the two concepts of *simplicity* and *strength* (Lewis 1994, p480). When considered as a realist theory of lawhood, reliance on the concepts of simplicity and strength implies that these, as well as the exchange rate between them, also are objective, mind-independent, measures.

Simplicity is relatively easy to define: the length taken, relative to a language L, to express the best system. 'Length' here could denote the number of characters in binary, or just the language L directly, just so long as propositions comprising all candidates for best system are expressed in the same language to make comparison meaningful. Unfortunately, this definition of simplicity is language relative, and thus fails the criteria for being objectively real. It will be shown in Chapter 6 however that an objective, mind-independent, language can be defined whose predicates designate PNPs, and thereby retains the BSA as a realist theory.

Strength likewise can be defined objectively as (for example) using the information content of the system, which could be quantified in terms of bits using Shannon's Information Theory (Shannon 1948). In this case the Humean mosaic of a world w will have associated with it a maximal quantity of information, corresponding to a complete description – in terms of the PNPs of the objective, mind-independent language of the BSA – of every point in some co-ordinate system encoding the Humean mosaic of w. Call this quantity of bits H_{MAX}. Strength can now be defined as that proportion of H_{MAX} that matches the predictions of a candidate best system. According to this measure, the strongest candidate would have strength = 1. But this system may require a complete description of every point, and would thus fail in terms of simplicity. A simpler system would likely have strength < 1, but would be so much simpler that the loss of strength is trivial in comparison. This trade-off between simplicity and strength is commonplace in science, particularly the statistical sciences.

Insofar that this is true, the BSA gains some credence by deploying the same kind of reasoning and logico-mathematical concepts as empirical science.

The trade-off can be depicted graphically with a hypothetical dataset considered to represent in abstract form the patterns of in the Humean mosaic of some world w. Such a dataset is shown in Figure 1.2 below. In this toy world, assume that the only regularity is the relationship between two variables X and Y. This relationship can be most simply described with the equation X = Y. But the degree to which this equation and its corresponding line in a Cartesian plane depicted in Figure 1.2 matches the actual data is low compared to alternatives; it is simple but not strong. In particular, there is some n-term polynomial (n>1), which would have maximum strength; it describes every data point in the Humean mosaic. Yet this n-term polynomial will be a far more complex equation than the simple X = Y. These two extremes represent the simplest and the strongest systematizations of this Humean mosaic. But neither could be considered the best systematization, which trades off decreased strength for greater simplicity. In the example shown in Figure 1.2, this is represented by a simpler linear expression of the form Y = cX; an expression slightly longer than 'X = Y', but far shorter than the n-term polynomial with maximum strength.

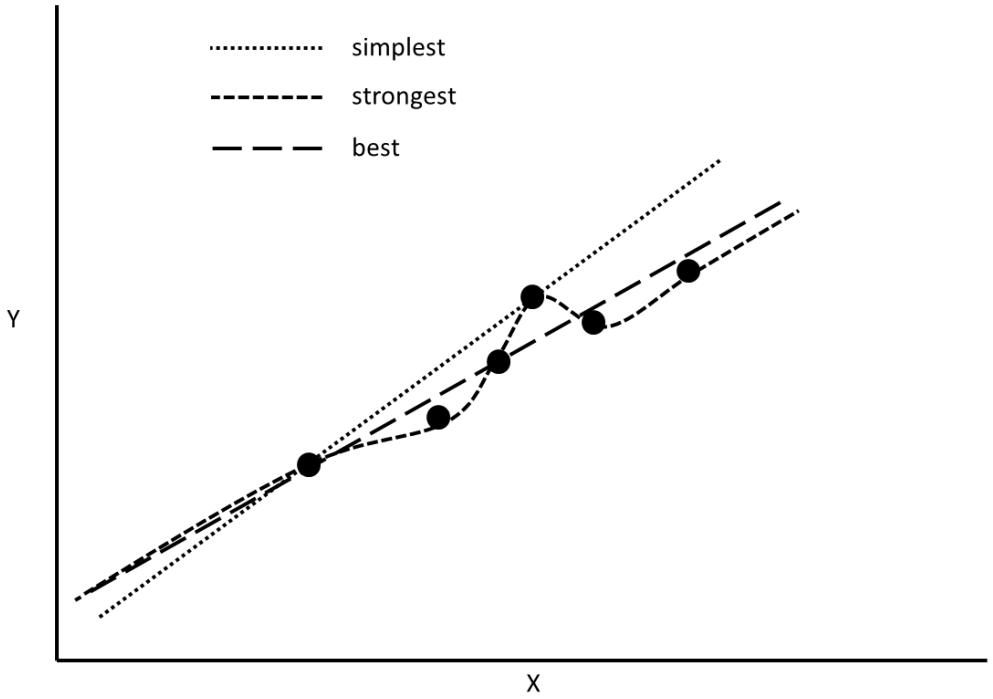


Figure 1.2. Graphical depiction of the trade-off between simplicity and strength for a two-dimensional Humean mosaic.

Lewis (1980) introduced fit as a third dimension to allow the BSA to work also for indeterministic worlds where at least some points in the co-ordinate system of the Humean mosaic evolve through time in a manner best described by a probability function. 'Best' here is essentially the same concept as above, except that the trade-off is between a deterministic lawlike expression and one comprising at least one (non-trivial) probability function.

Thus, in indeterministic worlds, the trade-off is not only between simplicity and strength, but also fit. Exactly what is being fitted, to what, and how fit is measured, will be discussed in chapter 3 of this thesis.

Lewis does not mention whether he considers the inderminacy of a world to be supervenient or not, and therefore whether it is possible to establish a priori or not

whether fit needs to be added to the criteria to identify the best system. Lewis (1980) defines indeterminacy using the standard definition of it sharing some same initial segment with some other world, but diverging from that other world at a later time. But every world shares every one of its segments with an infinitude of other worlds before diverging. Only an infinitesimal proportion of those diverging other worlds can establish the indeterminacy of some index world. Namely, that proportion that are isonomic. But the property of being isonomic is determined by the best-systematization process. Therefore, the best-systematization of an indeterministic world would presume its indeterminacy prior to the systematization process involving fit as a third criterion.

This is problematic for the Humean, who would prefer not to introduce primitive properties of the Humean mosaic such as indeterminacy. In chapter 3 of this thesis it will be shown how the indeterminacy or determinacy of a world can be determined by the best-systematization process itself, rather than being presumed as a primitive property.

The BSA as it has been developed by Lewis as a central thesis in his Humean supervenience program has the specific task of identifying only the fundamental laws of nature, and thereby the relations of nomic necessity obtaining in each world by the *No necessary connections* thesis above. Thus it plays a role beyond that of identifying the fundamental laws of any given world by establishing the supervenient modal properties of that world.

In addition to establishing nomological modality, the BSA presented by Lewis is also an account of fundamental chances: the best systematization of a world may comprise only deterministic laws consistent with only a single course of history given a single initial state, or at least one indeterministic law that is consistent with multiple courses of history. Fundamental chances are then defined by the probability functions appearing in – or implied by – any indeterministic laws.

The BSA may be considered in historical context as a refinement of the so-called naïve regularity account (NRA) of lawhood, developed by Hempel (1965), among others. In essence, the NRA acknowledges that there are regularities and order in the world, but denies that these are the result of some ontologically independent set of laws that determine those regularities and order. Instead, the NRA equates the laws of nature with those regularities. That is, a law of nature is any regularity in the Humean mosaic. For example, the observed regularity that all ravens are black, is thereby a law, which states that all ravens are black, and so on.

Multiple arguments from Armstrong (1985), Hall (2015), Tooley (1993) and others, have been brought against the NRA and have conclusively led to its rejection as a viable theory of lawhood. Among these are the problem that the NRA does not allow for any laws that are not instantiated as regularities, that it allows for 'accidental' regularities to be counted as laws, and that it does not support counterfactual analyses in a satisfactory manner. These problems are worth mentioning because the BSA overcomes many of these problems, and as another Humean theory of lawhood, may be considered as a successor to the NRA. The BSA overcomes many of the problems of the NRA because it is not equating laws with regularities, but instead only those regularities – whether instantiated or not – that would go into the best system for a world. Importantly, the regularity need not be instantiated, but nonetheless considered as a law if it is implied by other regularities, for example, or if its inclusion completes a pattern of other regularities, and so on.

Lewis requires that the theorems of the best deductive system are all true. But this would exclude almost all of the propositions that have the best balance between information content and length ("strength" and "simplicity" in Lewis's terms), a point that has also been made by Hoefer (2007). This is because the specification of boundary conditions or the encoding of endless exceptions and exclusion of confounders and outside influences are

frequently the source of complexity in (deterministic) lawlike statements. The lawlike statements most common in science are best understood as idealizations of isolated systems where there are no perturbations from the environment and measurement is perfect. But this never pertains in the actual world, and consequently almost never will any lawlike proposition be true in the literal sense; it will at best have some degree of verisimilitude or fit with the patterns in the actual world that is better than chance.

The fact that the BSA as conceived by Lewis, as well as throughout this thesis, concerns itself exclusively with only the fundamental laws of nature is important to note. This version of the BSA says nothing about the laws of the special sciences, so long as they are reducible to the fundamental laws of nature. For convenience, the modifier "fundamental" will often be dropped for brevity, and so it should be assumed that whenever laws of nature or lawhood is discussed in this thesis it is the fundamental laws of nature or fundamental lawhood that is intended.

Although the BSA solves many of the problems that have dogged the NRA, many problems remain to be properly addressed. The four problems introduced in section 1.2 that are the main focus of this thesis have specific chapters dedicated to their resolution, and shall be referred to as the zero-fit problem, the undermining problem, the non-locality problem and the immanent comparisons problem. Each of these problems will be expanded upon in chapter 2.

1.5 Other problems with the BSA

In this section, some other problems with the BSA beyond the scope of this thesis will be introduced.

1.5.1 The empiricist loyalty test

The empiricist loyalty test as specified by Earman (1986), is for philosophical theories to be epistemologically accessible. 'Accessible' here means open to empirical verification, at least in principle. The BSA theory of lawhood developed by Lewis is not epistemologically accessible in principle. As van Fraassen (2002) argues, it is conceivable for a world to have a best systematization using predicates that are not perfectly natural properties (PNPs), and we would be unable to tell the difference. More generally, we can never know, at any point in human development, if the best systematization that real world scientific endeavor has achieved is the best systematization as defined by Lewis's theory. This is because, for any predicate appearing in any candidate best systematization, we can never know if the predicate designates a PNP.

This is not a fatal argument against Lewis's BSA theory, but to have epistemological accessibility is a desirable property of any theory of lawhood. Compared to non-reductionist accounts, Lewis's BSA theory still wins out since it at least conforms to the current ideals of human science: to express in as simple manner as possible the most widely applicable generalizations as possible. But for the strict empiricist this is not good enough.

Cohen and Callender (2009) rejects Lewis's BSA theory on this basis. But their rejection appears emotional, arising from some unstated principle of meta-philosophy, rather than any direct appeal to reason. Failing the empiricist loyalty test is not sufficient reason to reject it, although passing the test is desirable. They reject non-reductionist theories of lawhood such as Armstrong's (1985) account for the same reason: under both accounts entities are postulated – necessitation relations in the case of Armstrong, PNPs in the case of Lewis's formulation of the BSA – as necessary conditions for scientific laws, yet remain necessarily inaccessible to any epistemic agent.

Whether any realist metaphysical theory can pass the empiricist loyalty test is beyond the scope of this thesis. Also beyond the scope is a more detailed assessment of whether or not Lewis's original formulation of the BSA, or variants such as Cohen and Callender's (2009), or that developed in the subsequent chapters of this thesis, would pass such a test. For present purposes, it is noted that while these variants of the BSA fail the empiricist loyalty test, such a failure is not fatal to the theory itself.

It is worth noting that Lewis himself accepted the epistemological problems with his theory. His solution was to propose a humility thesis (Lewis, 2004). This applies not only to knowledge of the PNPs of his original theory, but also of his wider belief in modal realism. As he said, the possible worlds of his modal realism are not things that can be viewed with a telescope (Lewis, 1986).

1.5.2 The non-convergence problem

Empirical indistinguishability in the philosophy of science refers to two or more theories that describe the same phenomena with the same degree of precision and with equal predictive power, but where the theories are not mutually reducible to one another. That is, there is at least one theoretical term in at least one of the theories that has no counterpart in any of the other theories. 'Counterpart' here means theoretical equivalence, demonstrable algebraically, for example.

The empirical indistinguishability of relevance to this discussion is of a different form than theoretical irreducibility. In a competition for the best systematization with two candidate theories, T1 and T2, that are empirically indistinguishable in the sense described above, there may still be an outright winner. Namely, that theory that is expressed more simply. But it may be metaphysically possible for T1 and T2 to be equal also in the simplicity of their expression. In such cases, there would be a tie for the best systematization and the BSA theory of lawhood would not converge to a single best systematization. Such

situations represent incommensurability, in the Kuhnian sense (Kuhn, 2012), whereby two theories have equivalent logical strength and simplicity of expression, yet remain incommensurable.

In such cases of non-convergence rules could be introduced to pick a winner. But, given the equivalence in all relevant respects specified by the BSA such rules would be arbitrary. Assuming that in such cases a definitive single set of laws would be desired, a better response to the problem is required.

Lewis himself admits of the possibility of there being no clear winner of the best systematization contest (Lewis, 1994). He puts blind faith in the 'kindness' of nature to provide a clear winner. Implicit in this act of faith is that worlds may not admit to unique best systematizations, resulting in perhaps a multiplicity of systematizations, each one defining a set of laws of nature on equal footing.

Cohen and Callender (2009) claim that there are real-world instances of empirical indistinguishability. But the example they provide to support this is the multiple interpretations of quantum mechanics. This is not a good example, since the differing interpretations are not equivalent in theoretical simplicity: the Everettian interpretation is simpler than the wave-collapse interpretation, for example, in terms of not positing additional axioms to calculate probabilities of measurement outcomes. A better example is a reformalization of Newtownian gravity (Newton-Cartan Gravitation; NCG), which Bain (2004) has shown to be empirically indistinguishable and formally irreducible to classical Newtonian gravitational theory. Still, the issue remains of whether NCG has the same level of formal complexity as the original formulation.

None of these results demonstrate the truth of non-convergence as a metaphysical possibility. The evidence from the actual world is weak. But the BSA theory of lawhood has

to be a metaphysically necessary feature of all possible worlds. Insofar that non-convergence is a mathematical possibility, it may be a problem in at least some possible worlds. A proponent of the BSA must therefore bite the bullet and accept that, in some worlds at least, there may be more than one best-systematization, and therefore more than one set of fundamental laws. This is effectively Lewis's position, although he does not commit to the strong version of Humean supervenience defined above.

1.5.3 The non-denumerability problem

By 'non-denumerability' it is meant the infinite divisibility of space or time (either would do). Or equivalently, the measure of space or time along the continuum of real numbers. The zero-fit problem, which will be discussed in chapter 3, as it applies to non-denumerable worlds may be considered one manifestation of the various problems non-denumerability causes for the BSA.

The infinite divisibility of space or time is a problem because, at whichever spatial scale the systematization is made, there will always be a smaller scale at which an alternative – and perhaps better – systematization could also be made. As Schrenk (2009) has noted, in such a non-denumerable world, physics can never be complete.

Of note, there is no obvious reason why alternatives to the Humean view would not also suffer from non-denumerability. One such alternative is Ladyman and Ross's (2007) version of ontic structural realism. This involves 'relations all the way down' and could be used as the basis of a structuralist version of a BSA.

Without denying the possibility of non-denumerable worlds or accepting defeat and proposing the BSA as a theory only for denumerable worlds, two arguments may be made which support the applicability of the BSA in non-denumerable worlds.

The first argument relies on inherent limits to measurement. In the actual world, Heisenberg's uncertainty principle implies a minimum spatial length and length of time beyond which the level of uncertainty exceeds the length. The resulting Planck length and Planck time thus sets an inherent limit on the measurability of the actual world. It is an interesting, though perhaps pointless, question to ask whether there can be any objects or time periods smaller than these measurable limits in the actual world.

It remains an open question – and one beyond the scope of this thesis – whether such physical limits to measurement exist in all metaphysically possible worlds. That is, whether Heisenberg's principle is a fundamental feature of physical reality, or just a peculiarity of worlds like the actual world. If it is fundamental, an empiricist could then quite reasonably declare that all metaphysically possible worlds are denumerable, since what cannot be measured in principle does not exist to the empiricist.

Barring the possibility that all worlds have some limit on measurement, Schaffer (2003) suggests another solution for at least some non-denumerable worlds. He has argued for a supervenience-only form of fundamentalism, where every possible world is either composed of atoms (in the mereological sense where some parts of the Humean mosaic do not contain proper parts), or there is an infinitely 'boring descent', where beyond some level of magnification or subdivision nothing interesting happens scientifically. So the fundamental level at which the laws of nature are to be formulated would be at that level beyond which nothing interesting happens.

1.5.4 The 'special sciences' problem

A special science is here intended to be any science other than the fundamental science. That is, all other sciences that study more macroscopic patterns in the Humean mosaic that are described by predicates other than those denoting PNPs. In other words, special sciences are reducible in some meaningful sense to the fundamental science. The 'special

sciences' problem as it relates to the BSA may then be characterized as the opposing view, that the special sciences are not reducible in some meaningful sense to fundamental science.

The literature on reduction in science is vast, and it is beyond the scope of this thesis to argue for the reducibility of the special sciences. One important thing to note however is that 'special sciences' are only problematic for the Humean if they can be demonstrated to be non-reducible in a very strong sense of not being supervenient upon more fundamental science. If they are irreducible in some weaker sense, then the Humean should have no problem as the BSA is proposed only as a theory of fundamental science and makes no claim on the status of the special sciences except that they should be supervenient on fundamental science. Specifically, the global supervenience relation adopted by Lewis (1980, p290) does not obviously necessitate reduction. Kincaid (2004) has claimed that the special sciences (as we intuitively conceive of them) have distinct causal patterns for example, and that this warrants their autonomy.

Ladyman and Ross (2007) go further, suggesting it may be the case that in fact only the special sciences have the concept of causation – that as a supervenient concept it does not itself exist in the supervenient base. Irrespective of whether or not having a distinct notion of cause, this does not itself imply metaphysical independence. This is particularly true in Humean metaphysics, since causation is – as with Ladyman and Ross – a supervenient concept.

It may even be argued that something like the BSA theory of lawhood proposed by Lewis may be applied also to the special sciences. For example, Cohen and Callender (2009) and Schrenk (2009) have suggested that each special science can have its own best systematization. This is reasonable, if each special science can be objectively delineated.

Unfortunately, it is not obvious how such a delineation may be made. Ladyman and Ross (2007) make two independent claims regarding how special sciences may be individuated from each other. This first claim is that special sciences can be individuated on the basis of their 'scale'. They claim that each special science operates at the same scale: at the level of the economy, or the organism, for example. But is clearly not the case: the science of population biology could study elephants moving across the plains of Africa or the spread of a virus through a single host; the science of microeconomics can study the sole trader as well as a global corporation. These scales range over many orders of magnitude. Moreover, there is no clear way to individuate special sciences even if they are at the same scale: epidemiology and population biology both involve the study of the growth of microbial populations in their host environments, for example. But they remain separate fields of research for no other reason it seems than the intention and context of the research.

Ladyman and Ross's second criterion of individuation is of the institutional or administrative autonomy of fields of research. But this is too mind-dependent to be considered in the context of a realist theory of special science lawhood.

It may be argued that a slight modification to Lewis's formulation of the BSA theory of lawhood could allow for both special science laws and a defensible individuation of special sciences that is grounded in the best system itself. The modification is to replace the term "fundamental laws" that Lewis exclusively considers as being grounded by his account of the BSA with the term "laws". Thus unrestricted by Lewis's requirement to consider only the phenomena of fundamental physics, the best system can encompass all phenomena as a whole. Crucially, if any special science is then genuinely autonomous in some meaningful sense from the fundamental laws, it would have its own subset of laws within all those defined by the best systematization of all (both fundamental and emergent) patterns in supervenient base.

Take the two examples of nuclear chemistry and economics (the former being considered a distinct special science, following Ladyman and Ross (2007)). In the case of nuclear chemistry the derivation of its laws from those of the more fundamental quantum physics is quite straightforward. But in the case of economics, there appears to be no tractable derivation from fundamental physics. If there are bridging laws, in the Nagelian (Nagel, 1961) sense, then they would likely be extremely complex or uninformative, comprising, for example, a huge disjunction of transaction events to define the transaction relation.

But it is in precisely these situations that the trade-off between strength and simplicity was used in the case of the fundamental laws. That is, a best systematization of both economic and physical phenomena, would in such circumstances require economic laws in addition to physical laws since the incorporation of the former would result in a system with great explanatory power that is far simpler than one which attempted to reduce economic laws to those of fundamental physics.

This argument further requires that there would need to be analogues of perfectly natural properties for each special science, just as they were posited for the fundamental science by Lewis (1986, p61). Without them, the best-systematization of special science laws would either result in laws that turn out to be written in the language of perfectly natural properties (of the fundamental science) after all, or they would be trivialized in the same manner as the BSA-defined laws of the fundamental science would be without PNPs.

The open question then is whether it is conceivable to introduce, without making the Humean ontology unnecessarily extravagant, PNPs for the special sciences (thus taking away their specialness in the sense defined above)? Given the example of economics above, it seems conceivable, although as yet unproven.

In chapter 6, a conception of PNPs will be developed that makes this issue moot, since they rely on a supervenient conception of perfect naturalness that makes no assumption about scale or reducibility. Rather, the conception requires only that PNPs imply the least amount of empirical information – defined in a precise sense in chapter 6 – necessary to formulate the fundamental laws of nature (even if they turn out to include laws of economics). In contrast, within Lewis's original (1986) formulation with its primitive notion of perfect naturalness, PNPs are conceived of as applying only to points or point-sized objects.

In conclusion, arguments that the formulation of the BSA theory of lawhood in terms of natural properties cannot account for laws in the special sciences requires a demonstration of autonomy of those special sciences, or at least intractability when required to be expressed using fundamental perfectly natural properties. Alternative formulations of the BSA theory of lawhood that allows for best systematizations for each special science requires objective individuation of each special science, as well as some restriction on the language used to describe the candidate special science laws to avoid the trivialization problem.

1.5.5 The meta-laws problem

A meta-law is a law that generalizes over other laws. They are quite common in contemporary science. For example, the theory of special relativity is effectively a meta law saying that the laws of physics are the same for each observer. More generally, the principles supporting the BSA theory of lawhood may be considered meta-laws too, of the form that laws are informationally efficient summaries of patterns in the Humean mosaic, for example.

There is no principled reason for meta-lawlike statements not to be involved, along with object-level lawlike statements, in the best systematization process, so long as they contribute to the simplicity-strength trade-off in terms of the properties of perfectly

natural properties. However, when framing the meta-laws problems, they have been assumed to be supervenient on the BSA laws rather than a component of them (Lange 2011).

Lange (2011) identifies two problems associated with meta-laws in the BSA. The first is that there may be meta-laws of the form: All laws hold everywhere and everywhen. This sounds reasonable. But, Lange claims, in order to support counterfactuals, small localized miracles must be introduced which contradict this meta-law. Therefore, either the meta-law must be wrong, or the BSA theory of lawhood cannot support counterfactuals.

Perhaps it is best to drop the requirement that the laws and the set of possible worlds that they define must support counterfactuals. In any case, whether the panoply of possible worlds assumed by the BSA is sufficient to support counterfactuals has no bearing on the BSA as a theory of lawhood. Besides, it has been independently argued that they cannot support all counterfactuals in any case (Barker, 2011).

1.6 Non-reductive realism

An appealing feature of the BSA is its concordance with what is presumed to be the ideal of scientific theorizing: to discover scientific laws that explain many different types of observation using a minimal set of parameters and mathematical or logical operators on those parameters. This appeal to the simplicity and explanatory power of real-world scientific laws can be appreciated by both reductionist and non-reductionist theorists. The non-reductionist could even endorse these as epistemic standards of laws. They differ from the BSA theorist in their not endorsing them as metaphysical standards.

The core argument of the non-reductionist theorist is not directed against the BSA specifically, but more generally against all reductionist theories. The purpose of this discussion is not to survey these more general arguments, but since the anti-reductionist

position is the main contender to BSA theory, brief mention of the non-reductionist argument provides context.

In fact, only brief mention is required, since the core non-reductionist argument template is quite simple as has been shown by Beebee (2000) to lie at the heart of many anti-reductionist arguments:

1. Our intuitions tell us that laws govern or explain phenomena, or support counterfactuals, or may not be instantiated, and so on.
2. Reductionist theories of law are inconsistent with these intuitions.
3. Therefore: Laws are primitive.

There are multiple instances of this argument template. For example, Tooley (1993) provides the example of a simple finite particle universe where all but one type of particle interaction is ever instantiated in the universe. Were such an interaction to be instantiated, it would follow the same pattern as other interactions that are instantiated. But, under many reductionist theories there would be no such law since such interactions do not exist in the supervenient base.

This family of arguments against reductive theories of lawhood appeal to our (non-Humean) intuitions that laws of nature are something ontologically independent of the patterns of instantiation that results from them. These intuitions seem to be strong, perhaps as the result of evolution in the case of human cognition. But to claim that human intuition is an accurate guide to truth is not supported by the evidence. There are many examples – for example from behavioural economics and quantum physics – that require us to disbelieve our intuitions.

Moreover, such arguments for primitive laws open themselves up to an infinite regress since it is then reasonable to ask what explains the laws, for example. If the response of the non-reductionist to this challenge is to postulate that they are not in need of explanation

because they are primitives it is reasonable to question why the underlying regularities or other empirical features of laws could not instead be defined as primitive.

Tooley's particle universe trades on our intuitions about what laws should be. But there are stronger arguments still. One such is by Esfeld (2010), where he imagines a single particle universe. This single particle is massive, but because there are no other particles in the universe there is no way for this mass to be detected or made manifest. In the Tooley example, it is possible that the BSA may infer some law about the non-interacting particle pair. But in the Esfeld single particle example, it is unlikely that the BSA would have any law governing mass; to do so would complicate things unnecessarily and thus be excluded from the best system. But this seems to conflict with our intuitions about the counterfactuals in such a single-particle universe. Namely, the resulting BSA laws would not accurately describe what would happen were another particle to be introduced into the universe. But this intuition, like that generated in Tooley's example, is question-begging: we must assume the truth of a nomological counterfactual, but it is the basis of such nomological possible worlds that are established by the laws of nature (for Lewis at least).

These examples should not be dismissed out of hand, since they clearly show the conflict between our intuitions and what is logically and metaphysically consistent and non-question-begging. But this is all that such arguments do. If someone believes steadfastly in a philosophical theory being required to satisfy such intuitions, that is all that can be said; there is no further reasoning that can be brought to bear on the situation. Lewis (1994) sums the situation thus:

> If you're prepared to agree that theorems of the best system are rightly called laws, presumably you'll also want to say that they underlie causal explanations; that they support counterfactuals; that they are not mere coincidences; that they and their consequences are in some good sense necessary; and that they

> may be confirmed by their instances. If not, not. It's a standoff – spoils to the victor. (1994, p.478)

However, the BSA theory of laws is not defenseless against such attacks. In terms of meeting all but the most extreme requirements of the non-reductionist, it is well placed in appealing to simplicity and strength as justification for providing the basis of explanation, supporting counterfactuals, and so on. This appeal should be more convincing than merely citing the underlying regularities as in the case of the regularity account of lawhood.

1.7 Humean supervenience and causation

Any theory of lawhood is incomplete without some accompanying theory of causation. Along with laws of nature, chance, and nomological modality, causation is the other key nomological concept. Many who have claimed laws to be primitive likewise also claim the same status for causation. Namely, that the relation "C causes E" is real and irreducible. Similarly, those who claim laws to be non-real typically also write-off causation as a folk concept. Most famously, Hume himself (Hume, 1955), as well as Russell (1912), both persuasively argued that our everyday and scientific thinking and theorizing can proceed by replacing "C causes E" with something such as "Whenever C happens, E happens".

That this replacement can be made implies little regarding the reality of causation. For, as with the debate regarding the reality of laws of nature, a distinction can be made between reductive and non-reductive theories, and among the former those that presume causation as something real and those that claim it as nothing more than a psychological phenomenon.

Among those that consider it to be something real and reductive, a sister theory to the naïve regularity account of lawhood may be cited as a predecessor of the theory proposed by Lewis. This is the covering law theory of Hempel (1965), whereby an event c causes an effect e iff there is a regularity linking c and e, and since – according to Hempel – all laws

are regularities, this means that c causes e iff there is a law which covers such causal instances (so-called covering laws).

As in the case of the other nomological concepts, Lewis considered causation to be Humean supervenient (Lewis, 1973a), corresponding to patterns of counterfactual dependence between events using the Stalnaker-Lewis possible worlds account of counterfactuals (Lewis, 1973b). This made it into a broader concept – for Lewis at least – than lawhood, since it allowed for an event to cause another event in the absence of an overarching 'covering' law of the type proposed by Hempel (1965). Therefore, Lewis's theory of causation is largely independent of the BSA. Nonetheless, as a key nomological concept, it is important to introduce Lewis's theory of causation and the problems associated with it to provide context for the wider discussion on the BSA and Humean supervenience.

Lewis's theories of counterfactual causation (Lewis, 1973a; Lewis, 2000) are all attempts to define the singular causal relation (as opposed to general causation, for which Lewis (1973a) adopts a modified regularity approach). Ultimately, this amounts to formalizing in possible world semantics, the following expressions, where c and e are singular events:

1. Event e causally depends on event c iff if c had not occurred, then e would not have occurred, and

2. If e causally depends on c, then c causes e.

In possible world formalism the causal dependence relation between events is expressed as counterfactual dependence between propositions about those events, where 'A □→ B' = 'if proposition A were true, then B would also be true', and 'O(x)' = 'event x occurs', with the following two statements:

1. Event e causally depends on event c iff ~O(c) □→ ~O(e), and

2. If ~O(c) □→ ~O(e), then c causes e

The causal relation is transitive under this interpretation. So, if c causes d and d causes e, then c causes e. Such transitive relationships, when strung together, are referred to as a causal chain. So, in the counterfactual analysis – since a causal chain may be composed of just two events – c causes e iff there is a causal chain connecting c and e.

However, the relation of causal dependency is intransitive because c may cause d, which may cause e, but e may still have occurred in the absence of c.

The intransitivity of causal dependency under this counterfactual theory is thought to solve two related problems of the competing regularity account of causation. These are the problem of effects and the problem of epiphenomena. The problem of effects occurs when an event of type C always causes events of type E, but not vice-versa, a situation which the regularity account cannot distinguish from C events being causally dependent upon E events. The related problem of epiphenomena results when events of type C causes both events of type E and F, but F events occurs after E events, and thus the regularity account cannot distinguish E events not occurring F events.

Although a generalized counterfactual theory of causation may solve the problems of effects and epiphenomena associated with naïve regularity accounts, in its singular form it also results in problems of spurious causation (Hall, 2000). An example from Hall illustrates: consider a terrorist who plants a bomb under a philosopher's desk. The authorities discover the bomb and defuse it. As a result, the philosopher survives. According to the counterfactual theory, there is a transitive causal chain between the terrorist planting the bomb and the continued survival of the philosopher. Thus, we could conclude – spuriously according to Hall – that the bomb caused the continued survival of the philosopher. There are many more examples of this kind of spurious causation in the literature. That these examples reflect only the weaknesses – let alone variability – of our intuitions as arbiters of

philosophical dispute seems the most that can reasonably be said. There are however much stronger arguments against this simple counterfactual theory of causation, all of which are generally considered fatal, and all of which are centered around various notions of redundant causation.

Redundant causation may be defined formally (in the case of two competing causes), using the formalism above, as holding when, for three events c, d, and e, where '$\Box\!\!\rightarrow$' denotes counterfactual implication, three propositions are true:

1. c causes e, and
2. d causes e, and
3. $\{\sim O(c) \wedge \sim O(d)\} \Box\!\!\rightarrow \sim O(e)$.

That is, both c and d may cause e. For example, if Billy and Suzy both shoot someone, either bullet would have killed their target, and if neither Billy nor Suzy would have fired their gun, nobody would have been killed.

This is a problem for counterfactual theories of causation since it can be seen from the above statements that $\sim O(c) \Box\!\!\rightarrow \sim O(e)$ nor $\sim O(d) \Box\!\!\rightarrow \sim O(e)$ is implied by $\{\sim O(c) \wedge \sim O(d)\} \Box\!\!\rightarrow \sim O(e)$. That is to say, it cannot be inferred that if Billy would not have fired his gun, then nobody would have been killed (because Suzy could have fired her gun), nor that if Suzy would not have fired her gun, then nobody would have been killed (because Billy could have fired his). Counterfactual statements of the form "$\sim O(c) \Box\!\!\rightarrow \sim O(e)$" are therefore not true in such circumstances, and therefore cannot be equivalent to statements of the form "c causes e", which in the example above was assumed to be true (i.e. Billy shooting a gun at someone would kill them).

Of note, regularity theories also have difficulties with causal overdetermination, since the minimally sufficient set of conditions necessary to substantiate the generalization from cause to effect includes c or d, but not both.

Multiple examples have been given in the literature of redundant causation (see Hitchcock, 2003 for a review), with variation on when the causal overdetermination happens in the causal chain and the dependency relations that hold between the overdetermining events, for example. While some modifications to Lewis's original (1973a) counterfactual analysis can solve some of these variants, it cannot resolve them all. The solution that Lewis eventually came to is that of causal influence (Lewis, 2000).

According to the causal influence theory, instead of the counterfactual being based on so-called whether-whether counterfactuality – whether event e occurs is dependent on whether event c occurs – it is generalized to all forms of counterfactuality, such as:

- How-when: How event e occurs depends on when event c occurs
- How-how: How event e occurs depends on how event c occurs
- ….

There seems no restrictions on the possible combinations, at least when limiting to whether, when and how.

This account of causal influence is a generalization of a modification to the simple counterfactual theory introduced by Paul (1998), where only when-whether alterations were added. In the presence of slight modifications to the timing of an event depending on the success or failure of a pre-empting cause that are universal, this is all that would be necessary, but there are conceivable cases of counterfactual simultaneity that need also to be considered.

Formally, causal influence states that for an event c to influence an event e:

Let $\{c_1, c_2, \ldots, c_n\}$ and $\{e_1, e_2, \ldots, e_n\}$ be possible 'not-too-distant' alterations of c and e, respectively, where for each e_i there is a c_i, such that:

1. When alterations c_i and e_i do not actually occur: $O(c_i) \mathrel{\square\!\!\rightarrow} O(e_i)$, and
2. When alterations c_i and e_i do actually occur: $\sim O(c_i) \mathrel{\square\!\!\rightarrow} \sim O(e_i)$.

In Lewis's (2000) account he does not equate causal influence with causation, but rather – as with the simple counterfactual theory – the causal relation is the chain of counterfactual dependency between influenced events.

By being able to establish counterfactual dependence between the manner in which the bottle smashes with Suzy's throw as distinguished from Billy's counterfactual throw, cases of redundant causation in which the set of counterfactual modifications $\{e_1, e_2, \ldots, e_n\}$ can be constructed are solved by the causal influence theory.

Yet the causal influence account of causation still fails to adequately deal with all cases of redundant causation. A counterexample, due to Collins (2000), is as follows: imagine a case where the wizard Merlin and the witch Morgana both cast spells to turn a prince into a frog at midnight. Merlin can cast the spell only at a single point in space and time and using a specific incantation that does not vary. Morgana can cast her spell in any manner and at any point in space and time after Merlin and before his spell takes effect. The one difference between the spells is that Merlin's spell trumps Morgana's. In this case, argues Collins, the effect of the spell in no way depends on any feature of Merlin's spell (since Morgana cast the same spell at a later time), nor does the event of the prince turning into a frog vary. Yet, we would still be inclined – according to Collins – that Merlin's spell caused the prince to turn into a frog, not that of Morgana.

There have been other criticisms of causal influence theory. The most substantial of these may turn out to be Collins' accusation that Lewis is actually answering a different question

with the theory of causal influence. Instead of answering the question, What is it for event c to cause event e? he is answering the question, What is it for event c to have a causal influence on event e?, he claims.

The problem of redundant causation seems fatal to counterfactual analyses of the style proposed by Lewis, and so it is not clear whether it is feasible to incorporate causation within the overarching thesis of HS. There does appear to be a viable logic of causation however, developed by McDermott (1995), that is consistent with our intuitions regarding redundant causation. But McDermott's theory requires a huge amount of metaphysical baggage and somewhat arbitrary stipulations as to the axioms of the theory. There are, moreover, many other problems that any philosophical account of causation needs to resolve, as enumerated by Hitchcock (2003), and many of these appearing to be dependent on rather arbitrary intuitions such as those supporting the counterexample of spurious causation above.

With all objections and responses considered, and the option to embrace a non-reductive realist conception of causation as recommended by Tooley (1993) unacceptable as a philosophical theory, the best Humean alternative seems to be an anti-realist stance towards causation. Namely, to adopt the view that causation is just a way human cognition works to segment the world and identify patterns in the Humean mosaic that are important for us.

This position remains consistent with Humean supervenience and the BSA, so long as the fundamental laws of nature can – in all possible worlds – be framed without reference to causation. That this is so is clear from the persuasive arguments of both Hume (1955) and Russell (1912) that expressions of the form "A causes B" can always be reduced to expressions of the form "A and B are in constant conjunction", or something similar.

Moreover, that causation is a mere folk concept is the better explanation given all of the arbitrariness and inconsistencies with our intuitions on the topic.

Therefore, wherever expressions such as 'nomological character' or 'nomological concepts' are used in this thesis, they should be assumed to exclude the concept of causation, which is to be understood as a concept of folk psychology having no metaphysical content or relevance to the formulation of the BSA.

1.8 Preview of upcoming chapters

In the next chapter the four problems that will be the focus of this thesis – the zero-fit problem, the undermining problem, the problem of non-locality and Armstrong's problem – will be introduced and summarized. Then, in chapters 3 through 6, respectively, each of these problems will be considered in more detail and resolved in turn.

With the exception of the non-locality problem addressed in chapter 5, in each chapter the problem will be defined, its proposed solutions that have been published in the literature reviewed and assessed. It will be shown that the published solutions fail in one or more ways to effectively or efficiently to deal with the problem. Finally, in each section a solution will be provided that solves the problem. Each solution will not only be consistent with Humean supervenience, but will also attempt to retain as much of the original Lewisian formulation of it as possible, or – in the case of Armstrong's problem – make it more Humean than Lewis' original.

The order of chapters is not intended to be logical in any sense; it is an arbitrary sequence, since the problems are mostly independent of each other. For the most part it is therefore possible to read them in any order other than that presented here, which is the order in which they are discussed above.

The problem of non-locality will be addressed in chapter 5, although unlike the other chapters addressing one of the four problems, chapter 5 will also include a discussion of the relationship between the quantum mechanics and the BSA. Apart from providing the context within which the particular problem of non-locality is set, the discussion also allows for the issue of whether or not the BSA – or at least the principle trading off strength with simplicity – can inform the current debate of which among multiple interpretations of quantum mechanics is most consistent with this principle.

Finally, in chapter 7, the four problems and their solutions will be reviewed and summarized, and the final, revised, BSA presented as the best philosophical account of lawhood currently available.

2. Four problems with the BSA

2.1 Introduction

The purpose of this chapter is to introduce each of the four problems with the BSA that have been discussed in the literature, and for which solutions will be proposed in the following four chapters of this thesis.

2.2 The zero-fit problem

The zero-fit problem specifically concerns those worlds where the notion of statistical fit is used as one of the evaluative criteria in selecting the best system. Such worlds will be all those where the use of indeterministic lawlike propositions have a better trade-off against simplicity than deterministic lawlike propositions. Namely, those worlds whose best systems of laws turn out to include one or more chance-like statements. In such worlds, those systems with better fit to the pattern of occurrence (or non-occurrence) of chance events are, ceteris paribus, more likely to be the best system than those with worse fit.

For the notion of fit employed in the BSA, one system has better fit than another if it more closely matches the actual history of the world than another. Conventionally, this is understood as implying that a candidate system $S1$ has better fit than another $S2$ if $S1$ results in the actual history of the world, $\mathcal{H}$, having a higher chance than the chance of $\mathcal{H}$ according to $S2$. Consider the instance of a simple coin-flipping world w. which consists of 10 coin-flips, 5 of which land head, 5 landing tails, $\mathcal{H}_w$ = {HTHTHTHTHT}, for example. Consequently, according to the best-systematization of w, the probability of landing heads is 0.5. But many alternative histories consistent with a $P(\text{Heads}) = 0.5$ law would include other sequences:

$\mathcal{H}_1$ = {THHTHTHTHT},

$\mathcal{H}_2$ = {THHHTTHTHT},

… as well as ones where, because of chance variation, the relative frequency of heads deviates from 0.5, such as

$\mathcal{H}_{n-1}$ = {TTTTHHHHHH}.

The fit of a $P(\text{Heads}) = 0.5$ law to a particular 1-of-n histories could then be measured as 0.5 x 0.5 x …. n times, or 2^{-n}, which so long as the relative frequency of heads was actually 0.5, and this was not a limiting frequency for an infinite sequence, comes out as a non-zero value greater than any other candidate system with a different chance assignment to the heads outcome. This non-zero value then can be thought of as measuring the likelihood of God putting his hand into a bag of possible worlds, and pulling out the actual world. Actually, this way of thinking is not strictly correct since chancy laws determine relative frequencies of chance events rather than actual histories of chance events. So, rather than an individual world being pulled from the bag, it should be bunches of worlds that all share the same relative frequencies of the chance events described by the candidate best system. For the moment this distinction matters little, for in both cases with $n = \infty$ chance events and $P(\text{Heads}) = 0.5$, then the likelihood of the actual world (or its corresponding equivalence class of worlds with the same relative frequencies) being pulled from the bag is proportional to $2^{-n} = 0$. Namely, the same value as any other candidate system with $0 < P(\text{Heads}) < 1$. This is the zero-fit problem. Similar zero-fit problems may also arise in cases of non-denumerable probability spaces, such as those ranging over continuous time or space. However, the focus in this thesis will be to address the problem as it arises with discrete chance events and their corresponding denumerable probability spaces.

Lewis (1994) was aware of this problem, and suggested the introduction of hyperreal values on top of the reals to allow for infinitesimal probabilities in the measurement of fit.

But Elga (2004) has claimed this option will not work in any systematic way to pick out only the best fitting system. In essence, the problem is that it is not possible to order infinitesimal values and therefore discriminate between one system with an infinitesimal probability value and another with (the same) infinitesimal probability value. In any case, the introduction of infinitesimals is far from an elegant, non-arbitrary, solution.

Briefly, the solution to the zero-fit problem that will be proposed in Chapter 3 is to first recognize that fit should not be measured against one particular history, $\mathcal{H}_w$, but for all histories that have the same relative frequencies of (independent) chance outcomes. Then, fit may be better thought of as a direct measure of similarity between the actual and predicted relative frequencies. This measure of similarity will be a measure of the degree of similarity between the relative frequencies observed, and those that are predicted by candidate best systems. If there is no difference between them, the measure is zero and uniquely identifies a best system when considering only fit; non-zero values identify worse fitting systems.

There is no reason, it will be argued, to measure fit using a probability measure – all that matters is the degree to which (in the parlance of the statistician) the observed relative frequency in outcomes of a sequence of trials matches the expected relative frequency, based on some probability function. The solution is further generalized to cases where there is more than one type of chance event.

As a consequence, standard statistical methods can be used to estimate this difference, without all of the mathematical machinery required to allow for an ordered field of infinitesimals to be defined in additional to the reals.

2.3 The undermining problem

While the zero-fit problem has never been considered as a fatal flaw in Lewis's overarching scheme of Humean supervenience, the undermining problem (Lewis, 1994) has become infamous as a death blow to both the BSA as a metaphysical theory of fundamental chances, as well as all other frequentist theories of chance , and the whole enterprise of Humean metaphysics more generally. A multitude of solutions have been proposed in the literature, although there is a worry that none of them addresses the problem directly and all of its consequences.

The undermining problem manifests itself as both a metaphysical problem for Humean supervenience, as well as an epistemological problem for Lewis's theory of chance. However, in the literature the focus has primarily been the consequences on the latter rather than the former. In this thesis, the focus will be shifted more toward the consequences of the former, since the dissertation is primarily concerned with the BSA as a metaphysical, rather than epistemological, theory.

Intuitively, the problem of undermining can be grasped simply: if the best system law regarding coin flips states that $P(\text{Heads}) = 0.5$, then this doesn't actually imply that all nomologically possible worlds would have equal numbers of heads landing upwards as tails – it is the nature of chance events that they can land either heads or tails at every single flip, whether that results in equal numbers of heads and tails or not. Ultimately, no relative frequencies of heads and tails could actually contradict the $P(\text{Heads}) = 0.5$ law. But, consider one of these 'deviant' possible worlds where, for example, the ratio of heads to tails is 10:1 rather than 1:1. While such a world wouldn't contradict the $P(\text{Heads}) = 0.5$ law, and so should be possible according to this law, the deviant world would likely not have a best systematization of its laws that would have resulted in the $P(\text{Heads}) = 0.5$

law. So, in that case, the world should not have been possible in the first place, since it would not share the same laws.

In the literature (for example, Hall, 1994; Lewis, 1994), such deviant worlds are often referred to as F-type worlds or undermining futures since they involve future sequences of chance events that would undermine the very laws of nature that were presupposed to describe them.

It is only 'likely' that such worlds would undermine the laws of nature due to the tolerance in the best-systematization between simplicity, strength and fit. If the value of the chance function turned out to be, based on the observed relative frequencies of the Humean mosaic, 0.50001, for example, it could be that the trade-off with simplicity resulted in this being rounded down to 0.5. But this tolerance is likely small, and so for any given chance function and F-type world it is probably the case that the Humean chances of F-type world would be different, but this cannot be determined *a priori*.

The problem with F-type worlds is that they contradict the chance-possibility connection: that to every nonzero chance value, there exists at least one (nomologically) possible world, and vice-versa. This is not fatal when the implication fails from possibility to chance because infinite probability spaces already ensure this. But F-type worlds block the implication in the other direction: from non-zero chance values to possible worlds. This latter implication is otherwise known as the Basic Chance Principle (Bigelow et al., 1993).

In terms of being able to consistently define the set of worlds with laws L, W_L, and thereby allow a reductive, extensional, account of nomological modality, F-type worlds are problematic: their different chance functions imply different best-systematizations, and therefore different BSA laws, and therefore gives different answers to the membership of

the set W_L. This contradiction is clearly serious as it lies at the heart of Lewis's reductive account of modality.

Lewis may have a way out of the metaphysical problem described thus far by denying the fundamentality of the Basic Chance Principle, which he has claimed to be something not fundamental to chance in any case (Lewis, 1994 p.483). But, there is a further, equally serious epistemological problem with F-type worlds. The epistemological problem is how F-type worlds relates to Lewis's Principal Principle (PP; Lewis, 1980) that establishes a link between chance and initial credence, or subjective degree of belief, denoted below as the '$\mathrm{Cr}_0(-)$' function:

1) PP: $\mathrm{Cr}_0\big(A|T, H_{tw}\big) = \mathrm{Ch}_{tw}(A) = a$

That is, the initial credence an epistemic agent should assign to the proposition A, given the total amount of elapsed history, H_{tw}, and the theory of chance, T, is equal to a (that is, that chance of A at t in w).

The term 'initial credence' here denotes credence in the absence of any other information regarding the chances of A; the epistemic agent is assumed to be in a *tabula rasa* state, given only the information admissible on the right side of the conditional: T, H_{tw}.

The PP has great intuitive appeal, it is assumed. But substitution of A with F, an undermining future that is inconsistent with the theory of chance for the reason given above, and the right hand side of the equation must be 0. This is because the credence of an inconsistency must be 0 and the assumption that credence obeys the laws of Kolmogorovian probability (i.e. $\mathrm{Cr}(X|Y) = 0$ when $\mathrm{Cr}(Y) = 0$). Thus, the contradiction that $\mathrm{Ch}_{tw}(F) = 0$ by the PP, and that $\mathrm{Ch}_{tw}(F) > 0$ according to the best-system theory of chance.

Thau (1994) has extended this problem of undermining to all frequency-based accounts of chance, since by definition under any frequentist account the non-zero chance of any future cannot imply deviation from the finite or limiting relative frequency observed in the Humean mosaic. Further, Briggs (2009a) has also shown that attempts to revert to an anti-realist or relativistic account of chance, such as that proposed by Ward (2005) or Halpin (1994) also suffers from the undermining problem.

Lewis (1994) and Hall (1994) proposed an amendment to the PP, effectively making all initial credence valuations conditional on the theory of chance itself. But this solution fails to solve the metaphysical problem about chance since the chance-possibility connection remains broken. It also fails to give a reductive epistemological account of chance, since – as will be shown in Chapter 4 – it invites an infinite regress as to where some idealized agent's (after Hall, 1994) credence in the theory of chance itself derives.

In Chapter 4, a solution will also be provided that solves the metaphysical problem of undermining, and thereby also solve its epistemological manifestation. This will allow Lewis's aspiration of the BSA to be a package deal for both fundamental laws and fundamental chances to be realized, albeit with some modifications to the underlying doctrine of Humean supervenient chances (and acceptance of a very strong form of modal realism).

In brief, the solution is to show that the problem of undermining can be solved if instead of the best-system chance functions being supervenient on the Humean mosaic of a single world $w_@$, they are instead supervenient upon the 'composite' Humean mosaic of all worlds with a common initial segment in the set W_L. Thus, instead of the best systematization being of the Humean mosaic of a single world, it is instead of a multiverse of possible worlds, all of which are isonomic, and all of which have a common initial segment. This solution works because, within the set W_L of isonomic worlds, are the F-type

worlds. Thus, any best systematization that is read off the Humean mosaic comprising both F-type worlds and non F-type worlds, must result in a system of laws that are consistent for all worlds. To avoid circularity, W_L is then defined as the set of all worlds that are isonomic with $w_@$ with respect to the deterministic laws of $w_@$.

Since 'isonomic' means sharing the same laws, there is a circularity worry in this solution. It will be shown however that this circularity can be overcome if the best systematization is performed in multiple steps. The first step identifies the deterministic laws of $w_@$ and whether the Humean mosaic of $w_@$ is able to be made without appeal to indeterministic laws. The second step then determines the chance functions appearing in the indeterministic laws. The reasoning here is that the first step would identify whether the world was indeterministic or not, and where any chance functions would appear in the fundamental laws. But this first step would not have sufficient information to parameterize any chance functions. It does however provide the information necessary to define the set W_L, and thus allows these functions to be parameterized in the second step.

This full solution, along with an argument that it solves the undermining problem, will be detailed in Chapter 4, along with a more formal presentation of the problem of undermining.

2.4 The problem of non-locality

Non-locality arises where a quantum state is not factorizable into two or more localized states. That is, it is possible via a quantum phenomenon called entanglement to create a quantum system of two particles located at some non-zero distance from one another where the properties of each particle depends on the properties of the other.

Non-locality is often interpreted as inconsistent with Humean supervenience as a whole, and that insofar as Humean supervenience contradicts results of empirical science,

Humean supervenience should be rejected (Loewer, 1997; Karakostas, 2009; Butterfield, 2011).

Strictly speaking however, this is not the case. Recall from Chapter 1, the definition of strong supervenience being assumed throughout this thesis:

> **Strong Humean supervenience**: In all possible worlds, laws of nature, dispositions, and chances supervene on the distribution of categorical properties in the world (past, present, and future).

This strong supervenience claim says nothing about the ability to localize fundamental properties to contiguous or point-sized regions of spacetime. This form of supervenience is an instance of global supervenience that does not require locality. That is, it is a supervenience relation between the entire Humean mosaic, not some local part or parts of it, and the laws of nature and other nomological characteristics of a world.

The problem is rather that it appears to contradict Hume's original prohibition on necessary connections in nature, expressed previously in Chapter 1 as follows:

> **No necessary connections (NNC)**: There are no necessary connections between fundamental particulars. That is, all particulars are fully specified by their intrinsic fundamental properties – properties that cannot differ between duplicates – and all worlds are built up from such particulars according to some principle of recombination consistent with this requirement.

The apparent contradiction arises not only from the Humean requirement that all particulars be fully specified by their intrinsic fundamental properties, but also from the principle of recombination cited in the No necessary connections principle. To see this, consider Lewis's (1986) principle of recombination which states, roughly, that every little bit of any possible world can be glued to any other little bit or bits from other worlds to

form a new world. In Lewis's words: "[A]nything can coexist with anything else, at least provided they occupy distinct spatiotemporal positions. Likewise anything can fail to coexist with anything else" (1986, p.88).

Non-locality appears to contradict this principle because there is some necessary connection between spatially separated points or point-sized objects: the non-local property applies to multiple points, spatially separated, meaning that the property of one point determines the property of another point.

In Chapter 5 it will be shown that this there is no inconsistency between quantum entanglement and NNC. This is quite simply because the notion of 'necessity' being appealed to is nomological necessity, not metaphysical necessity. And, it is the job of the BSA to define what is nomologically necessary and what is not. In particular, if the laws of quantum mechanics are scientifically fundamental (or even if they're not), they are supervenient, according to the BSA. However, Lewis's (1986) principle of recombination quoted above applies not to bits of stuff grouped together by such supervenient properties, but by the points or point-sized objects comprising the Humean mosaic. That is, according to Lewis, the principle of recombination applies only to those parts of the Humean mosaic that are ontologically fundamental, not supervenient. These objects in turn can instantiate only perfectly natural properties – those privileged properties that are stipulated (according to Lewis) to be the primitive vocabulary of the BSA laws.

This argument seems straightforward. However, it requires acceptance not only of Lewis's privileged set of perfectly natural properties, but also of the claim that such properties could not be non-local. Both of these issues will be addressed in Chapter 6.

At this point, it is worth clarifying the relationship between the BSA and NNC: some kind of best-system account can be developed irrespective of NNC, as has been shown by – among

others – Cohen and Callender (2009). This is because the ability to describe the ontologically fundamental bits of stuff (or structures, for the structuralist) by a system of laws is independent of whatever those bits of stuff (structures) turn out to be. That is, the defining principle behind the BSA is that the laws of nature is the best system of laws that describe whatever is ontologically primitive. It takes no position on what is ontologically primitive, and thus whether features such as non-locality should be permitted or not. However, if NNC is the defining principle of Humeanism, which is considered desirable in terms of ontological parsimony, then by implication a best system account of lawhood consistent with NNC would likewise be considered more desirable than one which is not.

Chapter 5 starts with an introduction of quantum mechanics and its main competing interpretations. The extent to which each of the various interpretations of quantum mechanics are compatible with the other principles of Lewis's version of Humean supervenience and the BSA is discussed. In particular, the apparent point of contact between the Everettian or 'many worlds' interpretation of quantum mechanics and Lewis's modal realism will be noted as both invoking multiple possible worlds, all of which are as real as the actual world. Each interpretation will also be considered in light of the concept of lawhood adopted by the BSA as being the best trade-off between strength and simplicity. From this discussion it will be concluded that the Everettian interpretation fairs best in invoking the fewest types of entity and simple laws, while remaining empirically equivalent to the other interpretations. Finally, the phenomenon of entanglement and the problems it suggests for Humean supervenience and the BSA are discussed.

2.5 Armstrong's problem

Armstrong's problem is so-called after recognition of the imagined need for primitive universals by Armstrong (1978). It has also been referred to as the immanent comparisons problem (Cohen and Callender, 2009). It states that the metric of simplicity or strength in

one possible systematization may not be the same as that in another systematization (Cohen and Callender, 2009). The problem implies that there may be no objective comparison between the two systematizations since there may be no common shared metric that can be guaranteed to obtain in both. Because the BSA theory of lawhood presupposes that such inter-systematization comparisons can always be made, this failure means the formulation of the BSA theory of lawhood – considered as a realist theory – is not possible.

There are two ways that the metric of simplicity may fail to be objective: one is a trivial way that can easily be resolved, and the second a more fundamental way for which no solution proposed in the literature is acceptable to a Humean realist.

The first way in which the metric of simplicity can fail to be objective is for it simply to be language-dependent. Recall that simplicity can be measured by parameters such as the length of the proposition. But this would mean that simplicity is language dependent. Clearly such dependence on language would be problematic.

The second more serious problem is that the predicates used to formulate the candidate systematizations of a world are language relative, irrespective of whether the language contains redundant information. Therefore, it is possible to have predicates such as F in language L, and predicates F’ and F’’ in language L’, where F is the conjunction of F’ and F’’. Predicates such as F’ and F’’ may be the more fundamental, but the language containing F may win in any simplicity-strength trade-off since propositions of the form Fx are syntactically simpler than those of the form $F'x \wedge F''x$.

The extreme case of this problem is the trivialization problem which is implied by Armstrong’s problem and which Lewis (1983, p.367) was aware. In the words of Cohen and Callender: “Consider the predicate F that holds of all and only things in the world where an

arbitrary system S holds. If allowed to compete, it's hard to see why the trivial generalization (x)Fx wouldn't be the Best System for all such worlds: it is very simple, and also very strong (it strictly implies all truths).” (2009, p.5).

Unlike the undermining problem, Armstrong’s problem has not been considered fatal to the BSA. Indeed, Lewis (1983) recognized the problem and in his formulation of the BSA proposed a solution to it. But, his solution comes close to contradicting at least the spirit of Humeanism and the principle of parsimony embodied by Occam’s Razor. This is because the solution appears to imply further ontological primitives beyond the points and point-sized objects comprising the Humean mosaics of all the worlds in modal space (Lewis, 1986).

Of those responses to Armstrong’s problem and the trivialization problem that do not consider it as being fatal to the BSA theory of lawhood, there are two overarching strategies: i) introduce natural kinds into the Humean ontology, and further stipulate that laws may only be framed using predicates that designate them, or ii) accept the subjectivity or language relativity of the strength, simplicity and balance metrics.

Lewis himself appears to have adopted the first strategy by stipulating the existence of some properties that are objectively special. These properties, referred to as perfectly natural properties, he describes thus (Lewis, 2004):

> They are not at all disjunctive, or determinable, or negative. They render their instances perfectly similar in some respect. They are intrinsic; and all other intrinsic properties supervene on them. They are not conjunctive or structural. (p.204).

This description makes perfectly natural properties (PNPs) sound very much like natural kinds (although Lewis nowhere introduces them as such). Nonetheless, they seem to be objectively special in just the same way natural kinds are often considered to be objectively

special: they 'cut nature at its joints' by objectively identifying scientifically meaningul or useful (when framing laws) types of thing out in the world. This is not a criticism of PNPs per se, for it seems to avoid Armstrong's problem introducing something like natural kinds would seem to be a productive way to proceed. Moreover, PNPs are consistent with Humean metaphysics insofar that they are required to be intrinsic, so their content is exhausted by the space-time points or point-sized objects at which they are instantiated and there are no necessary connections between those points of instantiation.

However, it is difficult to see how the introduction of PNPs as being objectively special cannot imply an addition to the Humean ontology of something that seems to be very non-Humean: some kind of second-order property that may not supervene on the Humean mosaic. To be clear, Lewis does not refer to the property PNPs all share as a second-order property, although it is unclear what else it could be, even if first-order properties are not primitive universals. So, the stipulated existence of PNPs and the implication of some second-order property that they instantiate sits uncomfortably with Lewis's overarching project of Humean supervenience as a parsimonious physicalist metaphysical world view. The points and point-sized objects that comprise the Humean mosaic must have fundamental properties of some description, for else every possible world would be composed of nothing by grey goo. But to stipulate some privileged subset of these fundamental properties is problematic for any parsimonious Humean ontology.

Further, as noted by Earman (2005a), Lewis's perfectly natural properties are also inconsistent with principles of empiricist epistemology. It is fundamental to empiricism that entities should not be postulated that are necessarily unknowable. But, perfectly naturalness appears to be introduced by Lewis not as some supervenient second-order property that could in principle be identified through empirical investigation of the properties instantiated by the Humean mosaic, but as an act of faith in their existence and

the belief that "nature has been kind to us" (1994, p477). But this stance cannot be consistent in the setting of strong Humean supervenience in which this thesis is set. It cannot be an act of faith that nature is "kind" in every possible world in the revelation of perfectly natural properties, let alone any scientist's ability to classify them as perfectly nature. Consider qualitatively indistinguishable but distinct possible worlds that have different types of instantiating properties of PNPs swapped around (for example, mass and charge, in the example of Hall (2015) following an argument originated by van Fraassen (1989)), then it follows that it is impossible to have any epistemic access to such second-order properties as PNPs. This is the problem of Earman's empiricist loyalty test.

Of note, it can be countered that the empiricist loyalty test assumes that all empirical knowledge requires distinguishing between possible worlds, and that this should not be a criterion to apply to second order properties such as PNPS. A Putnam-style argument from reference may be deployed, concluding that if PNPs are necessary for the theory to work, there must be something 'out there' as referent. Perhaps, but the notion that two physically identical worlds could be different contradicts a fundamental empiricist epistemological principle in any case: that everything there is to know about a world is empirical. Moreover, the argument from reference does not imply that such PNPs are primitive or supervenient, only that they exist.

Lewis's own response to the epistemic inaccessibility of ontologically primitive PNPs is, in effect, acceptance of the implied epistemic – 'Ramseyan' – humility (Lewis, 2009). Lewis's adoption of a humility thesis seems as unattractive as the alternative of adopting a more relativistic or subjective account that defines the best system in terms of something epistemologically more accessible than PNPs. This is essentially the approach adopted by Loewer (2007) and Cohen and Callender (2009). Both argue for ditching PNPs as a solution to Armstrong's problem.

Clearly the ideal solution will take the best from both of these extreme positions: accepting that there is order in the actual world that cannot be explained by some features of Human cognition, but that the order can be explained without resorting to the introduction of empirically inaccessible second-order properties such as PNPs.

In Chapter 6 a solution that lies between these two extreme positions will be proposed. This solution will be inspired by Shannon Information Theory (Shannon, 1948) and the notion that the nomological character of a world may be contained in the predicates used in the best systematization of it. It will be shown how this can be achieved using a concept based on Shannon's Information Theory that is defined over modal space, rather than some physical or communication system. This concept can be thought of as the degree to which a predicate is shared among all possible worlds in modal space – those predicates designating properties in fewer possible worlds contain more empirical information than those appearing in more possible worlds.

Intuitively, the idea is that those predicates with the richest empirical content (most notably those trivializing predicates denoting the best-systematization of a world w or the laws of nature of w themselves) are instantiated more frequently in w, and worlds with the same nomological character as w, than in less similar worlds. That is, such predicates would necessarily not be instantiated in worlds outside the sphere of nomologically possible worlds centered on w, but would necessarily – according to the BSA – be instantiated in all worlds within the sphere. In contrast, predicates with no empirical content would be instantiated in all possible worlds, irrespective of their nomological character, since they could only be used to construct logical or mathematical truths, all of which are defined as true in all possible worlds.

The proposal in Chapter 6 is that the predicates that should be used to express the laws of w should be at some optimal place in between these two extremes. More specifically, they

should conjointly contain the least amount of empirical content, while still containing enough empirical content to specify the laws. What is needed therefore is some way to measure this quantity of "empirical content".

It is the measurement of this "empirical content" where Shannon's Information Theory and a concept of entropy is introduced in Chapter 6. Briefly, according to Shannon's Information Theory, entropy is maximized for a message when it contains the least amount of information in the everyday sense of that word (or the maximum amount in the formal definition of the term in Shannon's Information Theory). That is, when it is the most difficult to determine the meaning of a message. In a similar vein, empirical content is maximized for a predicate when it is most difficult to identify the nomological character of the world given its pattern of instantiation across all possible worlds: those predicates instantiated in all possible worlds contain no empirical content; those instantiated in only those worlds that are within the sphere of nomological possibility have the most empirical content.

3. BSA chances and the zero-fit problem

3.1 Introduction

In section 3.2, the precise notion of chance as it appears in Lewis's so-called 'package deal' (1987a, postrscript C) formulation of the BSA will be introduced and defined as something metaphysical rather than epistemological. This metaphysical casting of the notion however does not affect its formal role, being still modelled using the probability calculus. Second, the notion of fit and the role it plays alongside strength and simplicity in the selection of the best system will be examined in a little more detail. Specifically, the BSA notion of fit as the probabilistic concept of maximum likelihood as it has typically been conceived in the literature will be examined in closer detail. This naturally leads on to the zero-fit problem introduced in section 2.2 which occurs in those cases with infinitely many possible worlds to choose from and an exploration of the implicit assumptions being made in the formulation of the problem.

After accepting the susceptibility of the maximum likelihood conception of fit to the zero-fit problem, the alternative notion of typicality will be considered before rejecting it on the grounds that it would not be generalizable to all nomologically possible worlds. Instead, a statistical measure, X^2, will be proposed in section 3.6 as a better solution to the zero-fit problem. X^2 may be used to measure the degree of difference between different sequences of chance outcomes sharing the same probability space (that is, possible outcomes). With alternative sequences of chance events denoting the alternative possible worlds in which those events occur, X^2 will be shown to be a better measure of fit than some probability measure. Moreover, this solution works in both the finite and the infinite case, and comes from re-thinking what is important about the measure of fit as it is used in the BSA selection process to identify the best fitting world.

3.2 BSA laws and chances

The best system account of lawhood can be used as a theory of both laws and – in the case of indeterministic possible worlds – chances too. This so-called 'package deal' (1987a, postrscript C) is attractive in that it provides the metaphysical underpinning for both a theory of lawhood, as well as a theory of chance. The purpose of this section is to explain what is meant by the term 'chance' in the context of the BSA as formulated by Lewis, and how the best systematization of chancy worlds is identified through the trade-off between strength, simplicity, and the notion of fit.

The term 'chance' as it occurs in Lewis's formulation of the BSA has a very restrictive and precise meaning (Lewis, 1980). First, assume that the probability calculus is the mathematical theory that models indeterminate truth values using continuous-valued functions (i.e. probability functions). Whether or not a truth value is indeterminate or not is here assumed to be made relative to some epistemic agent, A, that inhabits some world w, and A's knowledge of facts about w. A is omniscient if A knows everything about w's history and laws of nature and non-omniscient otherwise. A may be omniscient or not, w may be determinate (here assumed to be canonically defined as Laplacian determinism) or not, and these two possibilities are independent of each other. Thus, multiple combinations of A's omniscience and w's determinacy now arise. Two of these combinations in particular can be used to distinguish two philosophically distinct interpretations of the probability calculus:

1. If A is non-omniscient and w is deterministic, there are some propositions whose truth value A cannot assent to with complete confidence. A's degree of belief in the truth value of such propositions must therefore always be <1 given A's knowledge of w.

2. If A is omniscient and w is indeterministic, there are some propositions whose truth value A cannot assent to with complete confidence. A's degree of belief in the truth value of such propositions is thus, as in the case above, always <1 given A's knowledge of w. But, because A is omniscient, this is also the case with respect to any agent's level of knowledge. Such indeterminacy is therefore a feature of w, rather than of A or any other epistemic agent.

The two remaining possible permutations of A's omniscience and w's determinacy are uninteresting since the case where A being non-omniscient and w indeterministic cannot discriminate between the two possible types of uncertainty listed above, and the case where A is omniscient and w deterministic generates only trivial truth values ('trivial' here intending to mean taking only values 0 or 1).

The important point to note here is that the indeterminacy of both scenarios above may be modelled using the probability calculus. That is, the probability calculus itself does not imply a particular interpretation of its mathematical formalism. Metaphysically however, the two scenarios above are quite distinct. The first is metaphysically uninteresting (to a realist, at least), and is a consequence only of an epistemic agent's degree of certainty or belief. Applications of the probability calculus to these degrees of certainty or belief therefore have only an epistemic character, and the resulting probability functions that quantify such uncertainty shall be referred to as epistemic chances.

The second is metaphysically interesting, as it says something about w itself. Probability functions that quantify this type of indeterminism shall be referred to as metaphysical chances. The BSA theory of chances concerns itself exclusively with metaphysical chance, and has nothing to say regarding epistemic chances, beyond the special case of those modelling the epistemic state of an omniscient agent. Henceforth, the term 'chance' will be assumed to refer to this metaphysical concept of chance, rather than of the epistemic kind.

Notice that, in the way that metaphysical chance has been defined thus far, the BSA is committed to a realist theory of chance, as well as a realist theory of lawhood, since chanciness is defined to be a feature of the world rather than of a particular epistemic agent. Theories of chance that are conceived of as models of reasoning under uncertain epistemic states of non-omniscient agents – such as the theory of Ramsey (1990), for example, who is also one of the originators of the BSA theory of lawhood – are therefore at best orthogonal to that of Lewis's theory of chance, if not inconsistent with it.

These two forms of chance can be made distinct within the probability calculus. Using the formalism of Lewis (1980, 1994), let H_i denote a proposition describing all of the laws of nature and all of history up to some time i for some world w. For metaphysical chance, denoted by $\mathrm{Ch}_m(-)$, to exist in world w, the following condition holds for some future chance outcome p in w and for some prior times t and v:

$$\mathrm{Ch}_m(p|\mathrm{H}_t) = \mathrm{Ch}_m(p|\mathrm{H}_v)$$

That is to say, $\mathrm{Ch}_m(p|\mathrm{H}_t)$ remains the same, irrespective of how much historical data, H, is known. This condition for metaphysical chance fails to hold for epistemological chance in at least some – if not all – non-trivial instances. Metaphysical chance holds non-trivially only for future chance events. For past events, it is assumed that the chance of the event will be either 1 or 0.

It is possible to see these two types of chance in the actual world if some reasonable assumptions are made. Consider quantum mechanics first. Assume that the probability functions that model the chances (in the Copenhagen interpretation) of different measurement outcomes following collapse of the wave function are not reducible to some more fundamental process. Then, in the actual world those chances in quantum mechanics are metaphysical chances – it would not matter how much an epistemic agent knew about

the history of the universe up to the point of measurement, the chances of each measurement outcome would remain unchanged.

Thus, whether an omniscient agent at time t or time v ($t \neq v$) prior to a quantum measurement p – for example, a photon passing through one slit or another of a double-split experiment – the equality that $\mathrm{Ch}_m(p|\mathrm{H}_t) = \mathrm{Ch}_m(p|\mathrm{H}_v)$ would still hold.

It may be objected that this condition surely does not hold for all metaphysical chances, for they change over time as we come to know more about what has happened – or not happened – in some elapsed interval over which the chance measurement is made. For example, if p = "atomic nucleus X will decay before noon on 31st December, 2020", then if t is July 1st, 2020, at which point X has not decayed, and v is 30th December, 2020, then obviously the chance of X decaying within the one remaining day of 2020 is less than during the 180 days or more preceding it. True enough. However, the argument relies on p being poorly specified for it provides no information on the period over which the trial is to be measured. Revise p to "atomic nucleus X will decay between noon on July 1st, 2020 and noon on 31st December, 2020" and for any prior times (that is, before noon on July 1st, 2020) the chance remains constant.

More detail still may need to be added to p to make it fully watertight, but this risks losing the essential point: that some chance evaluations may change given more empirical knowledge, and others may not change. More will be said regarding the distinction between different notions of chance in quantum mechanics and how they relate to BSA chances in Chapter 5.

In contrast to the quantum mechanical example, consider one from epidemiology. Assume that smokers have a 10-fold increase in lifetime risk of developing lung cancer at 5% compared to non-smokers, and that those smokers with a certain genetic mutation have a

further 2-fold increase at 10%. Faced with a smoker whose genetic status is unknown, it would be rational for an epistemic agent to assume that, without further knowledge, their risk would be 5%. But with the further knowledge that they do in fact have the risk-conferring genetic translocation, it would then be rational for the epistemic agent to say that their risk would be 10%. These epidemiological risks are clearly denoting epistemic chances because the value of the risk varies by the amount of knowledge. In the extreme case, when an omniscient epistemic agent knows everything about the history of the actual world up to some point t, they would be in a position – in a deterministic world at least – to say that an individual will definitely develop or definitely not develop lung cancer in their lifetime. That is, the lifetime risk of lung cancer would be 1 or 0 and the need for any concept of chance disappears.

The distinction between what has here been called epistemic chance and metaphysical chance is not new, with metaphysical chance being essentially the same as C S Pierce's notion of absolute chance (Reynolds, 2002 p.93), dating back over a century. There is however one crucial difference between Pierce's absolute chance and the metaphysical chance of Lewis's theory: Pierce's absolute chance was apart from the laws of nature, effectively being a description of anomalous (according to the science of the day) phenomena such as the mind with an otherwise mechanistic view of the universe. Lewis's metaphysical chance on the other hand is defined as being part of the laws of nature in indeterministic worlds – metaphysical chance is denoted by those probability functions appearing in (or being implied by) the fundamental laws of an indeterministic world.

The epistemological notion of chance corresponds to Hoefer's (2007) notion of deterministic chance. That is, chance that can occur in deterministic worlds, due to some degree of ignorance on behalf of some epistemic agent. For the purposes of the BSA and this thesis, the notion of epistemological or deterministic chance will be of little interest, as

it has no metaphysical significance. Given the varying interpretations of quantum mechanics however, it should not be assumed that the distinction between the two is always clear-cut in practice. The position of compatibilism, for example, would dispute the distinction made above between metaphysical and epistemological chance as being meaningful (Glynn, 2010; Emery 2015). Chancy events thus can be as equally valid descriptions of statistical mechanics (Frigg and Hoefer, 2015) as quantum mechanics. This debate will not be discussed further however as it lies beyond the scope of this thesis, which is restricted to the definition of metaphysical chance given above.

What makes a world indeterministic according to the BSA theory of lawhood is when the best system contains at least one non-trivial probability function. The best system is the one which is optimal in terms of strength, simplicity and fit. As in deterministic worlds, strength may be measured by the number of true statements that can be inferred, simplicity by the length of expression, and fit – conventionally, at least – by the likelihood of the actual world (that is, the laws should make the actual world the most likely from all other possible permutations of chance events).

Clearly, in deterministic worlds, the best system will have very good, if not perfect, fit: the resulting laws describe almost every, if not every, observation of the Humean mosaic with absolute certainty. That is, there is very little, or no, difference between the world at any point in time and the description of it that is generated by the fundamental laws, given some initial state.

What makes a system best in worlds with chancy outcomes includes a measure of fit. In the literature fit has assumed to be measured as the likelihood of the actual world. That is, a system of laws A has better fit than some other system, B, if A makes the actual world more likely (in terms of its occurrence of chance outcomes) than B. The best fitting system is therefore that one which makes the actual world the most likely (Lewis, 1994):

> … some will fit the actual course of history better than others. That is, the chance of that course of history will be higher according to some systems than according to others (p.480).

Elga (2004), after Bostrom (1999), has correctly noted that this measure of fit fails in the case of an infinite field of isonomic possible worlds. It therefore remains an open problem how to address this failure of defining a measure of fit that, along with simplicity and strength, can identify the best fitting system(s), and therefore the laws, of a possible world. Later in this chapter, a solution to this problem will be given by proposing an alternative measure of fit which works in both the finite and infinite case. First however, the other measures of simplicity and strength will be considered, and how they trade-off against one another.

Simplicity is relatively easy to define using a measure such as Kolmogorov complexity (Hankerson, 2003): the length taken, relative to a language L, to express the best system. 'Length' here could denote the number of characters in binary, or just the language L, just so long as propositions comprising all candidates for best system are expressed in the same language to make comparison meaningful. The problem of defining a language with a so-called trivialization predicate (e.g. "is the best system") that describes a best system is recognized by Lewis (1994), and prompts the introduction of a subset of properties that share the same high-order property of being 'perfectly natural' properties. These perfectly natural properties may roughly be considered as natural kinds and are to be used as the sole predicates in the language to avoid such trivialization. Whether this solution is warranted will not be discussed here as the problem of trivialization is not directly relevant to the notion of fit. It will however be discussed in Chapter 6.

Strength can be considered as a measure of the information content of the system, which could be quantified in terms of bits using Shannon's concept of negentropy (Shannon, 1948). In this case the Humean mosaic of a possible world w will have associated with it a

maximal quantity of information, corresponding to a complete description of every point in some co-ordinate system encoding the Humean mosaic of w. Call this quantity of bits H_{MAX}. Strength can now be defined as that proportion of H_{MAX} that matches the predictions of a candidate best system. Clearly the strongest candidate would have strength = 1. But this system would require a complete description of every point, and would thus fail in terms of simplicity. A simpler system would likely have strength < 1, but be so much more simple that the loss of strength is trivial in comparison.

That there is a trade-off between simplicity and strength of the kind discussed is not disputed. It is not obvious however that either i) an optimal value of both exists, or ii) if such a value does exist, that it is either objective or unique. Ideally, the optimal value of strength and simplicity would be clearly revealed as some apparent join point or global minimum or maximum, depending on how the relationship is represented. Figure 3.1 below depicts a situation where the optimal value appears to be both objective and unique. Figure 3.2 depicts a situation where no obvious point of optimal trade-off is apparent.

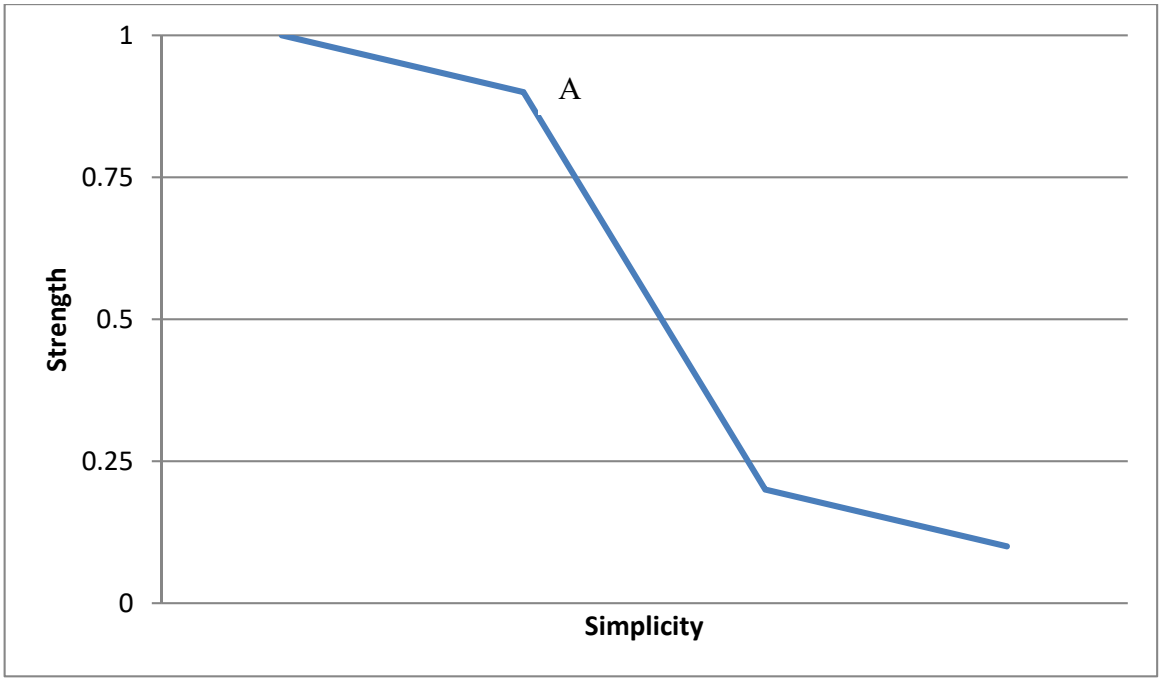


Figure 3.1. Chart depicting relationship between increasing simplicity and decreasing strength with an obvious and unique optimal value at A.

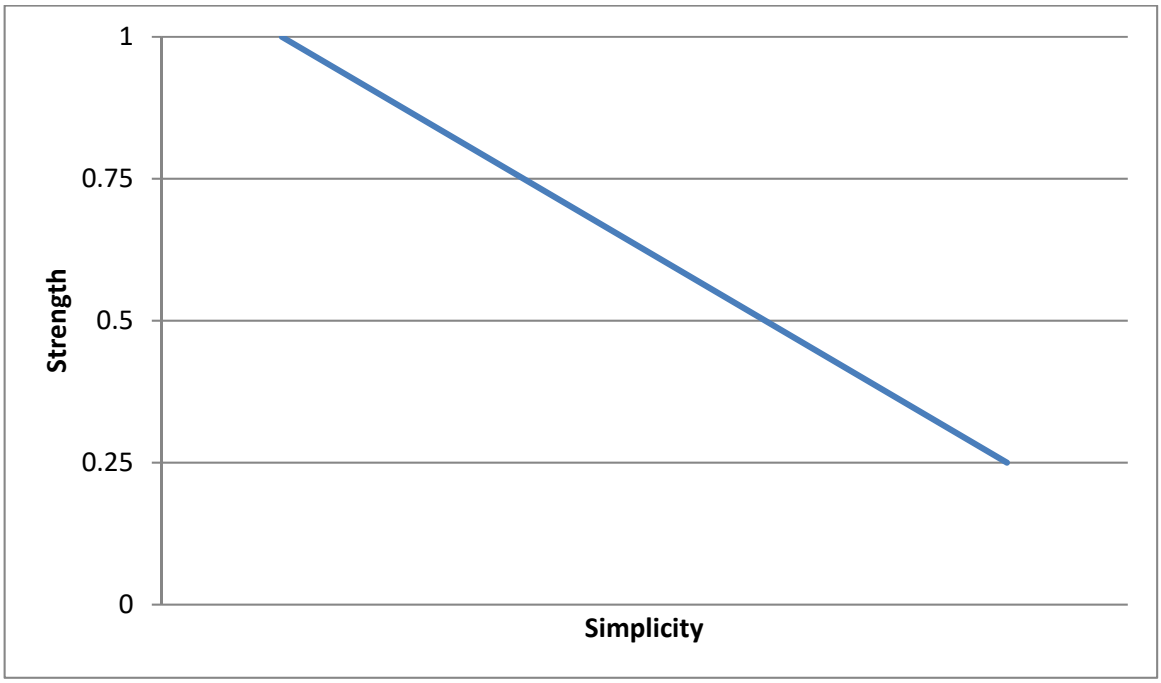


Figure 3.2. Chart depicting relationship between increasing simplicity and decreasing strength with no obvious or unique optimal value.

It is an open question whether, across all of $\mathcal{M}$, the ratio of worlds with a single trade-off point (and therefore a single set of fundamental laws) to those with more than one would

limit to some value greater than 1. Clearly the two scenarios depicted in Figure 3.1 and Figure 3.2 are huge simplifications. Most significantly, the issue of weighting is ignored. It may be, for example, that strength is more important than simplicity, or vice-versa. Ultimately, that it is meaningful to trade-off strength against simplicity and that any optimal trade-off would be the best systematization must be taken as an article of faith. This faith can be bolstered by the observation that scientific laws in the actual world do very effectively summarize numerous observations whose description would otherwise be hugely inefficient to describe individually. But the actual world is one of many possible worlds, so the most we could say is that the actual world seems to be one where the BSA trade-off works well.

In fact, to be a truly Humean theory, the BSA must assume a trade-off not only between strength and simplicity, but also fit. The argument for this conclusion is as follows:

> P1: Deterministic worlds are distinguished from indeterministic worlds by the appearance (or implied existence) of metaphysical chance functions in the fundamental laws.
>
> P2: Such chances are determined by the laws of nature (according to the BSA theory of chance).
>
> Therefore: Worlds cannot be partitioned into deterministic and indeterministic sets without first establishing their laws of nature.

It will be assumed that P2 above is uncontentious since it is nothing more than a statement of one of the assumptions of the BSA theory of chance. P1 however is contentious since it contradicts the prevailing Laplacian definition of indeterministic worlds as those instances where duplicate worlds up to a point in time diverge thereafter (van Strien, 2014 p.26).

That is, any two worlds $w_1$and w_2 are indeterministic according to this alternative definition when the following two conditions hold:

1. From time t_0 until some (future) time t_1, w_1 and w_2 are perfect duplicates, and

2. From time t_1 onwards, w_1 and w_2 are not perfect duplicates,

where the notion of *perfect duplicate* is defined in purely Humean terms after Lewis (1986 p.80). However, this inference is valid only when it is assumed *a priori* that the worlds share the same laws. Were this assumption not made, the definition above would not distinguish between worlds where the principle of the uniformity of nature holds and those where it does not. This can be seen with two example binary outcome worlds, w_1 and w_2 that both have the same sequence of outcomes from t_0 to t_1 thus: 0000. But then at some time $t_2$the worlds diverge and the remaining sequences are:

1. w_1 = 11111

2. w_2 = 22222

Note that, according to the principle of strong Humean supervenience above, neither world can be established *a priori* as being best described by a non-trivial indeterministic law (for example, Ch(0)=0.5) or a deterministic law (for example, for w_1: type-0 outcomes become type-1 outcomes halfway through the Humean mosaic). The indeterminacy or otherwise of any world must be read from the Humean mosaic. To object that the above kinds of worlds are impossible is to assume the principle of the uniformity of nature as primitive. But such an assumption contradicts the principle of recombination of the Humean modal realist. Therefore, assuming both strong Humean supervenience and modal realism, the property of being indeterministic is not a primitive property of possible worlds, but a supervenient one determined by the (supervenient) laws of nature.

For the Humean BSA theorist therefore, it is the real existence of chance that determines whether a world is indeterministic or not, and the chance functions of the world are implied and parameterized by the laws of nature, which are themselves defined as the best-systematization of the world. Therefore, whether the world is indeterministic or not cannot be answered a priori. It follows that in all possible worlds, the best systematization must be that which trades off not only simplicity and strength, but also fit, since it is this three-way optimization process that determines whether fit would be meaningful for every law or not (i.e. whether non-trivial chance valuations would exist in the world or not, respectively), and therefore whether the world would be indeterministic or not. In short, the indeterministic nature of the world is itself supervenient.

Lewis himself was not clear on this point, as he introduced the BSA using the special case of deterministic worlds, which gave the impression of some further ingredient being needed to account for indeterministic worlds (namely, some notion of fit). Nonetheless, Lewis's theory is consistent with the determinacy or otherwise of a world being pinned down by the BSA itself, rather than being an *a priori* property. Given that in deterministic worlds the tradeoff is between strength and simplicity, fit would have already been filtered out of the tradeoff process.

3.3 BSA chances and the probability calculus

One further assumption will be made here regarding BSA chances: that the probability calculus can be used to formalize chance, both as a general concept, as well as the chances that are produced by the BSA. BSA probability functions therefore provide the calculus to make chance valuations and are implied by any indeterministic laws of the best-system.

In the preceding section the notion of BSA chances was introduced as a metaphysical concept, rather than an epistemological one, that applied to those worlds deemed indeterministic by the selection of the best fitting system. Yet, while BSA chances remain

consistent with the probability calculus, their interpretation within that calculus must be distinct from other applications, such as purely epistemological conceptions of chance. The purpose of this section is to provide a distinctly metaphysical interpretation of BSA chances within the probability calculus. This interpretation uses the notion of modal space introduced in Chapter 1 and the language of possible worlds, as laid out by Lewis (1986). Use of these concepts facilitates the further discussion of BSA chance and the problem of undermining in Chapter 4, as well as the notions of fit used in the selection of the best system introduced and discussed in the succeeding sections of this chapter.

A more formal conceptualization of BSA chances within the canonical definition of the probability calculus can be made over modal space, as follows: Assuming that a best systematization exists for the Humean mosaic of the actual world, then any indeterministic law l that is a member of the best systematization of laws of the Humean mosaic, L, is therefore associated with a BSA probability space S_L. Such probability spaces will be relative to a set of isonomic worlds W_L of which the actual world is a member, where 'isonomic' refers to all worlds that have the same set of laws L.

S_L itself can be defined in terms of a space of possible outcome values Ω_L that are defined by L (for example, 'heads', tails' in the case of coin tossing), the set of isonomic possible worlds defined by S_L, W_L, and for each indeterministic law $l \in L$ a probability function $\mathrm{Ch}_l(-)$. Thus for a best systematization of laws, L, that defines the set of isonomic worlds W_L picked out by S_L, each indeterministic law l defines a probability space $S_l = < \Omega_l, W_L, \mathrm{Ch}_l(-) >$.

According to this interpretation of probability in terms of BSA chances, the set of isonomic worlds W_L constitutes the event space of all possible sequences of events for any outcome value within the outcome space Ω_l. To flesh this formalization out with an example,

consider the toy coin-tossing world $w \in W_L$ where the Humean mosaic of w comprises (only) the following sequence of coin-toss outcomes:

w = HTHHHTT

That is: 4 heads, 3 tails. Thus, for the BSA law of this world l that best fits this sequence, the corresponding probability space $S_l =< \{\mathrm{H}, \mathrm{T}\}, W_L, \mathrm{Ch}_l(-) >$, where the set of possible worlds W_L contains all worlds with some number of binary-valued chance events described by the chance function $\mathrm{Pr}_l(-)$. Thus to each such possible sequence we can associate a possible world $w \in W_L$. Each $\mathrm{Pr}_l(-)$ is then a function mapping each member of W_L onto the real-valued unit interval.

There are many things to note from this example. First, BSA chance valuations map to the real-valued interval, not merely the set of rationals within the unit interval, as it would be for naïve frequentist theories of chance. The BSA chance function may also deviate from the naïve frequentist chance function by not being an automatic calculation of relative frequencies. In the example above, the relative frequencies are multiples of $\frac{1}{7}$ and thus an infinite decimal value. But for the BSA, saying the chance of heads $\mathrm{Ch}_l(H) = 0.1428\ldots$ is not very simple, at least if simplicity is measured in terms of length. Far better to just say that $\mathrm{Ch}_l(H) = 0.14$ for example, or even $\mathrm{Ch}_l(H) = 0.5$. In this way the best-systematization trades off simplicity against fit and the chance functions describe distributions whose maximum likelihoods are not the observed relative frequencies.

Second, the term 'some number of binary-valued chance events' would almost certainly exclude those instances with a single coin toss, since in such worlds the best systematization should be a deterministic world where all tosses end heads up, for example. Similarly, for all worlds where only one outcome value ever occurs in the Humean mosaic of that world. This latter case is problematic and will be discussed in more detail in

Chapter 4. The problem is as follows: consider those worlds with multiple 'chance' events where only one outcome occurs. By construction, each one of these worlds is within W_L, a set of isonomic worlds that are indeterministic. But, if only one outcome occurs in some of those worlds, the best systematization of those worlds may not be one which contains non-trivial chance functions. That is, the BSA process would identify them as deterministic, not indeterministic, worlds, thus contradicting the membership criterion of W_L. This is an extreme example of the undermining problem, and at least a partial solution will be proposed in Chapter 4.

Finally, because the notions of simplicity, strength and fit are purely formal notions, as are the notions of a probability space and chance function, the 'binary events' need not be coin tosses without some further clause requiring that coin-toss counterparts exist in each $w \in W_L$.

In addition to defining the probability space associated with each best system law l of the best systematization L it is also possible to consider the probability space of L itself, S_L:

1) $S_L =< \Omega_L, W_L, \mathrm{T}_L >$

Where $\mathrm{T}_L = \{\mathrm{Ch}_1(-), \mathrm{Ch}_2(-), \ldots\}$ and denotes the entire theory of chance shared by all worlds in W_L.

The purpose of this section has been to develop an interpretation of the probability calculus specific to the BSA theory of chance. In many respects the key term in this interpretation is the set of possible worlds W_L. This set represents the set of all nomologically possible worlds relative to the set of laws L. To avoid circularity therefore, the BSA must define this set of worlds independently of the condition that all members of W_L have chance greater than zero. Key to achieving this is developing a notion of fit between each nomologically possible world and the best systematization of the actual

world that is both mathematically consistent while also ensuring that the degree of fit is maximized for the actual world. This will be the purpose of the following sections.

3.4 Fit as maximum likelihood

As discussed above, fit is one of the three measures involved in the selection of the best system, and needs to be brought into the selection process right from the beginning since, at that point, it cannot be determined whether a world is deterministic or not. A world is only deterministic if the best systematization of it does not contain or imply any metaphysical chance functions of the form proposed at the beginning of this chapter. Otherwise, it is indeterministic. That is, it is indeterministic if systems of candidate laws that contain or imply at least one metaphysical chance function have a better fit (all other things – simplicity and strength – being equal) than systems that do not contain or imply any metaphysical chance functions.

However, the notion of BSA fit in the philosophical literature, including that of Lewis's development of BSA, is often left vague (for example, Lewis, 1994 p.480). As will be shown below, Lewis and others appear to cite the standard notion of fit used in statistics as that being applicable to fitting a systematization to an indeterministic world. This standard notion measures fit is the degree to which the systematization makes the actual world the most likely. That is, given some random variable that maps a candidate systematization onto a set of nomologically possible worlds, the resulting distribution for the best systematization would have a maximum likelihood at a point corresponding to the actual world.

The main purpose of this chapter is to explore this notion of fit further, and then to raise the issue of the zero-fit problem for the BSA. This problem is a direct consequence of this standard notion of fit as maximum likelihood, and results in any case where the best

systematization would map on to an infinite number of possible worlds, thus making the maximum likelihood of each possible world, including the actual world, zero.

First, it is worth noting that in fact it is quite possible that the actual world would not be among those worlds with the best fit to the best system, simply because there is a trade off with strength and simplicity to be taken into consideration. This trade-off was demonstrated above with the example of a world comprising seven coin tosses. For such a systematization, a nomologically possible world with a relative frequency of exactly 0.5 for each outcome would fit the law $\mathrm{Ch}_l(H) = 0.5$ better than the actual course of history because of the trade-off.

The more general non-uniqueness problem appears not to have been a consideration from those who have discussed the notion of fit in the literature. Lewis (1994), for example, states:

> …some [candidate indeterministic laws] will fit the actual course of history better than others. That is, the chance of that course of history will be higher according to some systems than others. (Lewis 1994, p.480).

Similarly, Frigg and Hoefer (2010) state:

> The fit of a theory is measured in terms of the probability it assigns to the actual course of history. (Frigg and Hoefer 2010, p.105).

Williams (2008) has the most succinct definition of fit as the "probability that the world arises, given the laws of nature" (2008, p.406)

Thus, fit is measured in terms of the probability that the actual world be singled out from the set of all possible worlds. Thus, for maximum fit, that the distribution of chance outcomes in the actual course of history should be at the point of maximum likelihood for the probability distribution described by the BSA theory of chance of the actual world, T_L.

It is worth nothing that the fit of the actual course of history may not be uniquely the best. In fact, it almost certainly would not be in those worlds comprising independent identically distributed (IID) events within the same probability space. IID chance events are those – such as the roll of a die, the flip of a coin or the photon in a double slit experiment – where the outcome of each event (or 'trial' in the statistical literature) is causally independent of every prior event. Thus, for example, the chance of getting a head on the second coin flip is not influenced by getting a head on the first flip.

Given IID chance events and their corresponding chance functions appearing in the best systematization of indeterministic laws, the distribution of those IID events in the actual course of history would be as equally likely in all of the other possible worlds with the same relative frequencies of the chance outcomes. Thus, for example, in a coin toss world with four coin tosses, the world with the sequence HTHT would have the same likelihood as the world with the sequence HHTT. In fact, all permutations of the sequence would have the same likelihood, just so long as the relative frequencies of heads and tails remained constant. To assume otherwise would imply the events were no longer independent of each other, and thus not IID events.

However, this non-uniqueness is not a fatal blow to the notion of fit implied by Lewis, Frigg and Hoefer, and is indeed consistent with it. It simply means that the actual world should be at the point of maximum likelihood, albeit along with all other worlds that share the same sequence-independent distribution of chance events.

However, the definition runs into problems when there are an infinite number of chance events in the world (and therefore an infinite set of nomologically possible worlds with identical chance distributions), in which case the chance of the actual world, along with every other world in W_L, is the same. Namely, zero. The problem is summarized by Elga (2004):

> … when continuously infinite state-spaces are involved, a great many candidate systems … will ascribe zero chance to any individual history (Elga, 2004, p.68)

An intuitive rendering of the problem is as follows, for the case of chance functions with binary outcomes such as tosses of a coin. Assume that the coin is fair and that $\mathrm{P}(H) = 0.5$. Clearly, the most likely outcome for any sequence of coin tosses should therefore be 50% heads. For finite numbers of coin tosses, this is indeed the case. For 2 coin tosses, the probability of 50% heads is 50% (i.e. the sequences HT and TH), the probability of 0% or 100% heads is 25%. For 4 coin tosses the probability is approximately 38%, and for 100 tosses the probability is approximately 8%. The probability of 50% falls as the number of tosses (i.e. 'trials') increases, but it still remains the most likely outcome.

That describes the finite case. The infinite case is more problematic. The probabilities described above for various numbers of coin tosses may be thought of as being calculated by a product of two values. Let these two values be X and r. The X component need not be considered further (it is a co-efficient, estimated by Stirling's formula, for example). Rather the problem is with the value r, since for n coin tosses $r = 2^{-n}$. Thus, when n becomes infinite the probability of 50% heads is zero, as is the probability of 40% heads, or any other percentage. Therefore, in the infinite case, there is no outcome that has the maximum likelihood.

Translating this into the case of the BSA and fit ranging over possible worlds, first let the set W_L be the set of possible worlds defined by some candidate system of laws L and let $|W_L|$ be the cardinality of that set. Then, the probability of the relative frequency of some chance outcome in the actual world, $r_@$, being the same as that specified by L, r_L, will be

2) $\mathrm{P}(r_@ = r_\mathrm{L}) = \mathrm{X} \,.\, |W_L|^{-1}$

Where X is a coefficient as described above. Thus, where the cardinality of the set W_L is infinite due to an infinite number of possible permutations of the Humean mosaic, (2) is problematic as $\mathrm{P}(r_{@} = r_{\mathrm{L}}) = 0$. This is the zero-fit problem as it applies to discrete chance outcomes.

In the context of finite permuted sets of chance events, the zero-fit problem cannot occur. But (2), as well as any notion of fit involving a measure over the set W_L, does result in the non-uniqueness of IID-equivalent worlds as shown above. That is, all permutations of chance events with the same relative frequency would have the same value for $\mathrm{P}(r_{@} = r_{\mathrm{L}})$, and thus the same fit. This lack of uniqueness is not fatal to the measure of fit appealed to in the literature, since the actual world would still be the most likely. In fact, given the relative frequencies in IID-equivalent worlds would all the same, it is desirable.

Before proposing a solution to the zero-fit problem, another solution proposed by Elga (2004) will be discussed. This solution relied on the concept of *typicality* and will be discussed in the next section.

3.5 Fit as maximum typicality

The notion of typicality was first mentioned by Elga (2004) as a possible solution, and subsequently developed as a preferred measure in place of maximum likelihood by Williams (2008), when speaking of a toy coin flip world:

> What you can legitimately expect is that the outcome of the series of coin flips will have certain global properties. In short, you expect the outcome to fall within a broad class of outcomes: the typical ones.
>
> I submit that requiring a world to be 'typical' by the lights of its laws of nature is a legitimate explication of the constraint that the world 'fits' those laws. (Williams 2008, p.407).

This clearly advocates typicality as an alternative measure of fit, but the notion of typicality itself remains elusive. Williams (2008) attempts a definition in terms of randomness. That is, a sequence of chance events is typical if it is generated by a random variable. This alone is inadequate however as any sequence of coin flips can be generated by any (non-trivial) random variable. Williams (2008) appears to be adopting the naïve lay understanding of the term 'random'. Randomness, as it exists in its technical sense in statistics and probability theory, means nothing more than being generated or describable by a random variable, and a random variable denotes little more than a mathematical function with some kind of stochastic or probabilistic element comprising it.

Despite this, there appears something salvageable from Williams' notion. It is likely that he is confusing randomness with measures such as disorder, compressibility or entropy. That is, sequences such as:

3) HTHHTTHTTHHTT

Are more disordered, less compressible, and have higher entropy, than sequences such as:

4) HHHHHHTTTTTTT.

Assuming this to be closer to the intended meaning of 'typical' that Elga and Williams have in mind, call this measure T for a given sequence, and the maximum value T_{MAX}, where $0 < T_{\mathrm{MAX}} < \infty$.

The question then is whether the value T_{MAX} reliably identifies the actual world as being maximally typical, and thus solves the zero-fit problem. Clearly it cannot identify the actual world uniquely since the set of T_{MAX} worlds could include many disordered worlds, all of which are described by the same indeterministic law of nature, but all of which have the property of being maximally typical. For example, for a law of nature stating that $\mathrm{Ch}(x) =$

0.5, the coin flip sequence 3) above would be as typical as the reverse sequence, or the sequence where heads and tails are swapped, for example.

Moreover, a typical sequence for a law stating that $\mathrm{Ch}(x) = 0.5$ may be a very unusual sequence, such as:

5) HTHTHTHTHTHTH.

Which, according to the principles of the BSA would likely not correspond to any indeterministic law, but rather a deterministic law of the form "H always precedes T".

Perhaps some revised measure of typicality could be found that overcomes this problem. Yet, the biggest obstacle appears unsurmountable: there is no guarantee that the actual world is itself uniquely the most typical. Were sequence 5) to appear, we would find it very unusual. Were we to rarely encounter sequences of many heads in a row, we would conclude – with some justification – that the trials were not in fact IID events, but were somehow being manipulated.

Moreover, given that typicality must work over all of modal space, and the BSA work in each possible world, why should it be assumed that all indeterministic possible worlds would be maximally typical? For this to be the case, 'indeterministic' would need to be equated with 'typical', which not only contradicts the Basic Chance Principle (see Chapter 4), but needlessly classifies worlds that are not typical as deterministic, when clearly they are not. For example, the sequence HTTH may not be as typical as HTHT, but both seem like they should be within the same set of nomologically possible worlds.

Typicality also has the distinct disadvantage of being a continuous variable, with no clear threshold distinguishing a sequence that is typical from one that is not. For example, the two sequences below look very typical:

- HHTHTHTHTHTHTHTHTHTTHTH
- HHTHTHTHTHTHHHTHTHTHHTH

Yet, assuming them both to be generated by the same chance function, the second is less typical than the first. But only just. Without a clear delineation between typical and atypical, it could be that the second sequence is observed in some Humean mosaic, yet the first sequence would – in virtue of being more typical – have the better fit to the best systematization.

For all of these reasons, fit measured as typicality should be rejected in favour of something more objective and immune from the zero fit problem. Such a measure will be developed in the remainder of this chapter.

3.6 X^2 fit

The measure of fit that will be developed here avoids the zero-fit problem associated with fit-as-maximum-likelihood, as well as all the problems associated with the fit-as-maximum-typicality proposal. The measure will be based on the X^2 ('Chi-squared') statistic, commonly used in discrete statistics, and which itself is based on a measure of error between two chance distributions. It will be shown that this measure of fit, dubbed X^2 fit, serves the purpose of fitting the BSA to the actual world far better than fit as maximum likelihood by not only solving the zero-fit problem, but also being more intuitive. The possibility of using a X^2-based measure of fit in the context of the zero fit problem has also been independently proposed by Schwarz (2014).

In statistics, fit is conceived of as a measure of the likelihood that a set of observations are described by a particular mathematical function or vice-versa. In the case of discrete chance events, this amounts to measuring the likelihood that those events are randomly sampled from a theoretical population of such events defined by a probability space. In

effect therefore, statistical fit measures the likelihood that the observed sample and some theoretically expected population of chance events are described by the same random variable, or chance function. In this section it will be argued that the notion of statistical fit can be adapted to provide a measure of the fit of a system of laws to a possible world.

Statistical fit is founded on the idea of sum-of-squared error (SSE). A frequently used SSE-based measure of fit in discrete statistics is the X^2 test.

Consider a probability space $S =< \Omega, W, \mathrm{Ch}(-) >$ where $\Omega = \mathrm{O}_1, .., \mathrm{O}_m$ is a set of individual m outcomes. Let N be the total number of observations of Ω-type chance events, and $|\mathrm{o}_i|$ be the observed frequency of outcome O_i in the set of N observations, then:

6) $$\mathrm{X}^2 = \sum_{i=1}^{m} \frac{\left[|o_i| - \mathrm{Ch}(\mathrm{O}_i)\, N\right]^2}{\mathrm{Ch}(\mathrm{O}_i)N}$$

For a system of candidate laws containing m different chance functions. From (6) a probability ('p-value') can be assigned measuring the likelihood that the sample of observed values have relative frequencies of outcome values described by the probability space S.

In finite possible worlds for IID events, the mathematical form of the X^2 measure works just fine. An analogous p-value could also be assigned. However, both the X^2 measure itself and any p-value thus generated would need to be interpreted differently from the conventional (classical) statistical interpretation. Call the new interpretation for fitting indeterministic laws to the Humean mosaic the BSA interpretation and the interpretation used in statistics the statistical interpretation.

Under the statistical interpretation, X^2 measures the degree to which a finite set of observations fits a parameterized mathematical function (or vice-versa; the interpretation

of which 'direction' the fit is made is not relevant to this discussion). The parameters of the function admit of no error. This absence of error in the case of IID events with a uniform distribution is typically fleshed out by statisticians as due to the parameter measuring the properties of an infinite set, for example, or the relative frequency of a finite set of idealized outcomes. However, for the philosopher, this absence of error would be explained by some metaphysical theory of chance. For the BSA theory of chance, the absence of error is due to the parameter being ultimately derived from the entire Humean mosaic, of which there is only one per possible world, and thus each parameterization of it can have only one set of parameter values.

In the case of the BSA interpretation, the "$\mathrm{Ch}(-)$" term in (3) remains unchanged: fit is still being measured against what the BSA theory of chance says the output of the probability function should be. Likewise, the corresponding probability space remains unchanged in interpretation also. But the $|\mathrm{o}_i|$ frequencies and the total N observations in the sample now become something quite different. Under the BSA interpretation, what is being fitted is not a sample of observations extracted from the Humean mosaic, but the entire Humean mosaic itself. That is, it is not just a proper subset of, for example, all of the tosses of a coin that have ever happened or ever will happen in the actual world, but is the entire set of all tosses of a coin in the actual world. Another way of saying this is that the BSA probability function is being compared to the probability function of the actual world that would result from a naïve frequentist's metaphysics of chance, where (7) holds for all genuine chance functions:

7) $\mathrm{Ch}_F(\mathrm{O}_i) = \frac{|o_i|}{N}$

for any outcome value $\mathrm{O}_i \in \Omega$, where N is the count of all Ω-type chance events in the Humean mosaic. By substitution of $|o_i|$ by $\mathrm{Ch}_F(\mathrm{O}_i)\,N$:

8) $\mathrm{X}^2 = \sum_{i=1}^{m} \frac{[\mathrm{Ch}_F(\mathrm{O}_i)N - \mathrm{Ch}(\mathrm{O}_i)N]^2}{\mathrm{Ch}(\mathrm{O}_i)N}$

8) is measuring, for each i, the difference in the frequency of chance events between those expected by a candidate systematization and the observed frequencies of those same chance events. Thus, depending on N and m, $0 \leq \mathrm{X}^2 \leq \infty$, where $\mathrm{X}^2 = 0$ indicates a perfect fit and $\mathrm{X}^2 = \infty$ the worst possible fit given the size of the probability space and outcomes frequencies. In finite cases, 8) therefore works as an ordering function of candidate systematizations by similarity to the actual world – all that is needed for the purposes of identifying the best systematization in terms of chance functions. However, it cannot be guaranteed to identify the best systematization in the case of infinite N or m, because a trade-off against strength and simplicity would result in $\mathrm{X}^2 = \infty$ for all candidate systematizations that involved a trade off (that is, all situations where $\mathrm{Ch}_F(\mathrm{O}_i) \neq \mathrm{Ch}(\mathrm{O}_i)$). This is solvable by removing N from the equation and dividing by m to obtain a new measure $\mathrm{X}^2_{\mathrm{L}}$ such that for all cases $0 \leq \mathrm{X}^2_{\mathrm{L}} \leq 1$, as follows:

9) $\mathrm{X}^2_{\mathrm{L}} = \frac{\sum_{i=1}^{m}[\mathrm{Ch}_F(\mathrm{O}_i) - \mathrm{Ch}(\mathrm{O}_i)]^2}{m}$

The measure $\mathrm{X}^2_{\mathrm{L}}$ is proposed here as a better measure of fit between a possible world and a system of laws by avoiding the zero-fit problem, as well as any other cases of infinite valuations. That is to say, fit may be measured between a possible world and a system of laws as the mean squared sum of all discrepancies between the observed relative frequencies of chance events in the possible world described by the chance functions appearing in the indeterministic laws of the system of laws, and parameter values of those chance functions.

To adopt some statistical language which has some intuitive appeal: fit is measured as the mean difference between the observed and expected relative frequencies. The smaller the difference between the two, the better the fit, and vice-versa. So, perfect fit is obtained

when the relative frequencies for each O_i are the same as expected from the chance functions appearing in the system of laws L. That is, for all i:

10) $\mathrm{Ch}_F(\mathrm{O}_i) = \mathrm{Ch}(\mathrm{O}_i)$

In other words, perfect fit is obtained between L and some world where $\mathrm{X}^2_{\mathrm{L}} = 0$. Conversely, all worlds where $\mathrm{X}^2_{\mathrm{L}} > 0$ have imperfect fit, and because $0 \leq \mathrm{X}^2_{\mathrm{L}} \leq 1$ any trade-off against strength and simplicity cannot tend all systematizations with imperfect fit towards infinity. The process of selecting the best fitting systematization in the context of the BSA is therefore minimization of $\mathrm{X}^2_{\mathrm{L}}$ for each candidate system of laws.

Note that the complement of (9), $1 - \mathrm{X}^2_{\mathrm{L}}$, provides the basis of measuring the probability of the actual world. Namely, $\mathrm{P}(w_@) = f(1 - \mathrm{X}^2_{\mathrm{L}})$ for some probability distribution function f. Secondly, note that 9) is consistent with the assumption of all chance events being IID events, since two non-identical (with respect to the frequency or sequence of chance events) possible worlds could have the same fit. That is, $\mathrm{X}^2_{\mathrm{L}}$ is a function of relative frequencies, not chance sequences or any measure of typicality.

Further, if some L is a best-systematization, and defines a set of isonomic and indeterministic possible worlds W_L, where $\mathrm{X}^2_{\mathrm{L}} = 0$ for the actual world, this does not imply that for every $w_i \in W_L$ $\mathrm{X}^2_{\mathrm{L}} = 0$. In fact, inherent in the indeterministic nature of all $w_i \in W_L$ the value of $\mathrm{X}^2_{\mathrm{L}}$ will be greater than 0. That is, for any set of isonomic indeterministic possible worlds, there are some worlds that do not have perfect fit to the best systematization of at least one of those worlds (the world considered actual, for example).

This may be considered a trivial finding. It says nothing more than that some worlds will have different relative frequencies of chance events than others. But, it is a very serious problem for the Lewisian version of the BSA. This is because, for any world, the laws are supervenient on the Humean mosaic of the world, and this applies to indeterministic laws

as much as for deterministic laws. More specifically, the parameterization of chance functions is supervenient on the Humean mosaic of each world. Thus, the variation in relative frequencies may – beyond the limits of trade-offs against strength and simplicity – result in different parameterization of chance functions, and therefore different indeterministic laws, between different possible worlds. This is a contradiction if the BSA also defines a set of possible worlds with such variation as isonomic, where each world shares the same set of laws.

This contradiction is better known in the literature as the problem of undermining futures. This problem will be discussed in more detail in Chapter 4, where a solution will be provided that redefines the concept of a possible world to embrace the variation in relative frequencies.

3.7 Discussion and conclusion

The purpose of this chapter has been to elucidate the role that fit plays in Lewis's formulation of the BSA and, after introducing the zero-fit problem arising from its traditional interpretation in statistics, to propose an alternative interpretation using the measure $\mathrm{X}^2_{\mathrm{L}}$, where a world has perfect fit against a system of laws L, when $\mathrm{X}^2_{\mathrm{L}} = 0$.

There are a couple of objections that may be made to this alternative interpretation of fit. First, it can quite reasonably be claimed that this is not what fit is; fit – so the objection goes – is defined as being a probability measure. $\mathrm{X}^2_{\mathrm{L}}$ is not a probability measure, so it could not measure fit. The response is simply to state that 'fit' is just being used by the BSA as a label for whatever measure is used to pick out the best systematization that matches the distribution of chance events in the actual world. It does not matter if that measure is a probability lying in the unit interval, or some other measure. All that matters is whether or not it works.

The second objection is that, while the $\mathrm{X}_{\mathrm{L}}^{2}$ measure defined above avoids the zero-fit problem, it may not work well in those cases where the frequency of chance events is too small to provide a robust value for some $\mathrm{Ch}_{F}(-)$. In such cases it could happen that $\mathrm{X}_{\mathrm{L}}^{2} > 0$ for the best systematization. This objection is true, but this was equally the case with fit being measured as maximum likelihood – fit need not be perfect for the best systematization, since it is a trade-off with the other parameters of strength and simplicity.

The advantage of the measure $\mathrm{X}_{\mathrm{L}}^{2}$ to define fit is that it does not result in the actual world having 0 likelihood according to the laws of the actual world. Moreover, the issue of non-uniqueness is avoided, since the measurement of fit is between one world and a system of laws, rather than with some kind of probability measure over a set of worlds.

This concludes the discussion of the first of the three main problems being considered in this thesis. The next chapter will consider the much more well-known, and often thought of as intractable, problem of undermining futures.

4. BSA chances and the undermining problem

4.1 Introduction

The BSA as developed by Lewis is often considered to have a fatal flaw insofar as it is able to be both a theory of the fundamental laws of nature and of metaphysical chance (Briggs, 2009a; Thau, 1994). This presumed fatal problem is the problem of undermining futures, referred to hereon as the undermining problem.

Following a recap of the BSA theory of chance in section 4.2, the problem of undermining will be presented in section 4.3, first in an intuitive manner, and then in a more formal manner to highlight the contradiction that arises in Lewis's original formulation (Lewis, 1980). The problem of undermining will be presented as manifesting itself in both an epistemic as well as a metaphysical form. The presentation of the problem in section 4.3 will, contrary to its more typical presentation in the literature, focus on its metaphysical aspect and the resulting denial of a very intuitively appealing principle of chance: the Basic Chance Principle (Bigelow et al., 1993). Lewis (1994) and Hall (1994) have proposed essentially the same fix regarding the epistemic manifestation. It will be argued that the proposed solution of Hall (1994) and Lewis (1994), which will be rejected for leaving the metaphysical aspect of the undermining problem untouched, as well as opening the way to an infinite regress.

In section 4.4 a true fix will be proposed, called the *T-world proposal*. It fixes the problem by redefining the Humean mosaic from that of a single possible world to a set of branching possible worlds. This proposal is somewhat akin to the Everettian or 'many worlds' interpretation of quantum mechanics (Everett, 1957) and the interpretation of chance events that has been proposed to accompany it (Deutch, 1999; Wallace, 2012).

In parallel also with Everettian quantum mechanics, it will be shown that this set of branching possible worlds is not arbitrary, that the branching structure in particular can be well-defined within modal space, and its introduction as the supervenient base is well-motivated.

Following the presentation of the T-world proposal as a solution to the problem of undermining futures in section 4.4, its relationship to Everettian quantum mechanics (EQM) will be discussed. The T-world proposal will be distinguished from EQM as a distinct metaphysical theory, although both theories retain many structural similarities.

Finally, in section 4.5, some possible objections to the T-world proposal will be discussed.

4.2 Best-system chances

The BSA states that the fundamental laws of nature are those theorems of the best systematization of the Humean mosaic. Best-system chances (referred to simply as 'chances' hereon) may then be represented as any probability functions that appear, or are implied by, those fundamental laws of the best-system. This is Lewis's 'package deal' account of both fundamental laws and chance. In other words, the same Humean theory can account for both fundamental laws and chance by the claim that the latter is a concept defined by the former. The term 'chance' here is intended to denote the concept of metaphysical chance being identified as in the preceding chapter by conditional probability functions that satisfy the equivalence $\mathrm{Ch}_m(p|\mathrm{H}_t) = \mathrm{Ch}_m(p|\mathrm{H}_v)$, for some event p and two non-identical periods of history H_t and H_v. That is, the parameters of metaphysical chance functions remain constant irrespective of how much accumulated historical observations are made prior to the occurrence of the chance event. In simpler terms: no amount of information will allow an observer to 'hack' the chance function.

In Lewis's (1994) formulation, chance events are assumed to be both world- and time-indexed. This is to account for the chance of a non-time-indexed event evolving with time. Lewis provides an example by way of the chance of a coin landing heads not being the same a millisecond prior to it landing as when it was tossed. Lewis may have used a coin-tossing example merely to demonstrate his point, but the example is unfortunate, since coin-tossing appears not to be a fundamentally chancy event: when one considers the forces acting upon the coin, it's trajectory, and so on, it may turn out to be determinately heads or tails for an omniscient agent.

It can be debated whether the chances of all past events should be considered as 1. Past conditional statements, such as the unpicking of forensic evidence in a murder scene, attributing credence to various scenarios that led up to the murder, may be cited in such cases as reasonably interpreting the chance of any of the various scenarios being the actual course of history as being less than 1. But these are all instances where the probability calculus is being harnessed to model situations of ignorance and degrees of belief, not fundamental chance.

However, Hoefer (2007) has gone further and claimed that there is no in-principle reason to assume past events should have probability 1, and gone on to link this with the concept of admissible information that may be used in conditional chance statements, and in particular the Principal Principle to be discussed below.

While it could easily be conceded to Hoefer that there is no in-principle reason to assume past events should have probability 1, for the purposes of this thesis, it is stipulated that the chance of all past events is 1. More generally, world- and time-indexing seems appropriate for chancy events (that is, ones that satisfy the equality above), since it seems intuitively desirable for all chances of past events being 1 and necessary for them to be indexed to a particular world. In this thesis therefore, the expression '$\mathrm{Ch}(p) = x$' may be

read as an abbreviation of the proposition 'the chance of p at time t in world w is x', where the time and world indices are implied.

It follows from Lewis's package deal account introduced in section 1.4 that a world is indeterministic only if its fundamental laws include or imply probability functions. That is, since the fundamental laws are supervenient upon the Humean mosaic, the indeterministic status of a world is likewise also supervenient upon the Humean mosaic. This conclusion follows from the twin theses of Humean supervenience and Lewis's package deal account of laws and chances (1987a, postrscript C), coupled with the stipulation that indeterministic worlds are chancy worlds, and vice-versa. Humean supervenience regarding laws of nature states that laws supervene on the Humean mosaic; the package deal account claims that chances are whatever the laws of nature say they are. Therefore, if the laws of nature says a world is chancy by including or implying chance functions, and chancy worlds are indeterministic worlds, such worlds are indeterministic. Likewise, if a world is indeterministic, it is so because it is chancy, and the package deal account requires that such chance functions are included or implied by the laws of nature.

This may appear straightforward reasoning, but is a significant departure from the more intuitive account of the relationship between the indeterministic status of a world and its laws of nature: that the latter are a consequence of the former, rather than vice-versa. Moreover, the above reasoning also demonstrates that Humean supervenience and the package deal account also imply that indeterminism is a supervenient concept, since chance and the laws of nature are supervenient, and a world is indeterministic only if there are chance functions implied or included in the laws of nature.

This supervenient conception of indeterminism is also in contrast with the more standard Laplacian definition of indeterminism (as stated in, for example, Earman, 1986) when coupled with the BSA. Framed in more contemporary language, the Laplacian definition is

in terms of the divergence of worlds with common initial segments: those that diverge are classed as indeterministic; those that do not as deterministic (Laplace, 1951; Earman, 1986).

This conflates indeterminism with the Principle of the uniformity of nature: that the prevailing laws of nature are time invariant. This is so because, according to the BSA as conventionally understood, a world can be non-uniform in this sense without being indeterministic. For example, the best systematization could state that, up to some time, one subsystem of laws prevails, and then – following some anomalous discontinuity event – some other subsystem of laws prevails. But in such cases, neither subsystem of laws or the entire system need include or imply chance functions. It could just happen that at some time there is some discontinuity in the Humean mosaic and a new set of patterns appear, described by a new subsystem of laws. As discussed in section 3.2 above, such worlds would be Laplacian-indeterministic yet remain 'BSA-deterministic'.

The implication of this reasoning is that the Principle of the uniformity of nature is a contingent thesis for the Humean BSA theorist, not a necessary or metaphysical truth, since in such discontinuous worlds it no longer obtains. This implication is true, at least according to Humean supervenience. Non-reductive realists regarding laws of nature could stipulate the principle to be a necessary truth, just as they believe laws to be primitive. But for the Humean, such stipulation would be arbitrary.

The Laplacian conception of determinism could be tweaked easily enough by restricting the divergence to only those worlds that share the same laws of nature. Indeed, this may be an implicit assumption in Laplace's original statement. But this tweak implies the set of possible worlds that share the same laws of nature could be specified independently, which would necessitate their stipulation as primitive, and thus the whole Humean endeavor to provide a supervenient account would be foregone.

The proposal laid out below ensures that such divergent worlds would be classified as indeterministic and thus preserves the Principle of the uniformity of nature and all of the inductive reasoning that grounds empirical science.

With both the fundamental laws and chance being supervenient, there is no primitive notion of modality in the BSA. Nomological necessity and possibility may be reduced to an extensional, set-theoretic account of possible worlds according to Lewis's original (1986) formulation. Specifically, the best-system laws for a given world, L_w, defines a set W_{L_w} of nomologically possible worlds, defined as all worlds that share the same set of fundamental laws. An implication of this is that, for any $w_i \in W_{L_w}$, the best-systematization of w_i is the same as that of $w_@$.

This reduction of modality to extensional notions applied to the universe of all possible worlds should be considered one of the greatest advances in metaphysics, for it provides a Humean theory of laws of nature and chance, without relying on any primitive notion of modality, while also not resorting to some anti-realist or nominalist theory of laws and chances. Insofar that the problem of undermining lies at the heart of this theory, it destroys much of this theoretical work, at least in the case of indeterministic worlds. A real solution is therefore greatly needed and will be given in the latter parts of this chapter. First however, the problem itself will be introduced and examined in more detail.

4.3 Undermining futures

Put simply, the problem of undermining futures is this: the observed relative frequencies of a chance event in a world within the set W_{L_w} may not have the same relative frequencies of the same chance event in another world within the same set of nomologically possible worlds. But, because of the package deal account of laws and chances (1987a, postrscript C), this would imply that the two worlds may have different laws and if so, should not in fact be in the same set W_{L_w}, which is defined to be all worlds that share the same laws L_w.

This is because, for the BSA Humean theory of chance at least, chances are in some meaningful sense determined by the relative frequencies of chance events. That is, the BSA theory of chances is a frequentist theory of chance.

Lewis provides a succinct example using the random decay – according to the half-life law of nuclear chemistry – of a sample of the radioactive element tritium (Lewis, 1994):

> For instance, there is some minute present chance that far more tritium atoms will exist in the future than have existed hitherto, and each of them will decay in only a few minutes. If this unlikely future came to pass, presumably it would complete a chancemaking pattern on which the half-life of tritium would be very much less than the actual 12.26 years. … Could it come to pass, given the present chances? Well, yes and no. It could, in the sense that there's non-zero present chance of it. It couldn't, in the sense that its coming to pass contradicts the truth about present chances. If it came to pass, the truth about present chances would be different. (p.482)

The general form of this problem is as follows: Consider some BSA-defined chance function $\mathrm{Ch}(-)$ and associated probability space $P =< \mathrm{Ch}(-), \Omega, E >$. Consistent with this setup is some sequence of n future outcomes $\{o_1, o_1, \dots, o_n\} \epsilon E$. Call this sequence F, such that $\mathrm{Ch}(F) > 0$, and for some outcome $O_i \in \Omega$, world (that is not the actual world but is a member of the set W_{L_w} of nomologically possible worlds) w, and time t (both omitted for simplicity), the following inequality holds:

1) $\frac{|O_i|}{n} \neq \mathrm{Ch}(o_i = O_i)$

where $|O_i|$ is the frequency of outcome O_i in the sequence of length n and assume that $\mathrm{Ch}(o_i = O_i) = r$. That is, for a sequence of coin tosses for example, the ratio of heads to total tosses in w is not equal to the chance of heads for any given toss as dictated by the

chance function of the BSA. For example, the sequence F could be a long – possibly infinite – sequence of coin tosses where the proportion that have the heads outcome is 0.9. Call worlds exhibiting this inequality F-type worlds. Importantly, F-type worlds are nomologically possible because they are within the set W_{L_w}. Thus, they share the same laws as the actual world, despite the inequality above.

The problem now arises for any F-type world: in such worlds the best systematization of the Humean mosaic for that world may not have resulted in $\mathrm{Ch}(o_i = O_i) = r$, the value obtained from the best-systematization of the actual world. That is, the laws of F-type worlds, were their Humean mosaics to be systematized according to the BSA, would turn out to be different than the BSA laws of the actual world used to define the set W_{L_w}, because the relative frequencies of the same chance events differ. The reason is simply because, for the Humean (and any frequency-based theory of chance), the relative frequency of chance events has to be reflected somehow in the chance functions that appear in the laws of nature.

Note that this problem would not arise for all worlds exhibiting the inequality above. This is due to the tolerance built into the BSA resulting from the trade-off between simplicity, strength and fit: if the value of the chance function turned out to be, based on the observed relative frequencies of the Humean mosaic, to be 0.50001, for example, it could be that the trade-off with simplicity resulted in this being rounded down to 0.5. But this tolerance is likely small, and so for any given chance function and F-type world it is probably the case that the Humean chances of F-type world would be different. Nonetheless an existence proof can be provided by making the inequality so extreme that no degree of tolerance would suffice; namely, in those infinite F-type worlds where the ratio on the left side of the inequality above limits to 1 or 0, to which no chance function could approximate while still remaining a chance function (this assumes a deterministic law

is always going to be more simply expressed without a chance function than with some trivial chance function that takes on only 0 or 1).

As noted above, and in section 2.3, the problem with F-type worlds is that they contradict the Basic Chance Principle (BCP) (Bigelow et al., 1993) by breaking the implication from non-zero chance values to possible worlds.

Partly as a response to the problem of undermining, Lewis has claimed the BCP to be something not fundamental to chance (Lewis, 1994 p.483). This does get Lewis out of a problem lying at the heart of his theory, but there are two very good reasons to want to retain the BCP as axiomatic to any metaphysical theory of chance. First is simply that it is very intuitively appealing, to be able to assume that there exists – within the set W_L– for every nonzero chance outcome a world that instantiates it. Far more importantly, however, Lewis's reduction of modality is an extensional account defined over modal space: (nomological) possibility is defined in terms of membership of the set W_L. F-type worlds mean there are worlds that should be nomologically possible – at least according to our intuitive notions of chance and possibility – but are not in W_L, whether the BCP is retained or not.

The solution to the problem of undermining presented in section 4.4 below will rely on the fact that F-type worlds imply the falsity of BCP in a *modus tollens* that demonstrates that the solution solves both the problem of undermining while also retaining BCP.

To recap the reasoning so far: the possibility of F-type worlds implies that two worlds in the same set of possible worlds sharing the same laws, W_L could turn out to have different patterns of chance events. This in turn implies by the BSA that their indeterministic laws would be different, and therefore could not both be in the same set of possible worlds sharing the same laws, W_L after all. Thus, an apparent contradiction and the Basic Chance

Principle is violated: there are non-zero chances but no world within the set W_L that instantiates such a non-zero chance outcome.

This is problematic for the reductive account of nomological modality propounded by Lewis because the overarching program of Humean supervenience needs a supervenient account of modality. Without this the program fails, since its most attractive feature is to show how modality can be built up from nothing more than the properties of the Humean mosaic of each world.

The problems that F-type worlds have regarding reductive accounts of modality and nomological character are serious enough. But for those such as Lewis (1994 p.483), who are happy to ditch the BCP as an axiom, there is a further epistemological problem with F-type worlds. The epistemological problem is how F-type worlds relate to Lewis's Principal Principle (PP) introduced in section 2.3 above. This principle establishes a link between chance and initial credence, or subjective degree of belief, denoted below as the '$\mathrm{Cr}_0(-)$' function:

$$\text{Principal Principle (PP): } \mathrm{Cr}_0(A|T, H_{tw}) = \mathrm{Ch}_{tw}(A)$$

That is, the initial credence an epistemic agent should assign to the proposition that the chance at t of proposition A, given the total amount of elapsed history, H_{tw} , and the theory of chance T is equal to the chance of A, $\mathrm{Ch}_{tw}(A)$.

The term 'initial credence' here denotes credence in the absence of any other information that could change the chances of A. That is, the epistemic agent is assumed to be in a tabula rasa state in receipt of no observational data regarding the world in which the agent inhabits, except the information on the right side of the conditional T, H_{tw}, namely: all past history and the theory of chances in w. There is considerable debate about what actually constitutes admissible information for initial credence functions (see for example: Lewis,

1980; Hoefer, 1997; Ward, 2005; Hall, 2011), and the truth of PP depends in part on having a consistent and precise definition of 'admissibility'. However, for the purposes of this discussion this issue will be passed over to focus instead on the problem of undermining.

Although the PP may not turn out to be a fundamental principle of reason, instead being derivable from a principle coordinating chance and credence, coupled with a valid definition of admissibility (Schwarz, 2014), it nonetheless remains an attractive principle to retain, if not as an axiom then at least as a theorem.

The epistemological problem associated with F-type worlds is that substitution of A with F, an undermining future that is inconsistent with the theory of chance for the reason given above, means that the right-hand side of the equation must equal 0. This is because as argued above, F-type worlds imply the contradiction that they both are and are not members W_L, and the credence we assign to a contradiction must be 0. But according to the definition of F, $\mathrm{Ch}_{tw}(A) > 0$. Thus, the further contradiction that $\mathrm{Ch}_{tw}(F) = 0$ by the PP, and that $\mathrm{Ch}_{tw}(F) > 0$ by the definition of F.

In response to this, Lewis (1994) and Hall (1994) proposed an amendment to the PP, effectively making all initial credence valuations conditional on the theory of chance itself:

2) $\mathrm{Cr}_0(A|T, H_{tw}) = \mathrm{Ch}_{tw}(A|T)$

This new PP is claimed to have solved the problem of undermining by forcing $\mathrm{Ch}_{tw}(F)|T$ to equal 0, given that $\mathrm{Ch}_{tw}(F)$ would be inconsistent with T in the case of F-type worlds.

It will be shown here however that this apparently minor amendment is problematic. To see this, consider what is implied by making all chance evaluations conditional on the theory of chance T. For this to be meaningful, such conditionalization must have some effect on the final chance evaluation: sometimes T is true, and sometimes (in the case of F-type worlds) it is not. If this were not the case, 3) would be equivalent to the original PP.

However, if there is a chance that T is true and some non-zero chance that it is not true, this implies the existence of some further chance function ranging over T itself. Call this further chance function $\mathrm{Ch}^{*}_{tw}(-)$. Note that this chance function is of a different form than the family of BSA-derived chance functions of the form $\mathrm{Ch}_{tw}(-)$. This is because the $\mathrm{Ch}^{*}_{tw}(-)$ variety applies to the theory of chance, T, itself, rather individual chance events.

Hall (1994) provides an illuminating interpretation of this amendment, as follows: Imagine some omniscient epistemic agent who knows all the fundamental laws of nature and chances and all of history. This agent tells you that $\mathrm{Cr}_0(A|T, H_{tw}) = \mathrm{Ch}_{tw}(A) = a$ for some A. As a rational agent, you would – Hall presumes – only believe this to the extent that you trusted the omniscient epistemic agent to know and tell the truth. That is, your degree of belief $\mathrm{Cr}_0(A|T, H_{tw})$ cannot exceed your degree of belief in the omniscient epistemic agent to know and tell the truth. Thus, making the right side of the equation conditional upon this results in the above amendment.

Hall's interpretation may appear sensible. As an epistemic theory given the scenario of an omniscient epistemic agent, the amended version of the PP may well be correct. But this solution to the undermining problem is problematic. First, the theory of chance, T, now appears on both sides of the equation, and so this revised form no longer says anything about what chance is. In contrast, the original PP said that chance conformed to the credence of a rational agent with admissible information, and thus it was something knowable. It is true that the amended PP still constrains credence, to conform to chances given the theory of chance, and so remains an interesting epistemological principle. But so long as the theory of chance appears on both sides of the equation as the antecedent condition, and nothing is said about the boundary conditions of when $\mathrm{Ch}_{tw}(A)|T = \mathrm{Ch}_{tw}(A)$ – that is, when the theory of chance works and when it does not – it cannot be considered an explanatory principle about chance.

Second, because the new principle introduces the new chance function $\mathrm{Ch}^*_{tw}(-)$, it opens up a potentially infinite regress in explanation. This is because the function $\mathrm{Ch}^*_{tw}(-)$ cannot be implied by the BSA theory of chance, since it ranges over the BSA theory of chance itself. Moreover, as noted by Hoefer (1997), the function $\mathrm{Ch}^*_{tw}(-)$ must originate outside of the Humean mosaic, since it is unrelated to the chance functions of the BSA. Thus, both Lewis and Hall have let slip into their metaphysics non-Humean elements.

To summarize: the $\mathrm{Ch}^*_{tw}(-)$ function must denote some form of alethic or epistemic modality, saying something about what may or may not be true about the world or what we can know about it, rather than the distribution of chance outcomes in some real or imagined probability space. But its introduction says nothing about in what circumstances the theory of chance is correct and when it is not, so in explanatory terms, its introduction just pushes the problem back one step further.

Finally, and most importantly for the purposes of developing a fully consistent Humean theory of laws and chances, it does not solve the metaphysical problem. This is because Lewis and Hall's solution is directed only at the epistemological problem of undermining futures: F-type worlds still exist, it's just that they can no longer generate contradictions with the new PP. Thus, at least for those who prefer to retain the BCP, F-type worlds have non-zero chances, yet would not be within W_L.

Another proposed solution to the problem of undermining has been proposed by Halpin (1998, 2003), whereby the fundamental laws are made world-relative to allow cases where, from the perspective of one world, another world (specifically, some F-type world) has a different set of laws than it would have from its own perspective. This relativization thus makes the set of nomologically possible worlds, W_L, nomologically possible only relative to some member world of W_L, such as the actual world, which itself would be nomologically impossible relative to F-type worlds, for example. Under this so-called perspectivalist

account, the laws L_i of possible world i defines a set of possible worlds $W_{L,i}$. But crucially, this set represents the set of nomologically possible worlds when considered from the perspective of i and the laws L_i. From the perspective of i, all other worlds in $W_{L,i}$ have the laws L_i. But from the perspective of some world $j \neq i$, $j \in W_{L,i}$, j has laws L_j that define a set of worlds $W_{L,j}$, and there is no reason to assume that either the laws or the set of possibles worlds defined by either i or j are identical. In fact, given the existence of F-type worlds there is every reason to assume they are not.

Halpin claims his account avoids the problem of undermining. He claims it does so by allowing the laws to – in effect – change according to the perspective of the observer. F-type worlds can still exist, but only relative to the laws of some index set of worlds: the best systematization of the laws in F-type worlds would have different parameter values of any chance functions than those from the perspective of the index set of worlds.

There are two problems with Halpin's perspectivalist account. First is that it appears no longer to be a full-blooded realist account of lawhood. It is not obviously anti-realist, but if it is assumed that it avoids the contradiction involved in the problem of undermining, it does so by virtue of the relativization of laws to sets of possible worlds. That is, the laws of a world depend not only on the features of the Humean mosaic of that world, but also on its relationship to other worlds. As a Humean realist account this is counterintuitive, since it means that laws are not determined by the world or its Humean mosaic in isolation. The result of Halpin's account is therefore that there is no objective fact of the matter about what laws and chances are, because they are perspective-dependent. The cost of apparently avoiding the problem of undermining is resort to relativism, and possibly anti-realism.

However, there is a more serious problem for Halpin's account: it still fails to avoid the problem of undermining. Briggs (2009a), who identifies Halpin's perspectivalist account as

anti-realist, has claimed that our intuition on the meaning of chance is that it involves rigid designation. That is, she claims that our intuition about what chance means is that chance valuations do not vary between (nomologically) possible worlds: the chance of a coin landing heads in j should be the same as the chance of its counterpart doing so in i, irrespective of whether one of these is an F-type world. As Briggs states, for some chance statement A, "A is [necessarily] true just in case A is true at the actual world of every model which cannot be ruled out a priori" (Briggs 2009a, p.89). This rigid designation assumption implies that the relativization of chances to the perspective of one possible world does not work and Halpin's proposal would still be open to being undermined.

Briggs' argument does presuppose Halpin's account to be anti-realist, thereby accepting that the metaphysical problem of undermining is avoided. It is assumed by Briggs that the epistemological problem of undermining arises due to the perspective being dependent on an epistemic agent's location in modal space. Nonetheless, a more obviously realist, or at least non-relativistic, account of lawhood would be desirable. Such an account will be provided in section 4.4 below.

4.4 The T-world proposal

The proposal that will be developed here is of a modified BSA theory that retains chance as a reified Humean supervenient concept, but which can be shown not to be undermined by F-type worlds. As a result, the original formulation of the PP remains sound and the BCP is also retained.

Before the proposal is presented, some new nomenclature is required. The term 'possible world' will be disambiguated between three different meanings, and the meaning of 'actual world' made more specific, as follows:

- **Branch world ('B-world')**: a B-world is any spatiotemporal region that is not connected to any other spatiotemporal region, but which has some part of it that is qualitatively indiscernible from some other non-identical spatiotemporal region of some other B-world. Of the three notions introduced here, B-worlds most closely corresponds to the notion of a possible world specified in Lewis (1986).

- **Tree world ('T-world')**: a T-world is a set of B-worlds all of which have the same part in common and where the common part is the initial part (in time).

- **Modal locus**: This is the more general concept that includes both B-worlds and T-worlds. A modal locus is any set of B-worlds or T-worlds able to play a role in an extensional reductive analysis of some form of modality (broadly construed to include chance as a modal concept).

- **Our Tree**: the T- world that is actual (relative to an observer in some B-world part of a T-world, as in the indexical sense of 'actuality' proposed by Lewis (1986 p.92)).

- **Course of history**: A set of B-worlds that are all members of the same T-world that are all either the common initial segment of the T-world, or ordered by the 'is later than' relation.

In more familiar language, B-worlds may be thought of as being Lewisian possible worlds, or at least a subset of them: all those Lewisian possible worlds that are either deterministic in their entirety, or are finite time slices of indeterministic Lewisian possible worlds (and thus also Lewisian possible worlds in their own right). T-worlds in contrast may not be mapped onto Lewisian possible worlds in many cases. Specifically, in the case of indeterministic T-worlds (defined more comprehensively below), the mapping relation is many-to-one between T-worlds and Lewisian possible worlds. Therefore, T-worlds may not be physically identical to some Lewisian possible world, as is the case with B-worlds, simply

because they may comprise multiple B-worlds. The special case, as will be discussed below, is the case of deterministic worlds, where, because there is only one course of history, the T-world is a singleton set of B-worlds, and thus physically identical to some B-world.

Thus, for both deterministic and indeterministic T-worlds, a course of history may be thought of as segments (a single segment in the case of deterministic T-worlds) with non-zero duration of traditionally conceived spatiotemporally disconnected, Lewisian, possible worlds. Although spatiotemporally disconnected, the ordering relation may still place them on a timeline and thus, when next to each other on this timeline, be thought of as spatiotemporally contiguous in that sense. In the case of deterministic T-worlds of course, the issue is moot as the ordering relation is over a single B-world.

To disambiguate, henceforth throughout this Chapter and elsewhere in this thesis as required, B-worlds will be formally identified by a superscript 'B' thus: w_i^B; T-worlds similarly thus: w_i^T. Non-superscripted letters will denote the more general concept of modal locus. Further disambiguation of terms such as 'part' and 'in common' will be made below as needed. The term 'branch' above is not intended to denote spatiotemporal connectedness. Thus, in the figure below, in both cases a T-world composed of B-worlds is depicted. The nomenclature above does not distinguish between spatiotemporally overlapping and non-overlapping common initial segments, nor between qualitatively indiscernible but non-identical and numerically identical segments.

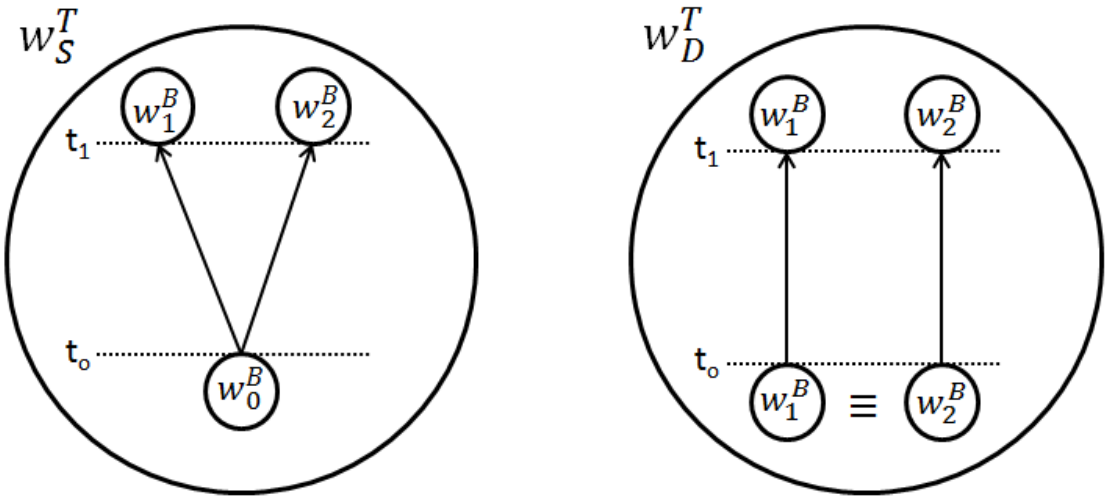


Figure 4.1: Left: a single T-world, w_S^T composed of a B-world w_0^B splitting into w_1^B and w_2^B, and right: another T-world w_D^T composed of two numerically distinct B-worlds w_1^B and w_2^B with an qualitatively indiscernible initial segment (denoted by the '≡' relation) that subsequently diverge.

Note that the notion of a B-world above is more restrictive than the broader notion specified by Lewis (1986) since each B-world has some part of it that is qualitatively indiscernible from some other non-identical spatiotemporal region of some other B-world. That is, for example, the final time slice of one B-world is qualitatively indiscernible from the initial time slice of some other B-world. It is this restriction that allows the ordering relation of B-worlds into different courses of history. Thus, there are Lewisian possible worlds that are neither B-worlds or T-worlds. Such worlds need not be considered here as they have no bearing on the broad applicability of the proposal below.

Note also that every B-world is – trivially – a T-world. To express a B-world being a part of a T-world, the conventional set-theoretic formalism will be used. Whether the relationship between a B-world and a T-world, or between a B-world and a course of history is better described using mereological concepts instead also need not be considered here.

An initial proposal to modify the BSA theory of chances to avoid the problem of undermining while still retaining the applicability and utility of the original PP without

resorting to anti-realism is the following: change nothing in Lewis's original formulation of best-system chances, with the exception that the best systematization is supervenient on the Humean mosaic be of a T-world rather than that of any B-world. That is, for any T-world, w_i^T, the entire map of space-time points of w_i^T constitutes the Humean mosaic of w_i^T over which best-system laws and chances are defined.

Let the '@' subscript indexically pick out Our Tree, or some B-world component thereof. Chances and other nomological concepts are therefore, according to this proposal, supervenient on Our Tree, $w_@^T$, not $w_@^B$ or some set of B-worlds, $\{w_@^{B1}, w_@^{B2}, \dots\}$ comprising the actual course of history.

More concisely: for any B-world, w_i^B, the laws and chances in that world are supervenient on some T-world w_j^T, where $w_i^B \in w_j^T$. Further, this set of laws defines a set of isonomic T-worlds, W_L, that share the same best system laws, L. Further, since each T-world is comprised of at least one B-world, it also defines a set of isonomic B-worlds that each w_i^B is a member of some $w_j^T \in W_L$.

Call this proposal the T-world proposal for abbreviation. It will be shown below how this change to the Lewisian 'package deal' account of laws and chances (1987a, postrscript C) resolves the problem of undermining. For the moment, the situation is depicted graphically in Figure 4.2 below for two sets of isonomic T-worlds.

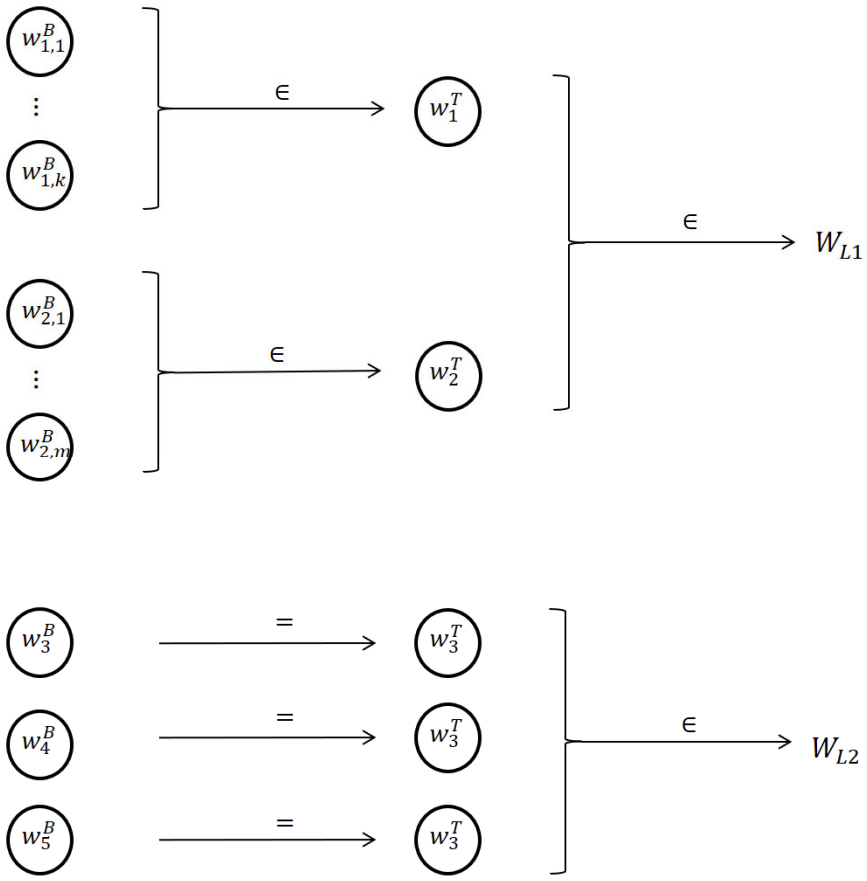


Figure 4.2: Two sets of B-worlds comprising two T-worlds, both of whom share the same laws (top half); multiple singleton sets of B-worlds, each one being identical to some T-world, all of whom share the same (deterministic) laws (bottom half).

As with the Lewisian package deal account of Humean supervenience with BSA laws and chances (1987a, postrscript C), a set of possible worlds is defined, corresponding to the set of all possible worlds that share the same laws. Apart from the obvious difference that the term 'possible worlds' refer to T-worlds in the T-world proposal (as opposed to a traditional Lewisian (1986) possible world in Lewis's original package deal account), there is also major changes needed in the process of identifying and defining the best-system and resulting parameters of chance functions also.

Specifically, the T-world proposal requires a more convoluted, stepwise, process of reading off the laws from any given T-world's Humean mosaic. In the original Lewisian version, the

process was simple: determine the best-systematization for a world, and you have the laws and chances of that world. For the T-world proposal a four-step procedure is required to avoid circularity, at least in the case of indeterministic worlds. This process is depicted in Figure 4.3 below.

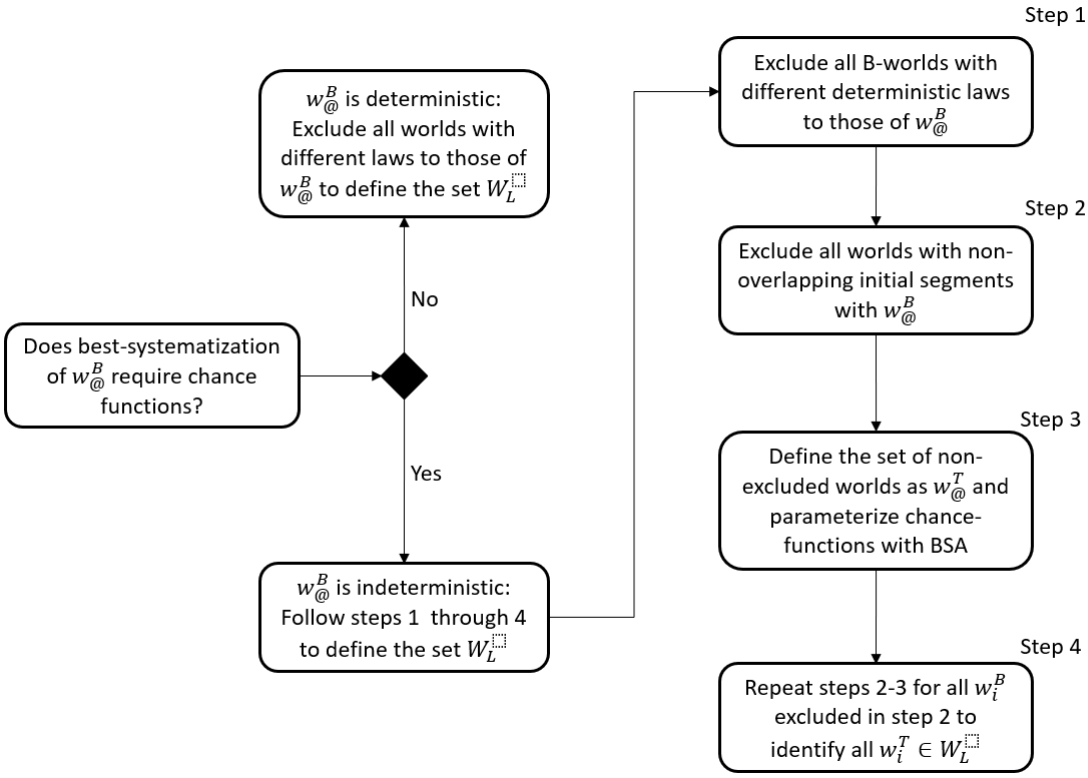


Figure 4.3: the four-step BSA procedure for determining chances in indeterministic B-worlds according to the T-world proposal using the actual B-world and T-world, $w_@^B$ and $w_@^T$ respectively, for illustration.

The first step in the best systematization process under the T-world proposal is to determine if the B-world is deterministic or not. That is, whether the B-world is a singleton T-world or not. As argued above, the indeterminacy or otherwise of a T-world is itself a supervenient property of the Humean mosaic of the B-worlds that compose the T-world.

Further, according to the package deal account and the notion of fundamental chance being used throughout this thesis, if a B-world or a T-world is indeterministic its fundamental laws will contain or imply chance functions. Likewise, if the best

systematization of a T-world does not contain or imply chance functions, it will be deterministic, and therefore be a singleton set of component B-worlds.

As mentioned above, whether or not a B-world is indeterminate or not can be addressed by identifying the nature of the best systematization of its Humean mosaic: if the best systematization would require or imply chance functions, it is indeterministic, if not, it is deterministic. To determine if the T-world is deterministic or not therefore, the BSA laws from any component B-world can be examined to see if they contain or imply chance functions: if they don't then the T-world is deterministic, is a singleton set of B-worlds, and the BSA process works in the same manner as originally proposed by Lewis. That is, the best systematization can proceed with optimization only of strength and simplicity. But if they do contain chance functions, the four-step BSA procedure for indeterministic T-worlds depicted above would be needed.

Notice that for any T-worlds subsequently identified as indeterministic, the first systematization of the Humean mosaic of a component B-world is not to determine all the laws of the T-world, but only to establish whether there are indeterministic laws and thus whether the B-world is part of an indeterministic T-world or not. Patterns of events or regularities in any component B-world's Humean mosaic can be identified as chancy (or, rather, more concisely described using a chance function) without needing to have a parameterized chance function. To see how this can be, consider the case of flipping a coin: the observed pattern can be seen to be best described by a chancy law rather than a deterministic one, without knowing the exact value of $\mathrm{Ch}(\mathrm{Heads})$.

It may be claimed that some indeterministic events may appear as deterministic in some worlds where, for example the sequence of chance events exhibit some form of regularity. In such cases, the B-world would be classed as deterministic. Notice in this line of reasoning however that there is a presumption of some kind of non-Humean nature to determinacy

or otherwise. That is, that the 'true' nature of an apparently deterministic world is indeterministic. The Humean can make no sense of this notion. Yet, such extreme worlds are clearly possible, and this issue will be discussed in more detail in section 4.5 below.

Notice also that there is no reason to assume that this first systematization should not establish the deterministic laws of indeterministic and deterministic T-worlds alike. It is assumed some indeterministic T-worlds can also have some laws that are deterministic, and that the parameters of any deterministic laws do not depend on any parameters of any indeterministic law.

This line of reasoning allows for the first step of the four step BSA procedure applicable to indeterministic T-worlds to be the exclusion of nomologically impossible B-worlds from the set of all possible B-worlds. These nomologically impossible B-worlds are determined to be nomologically impossible based on their deterministic laws alone, which have been established during the first systematization process. Those T-worlds that have different deterministic laws are thereby excluded through exclusion of their component B-worlds.

The remaining B-worlds are therefore those worlds with the same deterministic laws, but may have different initial segments. The second step is therefore to exclude those with different initial segments. The resulting set of B-worlds is a T-world. Specifically, if we start with the actual B-world, $w_{@}^{B}$, the resulting set of worlds is Our Tree, $w_{@}^{T}$.

According to the T-world proposal, in the case of indeterministic worlds, it is the Humean mosaic of this T-world that results from the exclusion of all other possible worlds not sharing the same deterministic laws or same initial segment from which the BSA laws of indeterministic B-worlds are read. This second systematization of a Humean mosaic allows for any chance functions contained or implied by the laws to be parameterized – something

that could not have been done definitively in the first systematization process since at that point it is unknown if the B-world is typical (in the sense of not being an F-type world).

The final step in the four-step process is to repeat steps two and three above for all B-worlds excluded during step two (that is, all B-worlds with the same deterministic laws but different initial segments). This will result in multiple T-worlds, all of whom share the same deterministic *and* indeterministic laws, but which have different initial segments. In other words, the set W_L. Once the set W_L is obtained, the process is complete, since the BSA laws are established, as well as the set W_L of all nomologically possible worlds.

The process above involves four separate steps. This is not to say that the process may be collapsed, possibly into three steps for example by combining the second and third steps into one single step that identifies all worlds sharing the same deterministic laws and overlapping initial segments.

It is a crucial assumption in the above process that the systematization for deterministic and indeterministic laws can be independent of each other. Specifically, that the deterministic laws are independent of the indeterministic laws, or at least of their parameter values. It may be argued that this claim is not sound. Consider, for example, the case of the Born rule in quantum mechanics. Briefly, the Born rule is a law (claimed to be an axiomatic law, in some interpretations – see Chapter 5) that specifies the chance of different measurement outcomes being realized. There have been numerous attempts to prove that the Born rule can be derived from other laws of quantum mechanics, and is thus not axiomatic (see, for example, Wallace, 2012). However, it is often claimed that all these attempts are in fact circular (Baker, 2007; Barnum, 2000). For the purposes of this discussion, both possibilities need to be considered, irrespective of which is right for the actual world.

Consider first the situation where the Born rule is indeed derivable from the other (deterministic) laws of quantum mechanics. This case is indeed problematic for the step-wise procedure above but can be easily corrected. Namely, in the first step, the process identifies those wholly deterministic worlds AND those indeterministic worlds where there are no chance functions whose parameters are NOT derivable from the deterministic laws. The remaining worlds – those indeterministic worlds where at least one chance function must be parameterized independently – then go through the four-step process depicted above.

This second situation, where some chance functions must be parameterized independently, is simpler to deal with: by definition, they must be parameterized independently of the deterministic laws (or deterministic components of mixed deterministic and indeterministic laws), so the issue does not arise.

To be any different from Lewis's original formulation, this proposal requires that there be non-trivial T-worlds. 'Non-trivial' in this context means having multiple courses of history (i.e. component B-worlds) within a single T-world. However, it does not require that such T-worlds have numerically identical common initial segments and is thus consistent with Lewis's preference for divergent possible worlds with counterparts and world-bound individuals. The other difference to Lewis's original formulation is that it avoids the problem of undermining, as will be shown in section 1.5 below

4.5 T-worlds and the problem of undermining

It will now be shown that the T-world proposal avoids the metaphysical problem of undermining associated with the BSA theory of fundamental laws and chances, and by implication, the conflict with the PP. The structure of the proof that the T-world proposal is unaffected by the undermining problem, U, will be as a *modus tollens* as follows: i) U

implies the falsity of the Basic Chance Principle (BCP), and ii) the T-world proposal ensures the truth of BCP, and therefore iii) by *modus tollens* the falsity of U.

It is a consequence of the problem of undermining, metaphysically conceived, that it contradicts the BCP (as defined by Bigelow, Collins, and Pargetter, 1993) when applied to Lewisian worlds, or B-worlds, as follows:

> Suppose $\mathrm{Ch}_{tw}(A) = x > 0$. Then A is true in at least one of those B-worlds, w', that matches w up to time t and for which $\mathrm{Ch}_t(A) = x$.

This is evident, since it is a consequence of undermining that for any undermining future F, $\mathrm{Ch}_{tw}(F) > 0$ by the definition of an undermining future. Yet, there is no B-world within the set W_L such that F is true because the theory of chance for such a B-world would have been different, and thus could not be a member of W_L.

Now assume modal realism (MR): that for anything that (logically, broadly construed) could happen, does happen in some Lewisian world or B-world (where B-worlds are assumed to be concrete entities, just as Lewisian worlds). If, for some T-world w_i^T, and some chance proposition A that is implied by the best systematization of w_i^T such that $\mathrm{Ch}_{tw}(A) > 0$ at some B-world w_m^B and time t, then for the BCP to be satisfied, there would have to exist some B-world w_n^B that is both a member of w_i^T and where A is true. This is guaranteed by MR since anything that is nomologically possible is also logically possible, and it is nomologically possible that A is true in some B-world and time t, since by definition $\mathrm{Ch}_{tw}(A) > 0$.

Therefore, MR guarantees BCP if the Humean mosaic is defined over the set w_i^T. Since some undermining future F implies the falsity of BCP under Lewis's account, but under the T-world proposal it does not, the problem of undermining is avoided by *modus tollens*.

Assuming that $\mathrm{Ch}_{tw}(A) = x, x > 0$, then because all B-world members of w_i^T are isonomic (defined as sharing the same best systematization of laws, and therefore chance functions), for any two distinct B-worlds i and j within w_i^T, chance proposition A and time t, $\mathrm{Ch}_{ti}(A) = \mathrm{Ch}_{tj}(A)$.

It could be objected that for some B-world w and time t, $\mathrm{Ch}_{tw}(A) \neq x$ because the relative frequencies in w are different from those in other B-worlds that are more typical in terms of their relative frequencies matching the best-system chances under the T-world proposal. But it is a stipulation of the T-world proposal that chance is defined over T-worlds, not B-worlds or courses of history, and by implication the relative frequencies in any particular course of history may diverge from the relative frequency predicted by the chance function as a limiting value. It is the advantage of the T-world proposal to allow such divergence and thereby have better fit to our intuitions about chance as well as avoid the problem of undermining. For example, if you believe a coin to be fair, yet after 9 tosses each landed heads, your distribution of credence between heads and tails on a tenth toss would remain unchanged, as would the chance of heads on the tenth toss. Yet, in a ten-toss B-world this may be an undermining future.

4.6 Discussion

The essential idea behind the T-world proposal – that chances should be parameterized from multiple possible worlds – has already been proposed by Tosh (2014), albeit in the context of a less general metaphysical thesis. His proposal is a correction of finitist frequency theories of chance, rather than of the BSA per se, and therefore does not involve the four-step process of determining the best systematization of laws as described above. Using his terminology, Tosh's proposal can be summarized as follows: instead of calculating the relative frequencies from the actual world alone, $c_@$, consider all counterpart worlds in the set C, such that $c_@ \in C$ and for all $c_i \in C$ the counterpart relation is one of exact

similarity of all past light cones. Relative frequencies for any future chance event then ranges over all worlds in the set C, and the desired result is obtained that F-type worlds are included in this calculation and thus undermining is avoided.

Tosh's proposal goes some way to a consistent account of BSA laws and chances by avoiding the problem of undermining, but two major problems remain. First, the account is specifically of chance functions, not of the package deal for laws and chances (1987a, postrscript C). Second, for reasons discussed below, Tosh formulates the account without the accompanying thesis of modal realism, which is problematic for these kind of accounts, as will also be discussed below. In other respects, however, Tosh's account has the same advantages and may provoke the same kind of objections as the T-world proposal.

It can be reasonably claimed that the T-world proposal and the BSA theory of lawhood depicted in Figure 4.3, while avoiding the problem of undermining, loses the elegance and simplicity of Lewis's original theory. Elegance is a somewhat subjective concept, but it is certainly true that the T-world proposal requires the introduction of additional steps and discrimination between different types of modal loci (B-worlds and T-worlds). But these steps are non-arbitrary and well-justified given the importance of solving the problem of undermining while retaining a realist and reductive theory of chances. In this sense, the solution is far superior to those who throw up their hands in exasperation and allow primitive chance into their ontology or go in the reverse direction and deny the real existence of chance. Moreover, in comparison to the fudge of Lewis (1994) and Hall (1994), it has the great advantage of retaining the Principal Principle in its original formulation.

The question arises of what is the truthmaker for chance statements under the T-world proposal. In the original Lewis (1980) formulation, chance statements were made true by a supervenient pattern of the Humean mosaic. It is the same for the T-world proposal, except

that the Humean mosaic is extended to that of a T-world rather than any individual course of history.

The truthmaker of chance statements may also be conceived subjectively in terms of some kind of credence measure, linked to the metaphysical conception by principles of reason such as Lewis's Principle Principal. When interpreted epistemologically, it could be some form of self-locating uncertainty or degree to which we care for alternative branches and their contents, as have been developed by Greaves and Myrvold (2010), for example, in the setting of Everettian quantum mechanics. In this sense, belief in chance statements could be interpreted as our perception of the branching structure of Our Tree. This interpretation of chance statements is complementary to that above, but is dependent on epistemic agents, and so for anyone preferring a realist, mind-independent truthmaker, the first interpretation may be preferred.

The T-world proposal solves the undermining problem for BSA chances and indeterministic laws, while retaining the BSA as an essentially realist theory of lawhood. However, there may still be some broader concerns on the whole Humean approach to lawhood and chances. From a purely intuitive perspective, an argument often used against Humean theories of lawhood is that they are unable to match our counterfactual intuitions regarding uninstantiated laws. The following simplified intuition-based claim from Carrol (1994) may be used in an argument against Humeanism: consider two worlds (Lewisian worlds or B-worlds), w_1 and w_2, both of which have ten particles of types p_1 and p_2 moving through a field; when particles of type p_1 exit the field they are spin down; when p_2 particles do so they all exit spin up. The only difference is that the single p_2 particle in w_2 happens not to enter the field and remains spin down. This claim is then made that humans are inclined to believe, given this setup, that had the p_2 particle entered the field in w_2 it would have exited spin up, just as in w1. This is a problematic conclusion for any

Humean supervenient account of laws, since it assumes that laws need not be supervenient on the Humean mosaic; the counterfactual event of the p_2 particle entering the field in w_2, does not exist (in w_2), and so cannot form part of the supervenient base.

As a matter of methodological approach, such intuition-based claims will not be considered as valid in this thesis without a clear demonstration that human intuition has veridical metaphysical insight. Nonetheless, two points regarding such arguments are worth making.

Firstly, Bebee (2000) has shown that arguments of the type proposed by Carrol are ultimately question-begging for a Humean by presuming a non-Humean stance at inception. More positively however, in so far that neurotypical pre-theoretic human intuition may be concordant with Carrol's, the proposed T-world proposal above can provide an alternative response. This is because, if it were by chance that w_2 in the above intuition claim had no p_2 particles entering the field, then if such chance is determined by the Humean mosaic of a tree world, and if w_2 is a branch world of an indeterministic tree world, then it follows that there would be another branch world w_2^* where the particle does enter the field and exits spin up. Because world w_2^* shares the same laws as world w_2, it should be more similar to some other world not sharing the same laws, since there would be no 'gross violation' of laws. Further, because at the time just before the chance event, worlds w_2^* and w_2 would share a common initial segment, and therefore have exact similarity. Using Lewis's similarity metric therefore, it is reasonable to assume that the existence of world w_2^* would support the counterfactual and make it true.

The above line of reasoning that the T-world proposal allows for a much broader set of chancy worlds to be considered when the fundamental laws are formulated goes some way to satisfying non-Humean intuitions. But it does not go all the way. In particular, it does nothing for the case of a deterministic world where, according to some non-Humean conception of 'chance', one variable considered by the best-system to be deterministic, is

'in fact' indeterministic. The Humean can make no sense of this intuition however, since any appeal to further facts beyond the Humean supervenient base that determine modality are either fruitless or ending in an infinite regress. So, while the T-world proposal does satisfy some of our non-Humean intuitions, it cannot satisfy them all. But no Humean accounts seems able to do this, and for those who consider human intuition not to be a source of metaphysical insight, this should not be a cause for concern.

Another advantage of the T-world proposal is that it introduces a subset of modal loci (i.e. all members of W_i^T), in addition to the set of all nomologically possible loci consistent with the BSA laws of Our Tree. Namely, the set of B-worlds that are part of Our Tree. These B-worlds play a modal role, but one that is particular to chance. Specifically, all modal loci within this set satisfy the following condition:

$$\mathrm{Ch}_{tw}(X|T, H_{t@}) > 0$$

for all X, and for any chance function.

In contrast, because of the problem of undermining, Lewis's formulation could not define the set W_L without contradiction. Moreover, because chance was defined over a single course of history under Lewis's account, rather than a set of B-worlds, the only modal distinction Lewis could make without further stipulation was that between the nomologically possible and the nomologically impossible, not between those worlds satisfying the above condition and those not, and hence some finer chance-based distinction.

One special case of F-type worlds is where the F-type world exhibits 'deterministic' behaviour. Consider a coin-tossing world. If the coin is unbiased, then the chance is 0.5 of landing heads. According to this set up there is one world where all tosses land heads. But the BSA laws of this world may have no need for any non-trivial chance function according

to the Lewisian formulation, since it would be a world that at least exhibited deterministic behaviour: the coin always lands heads. In other words, such a world is an F-type world according to the definition above. Such an F-type world would not only undermine the chances of the actual world, they would also undermine the metaphysical status of the world as being indeterministic.

Such extreme F-type worlds are problematic for the T-world proposal as they demonstrate that the T-world proposal cannot solve the problem of undermining completely. Thus, while the argument above works perfectly well for most worlds, in some cases the BCP may fail insofar that extreme F-type worlds may not be classed as being within the set W_i^T of some isonomic indeterministic T-worlds. That is, it cannot be guaranteed that in such extreme cases, step 1 (from Figure 4.3 above) in the best systematization of laws would not identify the world as wholly deterministic and therefore a singleton T-world. Moreover, this would preclude the world also being part of another T-world where one or more B-worlds exhibited more obvious indeterministic behaviour.

However, while the above possibility exists, there is reason to believe that in any interesting B-world such extreme worlds that demonstrate apparently deterministic behaviour would not be classed as deterministic under the T-world proposal. This is because describing the apparently deterministic behaviour by a deterministic law may be inconsistent with other laws. That is, other laws may imply that the apparently deterministic behaviour is actually indeterministic behaviour that just happens to be occurring in an extreme B-world.

To demonstrate this line of reasoning, consider double-path experiments in quantum physics such as the double-split experiment, where the path the quantum particle goes down is described by a chance function. Consider what would happen in the actual world if it turned out that, for example, all photons ever tested in such an experimental setup went

down one path rather than another. Physicists would find this strange behaviour, as it would be a peculiar asymmetry that would not be derivable from the deterministic components of quantum mechanics. In fact, such behaviour may even be contradictory if it were not interpreted as an extreme example of a chance distribution. Still, this is speculation, not proof, that the T-world proposal would not necessarily misclassify indeterministic B-worlds at the extreme end of a chance distribution. Despite this, it can be said that, given the very infrequent (0% of all B-worlds with an infinite space of possible chance outcomes!) occurrence of such worlds, the T-world does successfully deal with the overwhelming majority of F-type worlds.

One final and additional advantage of the T-world proposal in terms of concordance with intuitions is that there is a strong sense in which the appealing principle of the uniformity of nature is retained (at least for all B-worlds that are not part of some initial segment). Namely, that principle that nature does not exhibit discontinuities in its laws, with one set of laws applying to some part of the Humean mosaic, and another set in some other part (for example, from the cited scenario above, by having all coins landing head in one half of the world and tails in the other).

The principle of the uniformity of nature is retained since B-worlds that, when used in isolation as a Humean supervenient base, result in a set of BSA laws that contradict the principle, are no longer counterexamples when considered as a B-world member of a T-world. In such cases, they would not require a BSA law that implied non-uniformity, just a regular chance function with the extreme coincidence of all coins landing heads only in one half of the universe happening to be the B-world in question.

It may be objected that the T-world proposal solves the problem of undermining futures, but at the expense of eradicating any meaningful notion of chance by making all possible worlds deterministic. The reasoning would be that T-worlds are, like the multiverse

proposal of Everettian Quantum Mechanics (Wallace, 2012), just following a deterministic rule – a family of Schrodinger-like equations, all of which describe the branching of superposition-like things. In this picture, it has been argued that chance no longer denotes any physical phenomenon in the world (see, for example, Jansson, 2016).

In Everettian quantum mechanics, chance is typically considered to be an either epistemological or physical phenomenon. The epistemological interpretation is as self-locating uncertainty – that is, the credence an observer assigns to the various outcomes of a quantum event, where each outcome corresponds to an Everettian 'world' – as has been argued by Greaves and Myrvold (2010). Alternatively, it can be said to denote the actual branching structure of space-time forming the divergent superposition of worlds in the Everettian interpretation. Which interpretation one prefers seems arbitrary. A similar case can be made for chance as it occurs in the T-world proposal.

To begin a more constructive response for the specific case of the T-world proposal, it is worth noting that, unlike in the Everettian case, a very strong sense of indeterminism is already ruled out by Humean supervenience of the form advocated here and by Lewis. In its strong sense, indeterminism requires the future be open. But such an open future cannot be understood within the doctrine of Humean supervenience, since the Humean mosaic is intended as a manifold of space-time points where past, present and future all co-exist. Therefore, for any T-world, there is only one future just as there is only one past. This has been argued for quite comprehensively by Emery (2015). If it is assumed therefore that a substantive concept of chance presumes such openness, it does seem incompatible with chance as a supervenient concept.

One could argue that, despite Humean supervenience being made over a space-time manifold, the future can still be interpreted as being open since a many-valued logic could instead be adopted, whereby statements relating to future events are of indeterminate

truth-value. Unfortunately, this represents only an epistemological interpretation of openness, and one which remains essentially time-indexed, neither of which can be made sense of as a metaphysical thesis regarding a 4-dimensional object such as the Humean mosaic.

An alternative, and more precise, reading of indeterminism is as Laplacian indeterminism (see for example, Earman, 1986), where a world is indeterministic when the values of state variables at one point in time in some world w are consistent, under a closed system of laws governing system evolution, with multiple values of state variables at some subsequent time in w. In terms of possible worlds and modal loci, this amounts to the following condition:

> For a modal locus w with laws L, in state w_0 at time t_0, all of the following propositions are consistent with L:
>
> i) w may be in state w_1 at time t_1,
>
> ii) w may be in state w_2 at time t_1, and
>
> iii) $w_1 \neq w_2$.

However, this version of Laplacian indeterminism is framed in the general concept of modal loci. In the two variants corresponding to B-worlds and T-worlds, it can be seen that the T-world proposal is indeed inconsistent with T-world Laplacian indeterminism:

For a T-world w^T with laws L, in state w_0^T at time t_0, all of the following propositions are consistent with L:

i) w^T may be in state w_1^T at t_1,

ii) w^T may be in state w_2^T at t_1, and

iii) $w_1^T \neq w_2^T$.

This is easily shown: any T-world contains all B-worlds as parts that share a common initial segment and the same deterministic laws. If this initial segment is at time t_0 and labelled as w_0^T and at time t_1 it is labelled w_1^T, any other world at t_1, w_2^T, that is assumed non-identical to w_1^T must also share the same common initial segment as w_1^T – namely, w_0^T. But T-worlds can be distinguished from one another by not sharing a common initial segment. So for $w_2^T \neq w_1^T$ to be true, they would need to have distinct initial segments, which contradicts the premise. Therefore, w_2^T must be identical to w_1^T.

But replacing 'T-world' with 'B-world' in the formulation of Laplacian indeterminism results in a form of indeterminism that is consistent with the T-world proposal:

> For a B-world w^B with laws L, in state w_0^B at time t_0, all of the following propositions are consistent with L:
>
> i) w^B may be in state w_1^B at t_1, and
>
> ii) w^B may be in state w_2^B at t_1, and
>
> iii) $w_1^B \neq w_2^B$.

More concisely, from the Humean perspective of the BSA account of lawhood described here, a world is indeterministic iff its best-systematization of laws contains at least one chance function. That's it. It is not the branching structure of modal space, nor the openness of the future under this account that grounds indeterminism of the laws. The former is itself determined by the laws rather than being determined by them; the latter is not a presupposition of the theory.

However, it may still be objected that the B-world formulation of Laplacian indeterminism is not the correct version to use, since the laws themselves are deterministic and describe

the evolution of the T-world, not any individual B-world. This argument appears to take the Schrodinger wave equation of quantum mechanics and the Everettian interpretation of many worlds as its inspiration. It is true that in such an Everettian universe the Schrodinger wave equation is often taken to describe the evolution of the entire T-world. But the T-world proposal developed here is not committed to an Everettian-style interpretation of quantum mechanics any more than it is to another interpretation (apart from the fact that some interpretations may be simpler or stronger than others and so win the competition to be part of the best-system).

The laws of the best-system account are mathematical descriptions of observable phenomena; there would be no need for them to posit a particular interpretation regarding unobservables one way or another. Moreover, because it is assumed that the deterministic laws of a best-systematization must be derivable from a single course of history, the set of laws could never distinguish between an Everettian or a collapse-style interpretation within the laws themselves (although a Bohmian-style hidden variable theory would presumably be effectively excluded given that the role of the hidden variable may turn out better played by a simple chance function such as the Born rule).

Greaves and Myrvold (2010), and subsequently Wallace (2012), for example, have shown that a coherent concept of chance can be constructed for Everettian quantum mechanics. In this case, if Everettian quantum mechanics is true of the actual world, to deny that chance could exist in such worlds is to deny that in the actual world what is interpreted as chancy quantum phenomena are not in fact chancy. It will be argued in Chapter 5 that Everettian quantum mechanics need not be assumed true for the T-world proposal to be valid.

Another objection based on Everettian quantum mechanics is that, if Everettian quantum mechanics is true, it is as much a theory of physics as of philosophy. As such, it is

contingently true. There could be possible worlds that are indeterminate in some – as yet undefinable within the Humean supervenience developed by Lewis – stronger way where the BSA defines chance functions that are not part of some Everettian analogue. If Everettian quantum mechanics is to be contingently true, and other forms of indeterminacy exist, then the T-world proposal developed here would no longer be applicable since its notion of branching and the set W_i^T are consistent only with the Everettian interpretation of indeterministic phenomena.

Everettian quantum mechanics here is understood to comprise both the notion of Our Tree (i.e. the Everettian multiverse), as well as the Schrodinger equation and Born rule. It is true that the latter are contingently true. The argument could then proceed that while the T-world proposal may work as a theory of chance in our Everettian quantum mechanics world, it would not make sense in a non-Everettian quantum mechanics but still indeterministic world.

This argument is not sound because the T-world proposal does not assume the truth of Everettian quantum mechanics; it only assumes that Our Tree is not arbitrary insofar that it has a branching structure, and that the chance functions of the BSA range over all B-worlds comprising Our Tree. If the chance function of the BSA turns out to be the Born rule, this is a contingent truth, but it does not imply that the same chance function would apply to other indeterministic T-worlds, since they would have their own Humean mosaic and perhaps different BSA laws.

Another objection might concede that, while the T-world proposal does solve the problem of undermining, perhaps while retaining a substantive notion of chance and indeterminacy, it is not well-motivated and arbitrary. That is, the Humean mosaic is being redefined simply to solve the problem of undermining and serves no additional role and that the members of the set W_i^T share no metaphysically significant property.

The set of B-worlds that comprise Our Tree is not arbitrarily constructed insofar that all members of the set share a common property. Namely, the property of having a common initial segment.

However, it may be claimed that they are physically unconnected with each other, and that the common initial segments are spatiotemporally unrelated and therefore numerically distinct from one another. This is Lewis's position: that the branching of B-worlds within Our Tree is divergence of worlds in terms of some similarity metric, rather than physical splitting of worlds that at one point are numerically identical and at some subsequent point are numerically distinct (Lewis 1986):

> Two worlds have two duplicate initial segments, not one that they share in common. I, and the world I am part of, have only one future. There are other worlds that diverge from ours. These worlds have initial segments exactly like that of our world up to the present, but the later parts of these worlds differ from the later parts of ours. (p.206)

A response to this argument may proceed as follows: assume that branching is divergence of similarity, rather than physical branching (that is, the situation depicted in the right hand side of Figure 1, rather than the left hand side). Therefore, at every point in time of Our Tree, being composed of many divergent courses of history that happen to share a common initial segment, there is the same number of B-worlds. At time t_0 all B-worlds must therefore be qualitatively indiscernible – in the sense of being exactly alike in all respects – but non-identical. If world w_1 is qualitatively indiscernible to w_2, then w_2 is qualitatively indiscernible to w_1, w_1 is qualitatively indiscernible to w_3, and so on. Further, anything that is numerically identical is qualitatively indiscernible to itself, and so every world is qualitatively indiscernible to itself. Indiscernibility is, in other words, a symmetric, transitive and reflexive relation. Any relation satisfying all these properties is an equivalence relation. Therefore, any theory of chance or modality that assumes worlds

diverge from a common initial segment must be formally equivalent to a theory that assumes that they physically split. Therefore, to say any theory that assumes divergent worlds is arbitrary must be equivalent to saying that any theory that assumes splitting worlds is likewise arbitrary.

In short: whether it is assumed that branching involves diverging worlds with world-bound individuals, or overlapping worlds with transworld identity is of no relevance – the T-world proposal is consistent with both interpretations. Any claim of arbitrariness would therefore be applicable to both interpretations equally. This case has also been argued in case of Everettian quantum mechanics by Wilson (2012) already: which alternative is adopted depends on semantic preference, nothing else.

A further response can be made regarding whether B-worlds with common initial segments are physically connected or not, irrespective of whether such a common initial segment is considered as a single world or as multiple worlds in some kind of superposition (the nomenclature of Everettian quantum mechanics here is intuitively appealing).

An argument that such B-world(s) comprising a common initial segment of a T-world, whether in superposition or not, must be spatiotemporally connected proceeds as follows: consider modal space $\mathcal{M}$, the space of all possible modal loci. Lewis defines an ordering relation over $\mathcal{M}$ – the world similarity relation R. Placek (2010) has shown how this relation works in cases of branching worlds, as well as in divergent worlds. In either case (at least for the Humean), the following two assumptions are made:

i) R is the only ordering relation over $\mathcal{M}$, and

ii) for every modal loci in $\mathcal{M}$, R is defined in terms of the instantiation of properties of the modal locus alone.

Thus, because similarity between loci is measured as proximity in $\mathcal{M}$, any two loci whose instantiation of all properties are identical (i.e. any two qualitatively indiscernible worlds), must be located at the same point in $\mathcal{M}$, and vice-versa.

For such qualitatively indiscernible worlds sharing the same point in $\mathcal{M}$ to deny their spatiotemporal connectedness must be to deny either i) or ii) above since there must be some alternative way of ordering modal loci, or that R is defined in terms other than the instantiation of properties alone (for example, by introducing primitve haeccaeties of modal loci).

An objection regarding the T-world proposal is that it remains possible that a B-world member of Our Tree can have very different relative frequencies from what the chancy laws suggest they should be. For some such world and for some proposition A and time t, $\mathrm{Ch}_{tw}(A) = a$, $a > 0$, but the actual relative frequency of A-type events, r, is very different from a. In fact, a large number of the worlds will be such worlds where $r \neq a$. It can be argued that this means the link between frequencies and chances as required by the Principal Principle no longer holds when the Humean mosaic is assumed to be T-worlds, rather than the actual course of history alone.

The history is depicted in Figure 4.4 below for the case of two simple coin toss worlds where, for the actual course of history, $w^B_@$, $r \approx a$, but for some other world, $w^B_@$ that shares the same common initial segment, $r \neq a$.

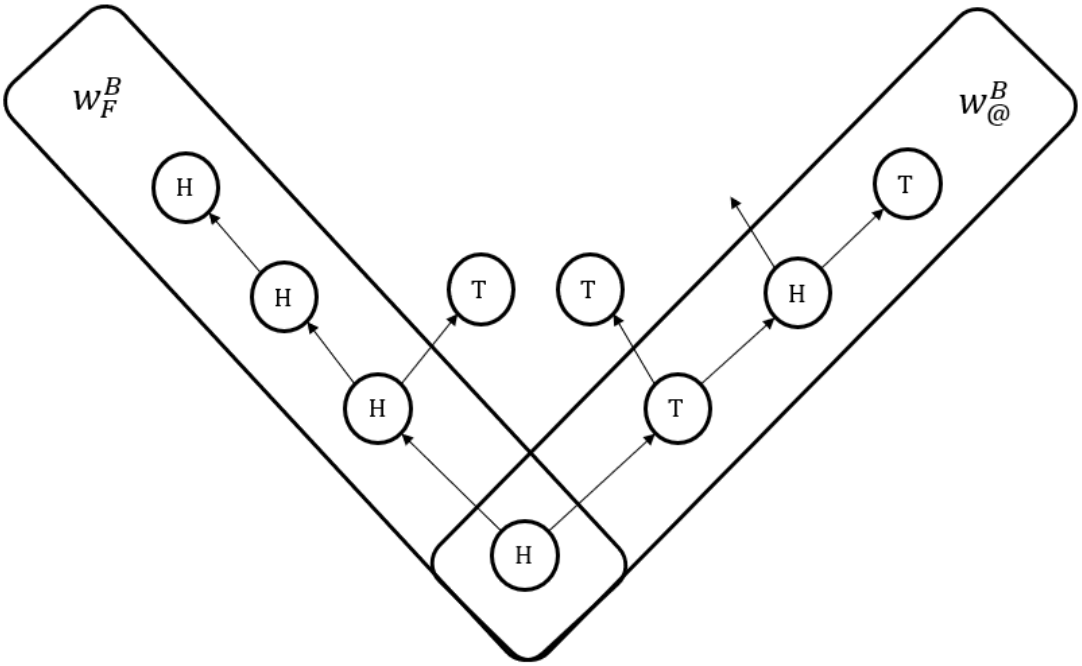


Figure 4.4: depiction of the branching structure of a hypothetical coin-toss T-world with divergent courses of history that do not share the relative frequencies of chance events.

In response to this objection, it is true that for such worlds where $r \neq a$ obtain and may be numerous. In those B-worlds it therefore follows that the PP does not hold. But to expect the PP to hold over every B-world member of Our Tree is unreasonable, for the same reason that it is unreasonable for it to hold for only a small, unrepresentative (in terms of the relative frequency of some estimated chance function), time period of the actual course of history. This is because the T-world proposal requires chance to be defined over a T-world, not a B-world. The PP should therefore be expected to obtain over the entire T-world or some representative periods of one or more member B-worlds. "Representative" here need not imply circularity: the requirement is only that, given a sufficiently large sample, the random selection of B-worlds will tend to result in $r \approx a$, as required by an empirical demonstration of the PP.

Perhaps the most serious objections to the T-world proposal come from a methodological standpoint. One such objection is that the proposal fails the so-called empiricist loyalty test. The empiricist loyalty test states, as introduced in section 1.5.1 above, after Earman (1986), that metaphysical theories should not assume the existence of in-principle empirically undetectable posits. Because all of Our Tree except for our actual course of history is in-principle empirically undetectable, and therefore the entire Humean mosaic on which the BSA theory of laws is based, the T-world proposal fails the test. This is true, but it is likewise true for Lewis's original BSA theory, since it presumed a potentially infinite Humean mosaic that no epistemic agent could ever observe or measure. Moreover, the thesis of modal realism also fails the test since other possible worlds are necessarily beyond the reach of any empirical investigation. It could even be argued that, by its very definition, any metaphysical theory is destined to fail Earman's test, since metaphysical theories are essentially non-empirical.

One's stance on the relevance of the empiricist loyalty test is likely dependent on how staunch an empiricist one is. One important point should be stated however: metaphysics is by definition non-empirical. While metaphysical theories should not contradict any empirical observations, by the very nature of metaphysics there should be no empirical observation that confirms one theory over another, so long as neither contradict the evidence. The empiricist loyalty test appears to require that metaphysics be an empirical endeavor and encroach on existing empirical science. This is not the position adopted here, and the empiricist loyalty test is not considered relevant for the evaluation of metaphysical theories.

A far better test for the worthiness of a metaphysical theory is whether it explains something without contradiction or the introduction of further inexplicable entities. Insofar that the T-world proposal avoids the contradiction implied by the problem of undermining,

and also does not posit any non-Humean ‘whatnots’, it succeeds where Lewis’s original proposal failed.

Finally, it can be argued that the T-world proposal assumes a very strong version of Modal Realism, namely Modal Realism with the Concrete Plenitude thesis: that all possible worlds exist and all exist as entities as concrete as the actual world. This is true. Both Modal Realism and the Concrete Plenitude are both assumed by the T-world proposal. It is a thesis supporting Humean supervenience that denies any primitive notion of necessity or necessary connection. Insofar that without any primitive modal notions, modal realism seems to be a natural and obvious result, it is assumed as a fundamental posit. To deny the T-world proposal on the basis that it assumes modal realism is therefore to adopt a contrary stance to Humean supervenience. It is not the purpose of this essay to argue for Humean supervenience as a metaphysical stance, but to show how it can be made consistent with a theory of chance that avoids the problem of undermining in a non-arbitrary manner.

5. The BSA and quantum mechanics

5.1 Introduction

The logical and metaphysical relationship between the BSA and quantum mechanics is important. Assuming that any metaphysical truth must be true in all metaphysically possible worlds, and assuming that quantum mechanics is true in the actual world, then any principle of Humean supervenience or the BSA that contradicts quantum mechanics implies the falsehood of Humean supervenience or the BSA.

The line of reasoning of course applies to any true (or even possibly true) scientific theory. In the case of quantum mechanics however, it is particularly pertinent because quantum events, at least according to our current knowledge, provide the only instance of metaphysical chance in the actual world (defined as in Chapter 2). Moreover, according to some interpretations of quantum mechanics, the theory of chance is analogous to that of BSA chances outlined in Chapter 4 using the T-world proposal to derive the best systematization of laws for indeterministic worlds. This interpretation, commonly known as the Everettian interpretation or the many worlds interpretation, invokes entities that seem much like the possible worlds of Lewisian Humean supervenience with modal realism and the concrete plenitude assumed in Chapter 4.

Three conclusions will be reached in this chapter. The first conclusion will be that the BSA theory of lawhood thus far developed is consistent with quantum mechanics, although it may not be consistent with all of the various interpretations of quantum mechanics that will be introduced below. Specifically, the theoretical virtues of strength and simplicity at the core of the BSA should prefer the Everettian, or ‘many-worlds’ (Everett, 1957) interpretation of quantum mechanics over the others.

The bulk of the chapter will focus will be on the Everettian interpretation to derive the second conclusion that, in the case of the actual world (that is, Our Tree) at least, the Everettian interpretation may be considered as a specific case of the T-world proposal. Alternatively: the T-world proposal is the generalized (i.e. metaphysical) variant of a specific (i.e. physical) theory of chance. That is, the many worlds of the Everettian interpretation may correspond to Our Tree as defined in Chapter 4.

The final conclusion of this chapter will be that the problem of non-locality in quantum mechanics, often thought to be a serious problem for Lewis's Humean supervenience program, does not contradict Humean supervenience and the BSA theory of lawhood.

This chapter will begin with an introduction to quantum mechanics and some of its more philosophically interesting properties, as well as three common interpretations of the observed or predicted phenomena of the theory. 'Interpretation' here will be intended to cover multiple interpretations that all have some overarching feature in common, but the differences between them will be philosophically uninteresting. Thus 'interpretation' may be better understood as referring here to a family of interpretations.

In drawing the first conclusion of this chapter, each interpretation will be briefly considered in turn and evaluated based on whether the theoretical entities they invoke result in consistency with the BSA theory of lawhood and fundamental chances. In doing so, two background assumptions will be made: i) that quantum mechanics is fundamental to the actual world, and ii) that all of the interpretations of quantum mechanics are – at least in their current form – empirically equivalent.

The first of these assumptions may of course be disproved by future science. But, irrespective of this, insofar that quantum mechanics may constitute the only known

exemplar of fundamental chance, it is worthy of discussion. The second assumption is important since it implies that theory choice must be made on non-empirical grounds.

From these assumptions it will be concluded that the resulting theory of chances from each interpretation is consistent with some interpretation of the BSA theory of chance. However, it will also be concluded that, because each interpretation is not equally parsimonious in terms of its set of axioms or invoked entity types, yet each describe the same phenomena to the same degree of accuracy, some will fair better in the trade-off between strength and simplicity that is at the core of the BSA. Following Deutsch (1999) and Wallace (2012), it will be argued that because the Everettian interpretation of quantum mechanics is the most simple, it would win in the trade-off against both of the competing interpretations.

After concluding that the Everettian interpretation would win out in the best systematization, the notion of a world as invoked by the Everettian interpretation is then examined to assess whether or not it denotes the same kind of entity as that invoked by those modal realists who believe each possible world to be as concrete as the actual world. It will be concluded from this that, while the answer may in fact turn on some possible future advance in physics, the current understanding suggests that Everettian worlds are special cases of the worlds of the modal realist. Insofar that this conclusion holds, it will be argued that the Everettian interpretation should be considered as a special case of BSA chances under the T-world proposal. That is, the Everettian interpretation is the theory of BSA chances under the T-world proposal for the special case of the actual world (i.e. Our Tree).

Finally, in section 5.6, the discussion will move away from chances and focus on a problem that some, most notably Ladyman (2007), have claimed to be fatal to Humean supervenience, and by implication to any theory built up from it, such as a supervenient

theory of lawhood or chances. This is the problem of non-locality. It is a phenomenon predicted by quantum mechanics and observed in the laboratory as well as the real world. Yet it appears to contradict the founding principle of no necessary connections of Humean supervenience.

Following Esfeld (2014), it will be shown that this apparent contradiction is not real, and that in fact the non-locality problem is not a problem for Humean supervenience. This demonstration therefore leaves the way clear for the Humean supervenient theory of lawhood being developed in this thesis.

5.2 Quantum mechanics

In this section the key concepts and principles of quantum mechanics will be introduced with the bare minimum of any formalism or equations. The purpose is not to examine the mathematics describing quantum events or to consider any of the various interpretations of those events. The philosophical issues arising from the theory can be elucidated without any of the mathematics, and the various interpretations of the theory will be considered in the succeeding section.

Quantum mechanics is the best currently available description of sub-atomic particles and their dynamics. At the core of quantum mechanics is the notion of a quantum state and its evolution through time described by the Schrodinger equation. Following convention, these states will be denoted by the 'ket' formalism, such as $|\varphi >$, denoting the quantum state φ. Thus, the Schrodinger equation (SE) links the evolution of a state at time t_1 with that at a subsequent time t_2. Using the ket formalism: $|\varphi t_1 > \rightarrow \ \ |\varphi t_2 >$

The exact form of the Schrodinger equation that describes this evolution is not important for this discussion, except to note that it is an experimentally verified description of the evolution of quantum states. What is important to note is that SE is linear. That is, it

describes a deterministic evolution of quantum states. Or, more simply: it does not contain any probability functions.

Quantum states may be expanded into a set of n vector states, $|b_1>, \ldots, |b_n>$, and associated $a_1, .., a_n$ scalar values representing wave amplitudes. For example, some state $|\varphi>$ may be expanded thus:

1) $|\varphi> = a_1|b_1> + a_2|b_2> \cdots + a_n|b_n>$

Crucially, a quantum state may be a superposition of many vector states ('eigenstates'), so that each $|b_i>$ corresponds to one of many possible measurement outcomes. That is, it corresponds to a multitude of alternative possible ways the world would appear to us when we make a measurement of the quantum state. However, when the quantum state is measured, we can only ever observe one of these states. So, for example, in the famous case of Schrodinger's Cat, the cat – while in the box – evolves to a superposition of alive and dead:

2) $|\text{atom undecayed}> |\text{cat alive}> \rightarrow (\frac{1}{\sqrt{2}}|\text{atom (still)undecayed}> |\text{cat alive}>) + (\frac{1}{\sqrt{2}}|\text{atom decayed}> |\text{cat dead}>)$

When the box is opened however, we effectively make a measurement of the cat's vital status (or, equivalently, the decay state of the atom) and appear to observe only one of these quantum states: $|\text{cat alive}>$ or $|\text{cat dead}>$, the chance of each observation being 50%.

Of crucial importance for the subsequent discussion is that, as far as current scientific understanding goes, this notion of quantum chance satisfies the definition of metaphysical chance introduced in Chapter 3. Namely, where H_i denotes a proposition describing all of the laws of nature and history up to some time i, then the following condition holds for some specific future chance outcome p in w and for all prior times t and v:

3) $\mathrm{Ch}(p|\mathrm{H}_t) = \mathrm{Ch}(p|\mathrm{H}_v)$

In other words, it doesn't matter how much the observer knows about the experimental set-up and all of the history that preceded the measurement, no amount of knowledge will change the chance of p. Likewise, that metaphysical chance does not evolve over the time period preceding the event. In short: metaphysical chance is invariant with respect to time and knowledge. This account leaves out a large amount of detail and conflates some concepts that are distinct, but none of this additional detail is needed for the subsequent discussion.

As noted above, the Schrodinger equation is linear: it describes a deterministic evolution of quantum states from one point to another. Yet, despite this we have a phenomenon that satisfies the conditions for metaphysical chance expressed by (3) above. Explaining how these two apparently contradictory situations can co-exist is where the alternative interpretations of quantum mechanics come into play. Three of these interpretations will be introduced briefly below. In the case of the so-called collapse theories and hidden variable theories, the introduction will be very brief, with the focus being on the Everettian interpretation. This will set the scene for a comparative analysis between chance in the Everettian interpretation and chance according to the BSA and the T-world proposal of Chapter 4.

5.3 Three interpretation of quantum mechanics

Given that the Schrodinger equation is linear, it may reasonably be asked how it comes about that observed quantum states appear chancy: there is 50% chance of observing $|\text{cat alive}>$ and 50% of observing $|\text{cat dead}>$. How the theory explains such apparent chanciness of measurement outcomes is where the different interpretations of quantum mechanics come into play. These three interpretations all share the above essential details in that they accept that the Schrodinger equation describes the evolution of quantum

states from one superposition to another. Where they differ is how they build the apparent chanciness of observing one measurement outcome rather than another onto this essential core.

The three interpretations that will be considered here are i) collapse theories; ii) hidden variable theories, and iii) Everettian quantum mechanics. Although often referred to as different "interpretations" of quantum mechanics, this is strictly speaking incorrect since, as noted by Wallace (2012), the former two variants add to the axioms of quantum mechanics and/or the primitive ontology of the theory, while Everettian quantum mechanics – according to Wallace – adds nothing beyond the Schrodinger equation. Further, the three variants should be seen as broad groupings of different variants rather than individual variants, since there are, for example, many variants on the collapse theory of quantum mechanics. Finally, for the purpose of this discussion, all three variants will be considered empirically indistinguishable. Strictly speaking, this may not be true as the hidden variable of hidden variables interpretations could be empirically measured as something other than an indeterministic phenomenon; the multiplicity of Everettian worlds might also be empirically detected somehow. But, such possibilities seem remote, and while all interpretations predict the same measurements given their current state of development and available measuring devices, it is reasonable to consider them here as empirically equivalent under all experimental setups.

First, collapse theories. This group of theories is characterized by introducing a projection postulate in addition to SE, of the form: $|\varphi> \rightarrow \text{collapse} \rightarrow \ |b_i>$. That is, when a measurement is made on some system in state $|\varphi>$, the wavefunction 'collapses' to some observed state $|b_i>$ from the n possible outcomes defined in (1) above.

The term 'collapse' is defined operationally: it is that process that satisfies the projection postulate. More intuitively, the idea is that SE defines a probability wave over each of the

states on the right-hand side of (1), and at the point of measurement, the wave 'collapses' to one definite state on the right-hand side. So, in the case of Schrodinger's cat, when the box is opened (measurement), the cat turns out to be definitely alive or dead.

Under all versions of quantum mechanics, the probabilities of each outcome are provided by the Born rule (BR), which states that each possible measurement outcome has probability of the square of the amplitude, a, of the corresponding wave function:

4) $\Pr(|\varphi > \ \& \ \text{measurement} \ ||b_i >) = \ |a|^2$

The second interpretation to be discussed here are hidden variable theories. So-called, because theorists who propose this option believe the apparently chance-like behaviour of measurement outcomes is to be explained by some underlying process or entity (the hidden variable), and so what appears to be chance-like is actually deterministic. The most often cited such theory is Bohmian mechanics, which introduces the notion of an additional variable that guides the motion of all other particles. The resulting probability distributions that are observed are then explained as the result of this additional variable.

With hidden variable theories, there is no longer any concept of metaphysical chance. The probabilities of the Born rule would, according to the theory, be some kind of statistical summary or approximate guide to the system behaviour, which would ultimately turn out to be deterministic.

Collapse-theories and hidden-variable theories of quantum mechanics can be grouped together under the term 'single-world theories', since they require only a single world, not a multiplicity of worlds, to be interpretable.

Finally, Everettian quantum mechanics or many-worlds quantum mechanics, denies the need to introduce either the collapse postulate or the Born rule as axioms. In essence, Everettian quantum mechanics denies that the wavefunction ever collapses. Instead,

Everettian quantum mechanics claims that SE obtains at all times, it is simply that the superposition of states depicted in (2) and the Schrodinger cat example really represents distinct worlds. Thus, for each measurement outcome of Schrodinger's cat, alive and dead, there is a corresponding cat-alive world and a cat-dead world; all other macroscopic features of each world remain the same, the only difference being in one world the cat is alive, and in the other dead.

It is claimed by proponents of Everettian quantum mechanics that the reason an observer never sees multiple worlds is that the observer is part of the quantum system, or rather rapidly becomes entangled with it at the time of measurement, and so the observer too appears in each world. Or rather, one version of the observer appears in each world, since personal identity is typically interpreted as a single trajectory or Lewisian time-worm through such worlds in Everettian quantum mechanics (Wallace, 2012). It is also implicit in Everettian quantum mechanics that in fact the term 'measurement' does not really have the same sense as in collapse theories. For, when a 'measurement' is made, rather than the system taking on a particular state as in collapse theories, it takes on multiple states, and thus all measurement outcomes obtain.

There has been much discussion of which among these interpretations of quantum mechanics is the best, either in terms of consistency with other physical theories (such as relativity), theoretical parsimony, or the ease with which the theory grounds some of our intuitive notions (such as chance and indeterminism).

Perhaps the greatest discussion is around the so-called measurement problem (Bell, 1987), as it occurs in each of the above interpretations. In brief, this is the problem as to why measurement should have any effect on a quantum system. More specifically, why the presence of a sentient observer of the measurement should have any effect, whether it be collapsing the wave function, changing the trajectory of some world particle, or splitting

the world into its component quantum states. However, this is not especially pertinent to the purpose of this chapter, which is more concerned with how quantum mechanics fits with the BSA and its associated theory of chance. Thus, the focus of this discussion will be how well these three variant interpretations of quantum mechanics fit within the T-world proposal of Humean supervenience and the BSA developed thus far.

5.4 Interpretations of quantum mechanics and the BSA

Regarding the consistency of collapse theories with the principles of modal realism and the concrete plenitude of worlds, this may appear to be a more obvious area of conflict, since the former invokes only a single world, whereas the latter posits multiple worlds. But there is nothing in the additional axioms of collapse theories to exclude the possibility of additional worlds. Rather, such theories do not commit themselves either way as they are theories of physics that aim to explain only the actual world, not the empirically unobservable worlds of the modal realist.

However, there is a reason to believe that collapse theories would be inconsistent with the BSA as it has been presented thus far. This inconsistency concerns the requirement of simplicity (the requirements of strength and fit are of no relevance here because all three version of quantum mechanics discussed here have equal strength and fit; they are all assumed to be empirically equivalent).

Namely, of the three versions, Everettian quantum mechanics is the simplest (at least in terms of qualitative parsimony). This parsimony is achieved because no additional postulates or hidden variables are added to the theory of quantum mechanics by Everettian quantum mechanics. That is, it is just the Schrodinger equation, and perhaps also the Born rule also depending on whether one believes it needs to be added as an axiom or derived from observations (this remains an open question – see Baker, 2007). This

argument for the parsimony of the Everettian interpretation over single-world theories has been made by Wallace (2012), among others.

It may be argued that Everettian quantum mechanics is not in fact more parsimonious than collapse-theories, since the former posits a multitude of worlds, whereas the latter posits only the actual world. Everettian quantum mechanics therefore goes against Occam's Razor in invoking many more entities. A response to this claim as it relates to quantum mechanics has been given by Deutsch (1999) and Wallace (2012) quite conclusively by making the distinction between parsimony of things, or quantitative parsimony, and parsimony of types of things, or qualitative parsimony. In both cases, it is the objective of science, and explanatory endeavors more generally, to invoke as few types of thing as possible. Alternatively, in the language of the physicist, to explain with as few variables as possible. This is quite obviously in the spirit of the BSA, and is the better interpretation of the 'entities' in the typical 'don't multiply entities without necessity' rendering of Occam's Razor. To this end, Lewis (1986 p.4) also makes a similar distinction in relation to the resulting parsimony of invoking a multitude of possible worlds. While Everettian quantum mechanics does indeed invoke many more entities (i.e. worlds) than single-world theories, they are all the same type of entity. Namely, they are all the same type of world as the world we see around us.

In contrast, the collapse interpretation invokes this same type of entity, albeit once only, but also the additional variables to describe the collapse of the wave function. Parallels to this kind of parsimony-based argument have also been made by Lewis (1986 p.4) in favour of modal realism and his reductive account of modality. In response however, it might be claimed that collapse theories are to be preferred ideologically, in being open to falsification, for example.

Finally, hidden variable theories are difficult to evaluate because by definition they claim that our current understanding of quantum theory is incomplete. It is unclear whether this incompleteness is supposed to be an inevitability or not. Further, if indeed the hidden variable(s) are inevitably beyond empirical investigation or mathematical description, their introduction seems to defy good scientific practice.

These criticisms aside, and assuming there is some additional variable(s) that would turn quantum mechanics into a deterministic theory, it is possible to discuss its subsequent consistency with the BSA as thus far developed. Such consistency turns on the complexity of the mathematical function that is parameterized by the hidden variable(s) and the added complexity that adds to the system, versus the additional predictive power it brings. Because quantum phenomena can be magnified to macroscopic outcomes, as in the case of Schrodinger's cat, being able to predict quantum behaviour deterministically seems like a very great improvement in strength. But equally, given the apparently fundamental chanciness of quantum measurement outcomes, it likewise seems like such prediction would come at a huge cost in terms of complexity. Where the trade-off lies is of course unknown, and so no further meaningful discussion beyond this is possible.

In summary, in the absence of any clue as to the nature of any hidden variables, the only meaningful comparison is between Everettian quantum mechanics and collapse theories. In this regard, according to Occam's Razor, it seems that Everettian quantum mechanics would win in any trade-off between strength and simplicity, since it is (qualitatively) simpler yet equally strong.

Turning now to the comparison with BSA chances and the notion of chance in the different interpretations of quantum mechanics. Clearly this comparison is likewise also only meaningful for the Everettian and collapse theory interpretations, since hidden variable

theories deny that apparently chancy observations are anything but examples of what has earlier been referred to as epistemological chance.

The case of collapse theories seems straightforward, since in fact it says nothing about what chance is or what constitutes the truth maker of chance propositions. It says only that there are chance events and the occurrence of those events is predicted using the probability calculus. Thus, insofar that it says nothing about what makes those chance functions true, it is consistent with any philosophical theory of chance consistent with the same probability calculus. The BSA theory of chance is such a theory, and therefore collapse theories are consistent with it.

This claim of consistency may appear odd, given that collapse interpretations are single-world theories that do not posit the existence of multiple worlds corresponding to different measurement outcomes (see the Everettian interpretation below), whereas the T-world proposal of Chapter 4 does require the assumption of modal realism. The resolution of this apparent inconsistency is straightforward: the T-world proposal and modal realism are justified on independent grounds, such as the ontological parsimony that it permits and the solution to the problem of undermining futures. The T-world proposal allows chance functions to be parameterized without the problem of undermining in the context of the BSA theory of lawhood and chance thus far developed. Single-world interpretations of quantum phenomena in the actual world are consistent with this.

Further, because the BSA theory of chance thus far developed avoids some of the problems with other frequentist accounts of chance, such as the problem of undermining, the reference class problem and the problem of irrational chance values, it could be claimed that the BSA theory of chance is in fact the best interpretation of chance in the collapse theory interpretation of quantum mechanics. To argue for this claim in more detail however, or to review the problems with frequentist theories, is beyond the scope of this

discussion. For the moment it suffices to demonstrate that the BSA theory of chance is consistent with collapse theories.

The most interesting case is that of Everettian quantum mechanics. It should be expected that Everettian quantum mechanics, with its many worlds, should not only be consistent with the BSA theory of chance, but somehow complementary, at least in spirit, to the modal realism that underpins the variant of Humean supervenience being presumed in this thesis and by Lewis.

As with hidden variable theories, and unlike collapse theories, the Everettian interpretation is a deterministic theory. Without collapses of the wave function, there is no concept of chance built into the theory of Everettian quantum mechanics. An analogous claim could be made for the T-world proposal: the branching structure of T-worlds are, like Everettian quantum mechanics, the result of following a deterministic rule – a family of Schrodinger-like equations, all of which describe the branching of superposition-like things. In this picture, chance would seem to be a pure epistemic concept. This is more of a problem for the Everettian theory than hidden variable theories, since the latter denies that anything we observe is chancy or chance-like. In the Everettian interpretation by contrast, there still remains a need to recover some concept of chance to describe the observations of an agent finding themselves in the world with a dead cat rather than a live one, for example.

Deutch (1999) and others have shown that an objective concept of chance as self-locating uncertainty can be constructed in Everettian quantum mechanics. Moreover, Wallace (2012) has gone further, to reify chance in the Everettian interpretation with the claim that it is supervenient upon the quantum state of the system. This reification of chance is effectively claiming that – in the language of B-worlds and T-worlds –the branching structure of an indeterministic T-world is the truth maker of chance propositions. In this

interpretation, the occurrence of quantum measurement outcomes and their relative frequencies is, in effect, our perception of objective chance.

Interestingly, Lewis (2001) himself argued against Everettian quantum mechanics and appeared to doubt whether quantum mechanics was an instance of the kind of metaphysical chance he developed. At the centre of Lewis's position was the so-called quantum suicide argument. This argument is directed against epistemological interpretations of chance in the Everettian interpretation, such as developed by Deutch (1999). It proceeds as follows: If an agent puts himself into a Schrodinger's cat scenario, one outcome involves the agent no longer existing. In such an event it makes no sense, so the argument goes, for the agent to ascribe any likelihood of finding himself in the corresponding world where he would no longer exist. That is to say, the agent cannot possibly find himself in the world where he no longer exists, since that involves a logical contradiction. Thus, the chance of the agent finding himself in the world where he no longer exists must be zero, not 0.5.

This argument seems to be relevant only to self-locating interpretations of chance in the Everettian interpretation however. Interpretations such as those of Wallace, where chance is not an epistemological phenomenon but is instead directly supervenient on the quantum state, are unaffected by Lewis's quantum suicide argument.

Whether either the self-locating uncertainty of Deutsch or the state supervenience of Wallace are actually tenable in the case of Everettian quantum mechanics matters little for the BSA theory of chance thus far developed, for in that case there is a manifest difference between chancy T-worlds and non-chancy T-worlds. Namely, only the former have more than one B-world as components and the branching structure between those worlds defined by the best-systematization (which, as argued in Chapter 4 must be performed on

the Humean mosaic of the T-world, not any one B-world). This is a real difference between chancy T-worlds and non-chancy T-worlds.

True, the 'branching structure' that defines T-worlds implied by the BSA theory of chance developed thus far is supervenient rather than primitive. But the BSA remains a realist theory of lawhood and chances, and its components and structures must likewise be equally real if they are real at all. Thus, there is a real difference between chancy and non-chancy worlds, and it is this difference that constitutes the truth maker of chance propositions. The details of how agents cognize this difference, and whether that alone is sufficient to act as truth maker for any given world is beyond the scope of this discussion.

Another interesting question is whether chance according to the BSA theory of lawhood thus far developed implies indeterminism. Before addressing this, a very strong sense of 'indeterminism' is already ruled out by Humean supervenience of the form advocated here and by Lewis (1986, p.206). In this very strong sense, indeterminism requires that there are branching possible futures at different points in time, and that these possible futures are only realized when branching occurs. But this extreme conception of an open future cannot be understood within the doctrine of Humean supervenience, since the Humean mosaic is intended as a manifold of space-time points where past, present and future all co-exist. Therefore, for any T-world, there is only one future and only one past. This has been argued for quite comprehensively by Emery (2013). If it is assumed therefore that a substantive concept of chance presumes such strong openness, it must be conceded that it is incompatible with the conception of supervenient chance as it has thus far been developed.

However, it could be further argued that, despite Humean supervenience being made over a space-time manifold, the future can still be interpreted as being open since a many-valued logic could instead be adopted for epistemic agents, whereby statements relating to

future events are of indeterminate truth-value. Such a many-valued logic can indeed be constructed and would be meaningful as a metaphysical thesis regarding a 4-dimensional object such as the Humean mosaic. Whether such a 4-dimensional, eternalist, universe is consistent with the collapse interpretation of quantum chance and undetermined truth values regarding setups such as Schrodinger's cat is a broader topic, and beyond the scope of this discussion.

These apparent contrary positions can be made sense of within the context of the BSA theory of chance thus far developed if the problems are instead recast in terms of B-worlds and T-worlds. For, according to this theory, it is true that all T-worlds are deterministic in the strong sense, but that some B-worlds are not deterministic in this strong sense. The crucial distinction is thus whether indeterminism is a property of T-worlds or of B-worlds. It will be demonstrated below that at least a weak sense of indeterminism can be made meaningful for B-worlds, although not for T-worlds.

A more precise reading of indeterminism is as Laplacian indeterminism, where a world is indeterministic when the values of state variables at one point in time in some world w are consistent, under a closed system of laws governing system evolution, with multiple values of state variables at some subsequent time in w (Earman, 1986). In terms of possible worlds, this amounts to the following condition, where 'state' refers to some n-dimensional world that is identical to some mereological part of an n+1 dimensional world, where the additional dimension in the latter world is a time dimension:

5) For a world w with laws L, in state w_0 at time t_0, all of the following propositions are consistent with L: i) w may be in state w_1 at t_1, ii) w may be in state w_2 at t_1, and iii) $w_1 \neq w_2$.

However, this version of Laplacian indeterminism is framed in the general concepts of *world* and *state*. In the two variants corresponding to B-worlds and T-worlds and their

corresponding states (where the state of a T-world is the composite state of all contemporaneous component B-worlds), it can be seen that the T-world proposal is indeed inconsistent with T-world Laplacian indeterminism:

6) For some T-world w^T with laws L and initial segment S, in state w_0^T at time t_0, all of the following propositions are consistent with L where $w_1^B \neq w_2^B$: i) w^T may be in state w_1^T at t_1, and ii) w^T may be in state w_2^T at t_1.

This is easily shown: any T-world with initial segment S contains all B-worlds as parts that share S and the same set of laws. If S is at time t_0 and the resulting T-world labelled as w_0^T and at time t_1 it is labelled w_1^T, any other world at t_1, w_2^T, that is assumed non-identical to w_1^T must also share the same common initial segment as w_1^T – namely, S. But T-worlds can only be distinguished from one another by not sharing a common initial segment. So for $w_2^T \neq w_1^T$ to be true, they would need to have distinct initial segments, which contradicts the premise. Therefore, w_2^T must be identical to w_1^T.

But replacing 'T-world' with 'B-world' in the formulation of Laplacian indeterminism results in a form of indeterminism that is consistent with the T-world proposal:

7) For a B-world w^B with laws L, in state w_0^B at time t_0, all of the following propositions are consistent with L, where $w_1^B \neq w_2^B$: i) w^B will be in state w_1^B at t_1, and ii) w^B will be in state w_2^B at t_1,.

However, it may still be objected that this B-world formulation of Laplacian indeterminism is not the correct version to use, since the laws themselves are supervenient upon the Humean mosaic of the entire T-world, not any one individual B-world, and therefore must describe the evolution of the T-world, not any individual B-world.

This argument appears to take the Schrodinger wave equation of quantum mechanics and the Everettian interpretation of many worlds as its inspiration. It is true that in such an Everettian universe the Schrodinger wave equation is often taken to describe the evolution

of the entire quantum multiverse. If it is granted that the fundamental laws are framed in terms of the T-world, then so be it. This does not mean the fundamental laws imply chance functions at the level of individual B-worlds, much as the Born rule remains valid under Everettian quantum mechanics as a rule by which B-world-confined epistemic agents may live by.

In any case, the T-world proposal developed here is not committed to an Everettian-style interpretation of quantum mechanics any more than it is to another interpretation (apart from the fact that some interpretations may be simpler or stronger than others and so win the competition to be part of the best-system): the laws of the best-system account are mathematical descriptions of observable phenomena; there would be no need for them to posit a particular interpretation regarding unobservables one way or another. Moreover, because it is assumed that the deterministic laws of a best-systematization must be derivable from a single course of history, the set of laws could never distinguish between an Everrettian or a collapse-style interpretation within the laws themselves.

Everettian quantum mechanics and the T-world proposal developed here have at least one thing in common: they both include the notion of world-like parts of a greater sum of such parts. That is, both have notions of a multiverse of worlds. The conventional description of world-like parts in Everettian quantum mechanics have them as splitting from a single common initial segment, rather than as diverging from indiscernible common initial segments, but this is arbitrary and both options are formally equivalent (although, in the case of Everettian quantum mechanics, there may be implications for conservation laws, for example). Both also include the notion of a sum of world-like parts, with the various world-like parts being components of some greater whole. However, Everettian quantum mechanics is not committed to the metaphysical theses of modal realism and concrete plenitude that are presumed by the T-world proposal developed here.

Further, if in the actual world quantum phenomena are the only fundamental phenomena to be indeterministic (in the sense of being best systematized by a chance function), then the set of B-worlds and the internal structure of Our Tree that would be defined by the T-world proposal may turn out to be identical to the set of world-like parts and internal structure of the multiverse defined by Everettian quantum mechanics. However, such apparent unification would be rejected by physicist proponents of the Everettian multiverse, for in such a universe the branching structure between world-like parts corresponds to something tangible. That is, there must be some physical communication between world-like parts. This is the crucial difference between a Lewisian world (i.e. a B-world), and an Everettian world: in the latter there must be (it is assumed) some form of transworld causation; in the former there can be no such thing. As Lewis (1986) says:

> So there isn't any trans-world causation. And not because I so stipulate as a principle of demarcation, but as a consequence of my analyses of causation and of counterfactuals. This is the real reason why there couldn't be a very powerful telescope for viewing other worlds. … Telescopic viewing, like other methods of gathering information, is a causal process: a 'telescope' which produced images that were causally independent of the condition of the thing 'viewed' would be a bogus telescope. No trans-world causation, no transworld telescopes. (1986, p.80).

Deutch (2016) has argued that the existence of such causation can be empirically tested. This remains an open question however. Wallace (2012), for example, has critiqued some observational claims for transworld causation as instead being claims for Everettian theory offering the more elegant explanation – for example, in terms of the power of quantum computing (Wallace, 2012 p.373) being the ability to use the information processing power of other worlds.

However, that there must be transworld causation between Everettian worlds appears no more necessary to the theory than notions of within-world causation. There appears to be

no difference in principle for example between describing the motion of billiard balls in a Humean manner than describing the evolution of the Everettian multiverse. The concept of *causation*, as was argued in section 1.7, is best considered a folk concept that is redundant in science, including quantum physics. In this case, Everettian worlds and B-worlds may turn out to be one and the same thing.

The purpose of this section has been to determine the relationship between the BSA theory of fundamental laws and chances thus far developed, and the various interpretations of quantum mechanics, considered as a possible candidate for such fundamental laws and chances in the actual world. Of the various interpretations of quantum mechanics, the BSA theory of lawhood and chances would seem to prefer the Everettian interpretation over the collapse theory on the basis of simplicity given empirical equivalence, and has little to say about hidden variable theories such as Bohmian mechanics.

Insofar that both the Everettian interpretation and the Humean metaphysics of the BSA theory of lawhood and chances developed thus far are many worlds theories involving worlds that are concrete, and from which a notion of chance results from the branching network of such worlds, they are natural bedfellows. Despite this, Everettian quantum mechanics remains only contingently true, whereas as a metaphysical theory, the BSA theory of lawhood and chances thus far developed, must be true in all possible worlds, if true at all. It is in this sense that Everettian quantum mechanics may be viewed as a possible instantiation of the BSA theory of lawhood and chances thus far developed in the actual world.

Everettian quantum mechanics and the collapse theory are both consistent with the BSA theory of lawhood and chances thus far developed. However, it has been claimed that a particular quantum phenomenon that has been empirically demonstrated to exist is

inconsistent with the Humean metaphysics that underlies the BSA theory of lawhood. This is the problem of non-locality, addressed in the section 5.5 below.

5.5 The problem of non-locality

Non-locality is one of the most peculiar phenomena of quantum mechanics, famously referred to as 'spooky action at a distance' by Einstein (1935). To introduce it, consider two subatomic particles which have a property that can take on exactly one of two values, say red and green. Each particle can be flipped between red and green by some experimental intervention. It is possible to set up these two particles in such a way that whenever one particle is flipped to green, the other particle automatically – and instantaneously – flips to red without any apparent direct intervention on the particle from an experimenter. Thus, the characterization of the second particle in terms of its colour is dependent upon the characterization of the first particle. Moreover, this dependency of the property of the second particle on the property of the first particle involves no obvious connection between the two particles, which may be separated from one another by many light years of empty space. It is as if the second particle magically 'knows' the property of the first particle and changes its own property accordingly. Further, this happens instantaneously.

Non-locality refers to the consequence that there is assumed to be no single, localizable, property had by either particle individually: the property that is clearly instantiated by both particles cannot be localized to either particle alone. It arises where a quantum state is not factorizable into two or more localized states. This happens via a quantum phenomenon called entanglement, which creates a quantum system of two or more particles located at any distance from one another where at least one of the properties of each particle depends on the properties of the other.

Non-locality is often interpreted as inconsistent with Humean supervenience, and that insofar as Humean supervenience contradicts results of empirical science, Humean

supervenience should be rejected (Loewer, 1997; Karakostas, 2009; Butterfield, 2011). Briefly, the basis for this claim is that the phenomenon of entanglement and the resulting non-localizability of particles in the Humean mosaic implies that their properties cannot be locally supervenient because the properties are likewise not localized to specific contiguous points or regions in the Humean mosaic. Of note, Lewis himself (Lewis, 2004) cast doubt on the claim that such non-local properties would turn out to be fundamental, and so obviated the objection.

This inference that Humean supervenience contradicts results of empirical science is not strictly true however. Recall from Chapter 1, the definition of strong supervenience being assumed throughout this thesis:

> **Strong Humean supervenience**: In all possible worlds, laws of nature, dispositions, and chances supervene on the distribution of categorical properties in the world (past, present, and future).

This strong supervenience claim says nothing about the ability to localize fundamental properties to contiguous or point-sized regions of spacetime. This form of supervenience is an instance of global supervenience that does not require locality. That is, it is a supervenience relation between the entire Humean mosaic, not some local part or parts of it, and the laws of nature and other nomological characteristics of a world.

The problem is rather that it appears to contradict Hume's original prohibition on necessary connections in nature, expressed previously in Chapter 1 as follows:

> **No necessary connections (NC)**: There are no necessary connections between fundamental particulars. That is, all particulars are fully specified by their intrinsic fundamental properties – properties that cannot differ between duplicates – and all worlds are built up from such

> particulars according to some principle of recombination consistent with this requirement.

NC is formulated in terms of intrinsic properties. But in the case of non-locality, for any x_i entangled with some x_j, there is some two-place F (that is, the non-localized property) such that Fx_ix_j. But because F is a two-place predicate, at least one of x_j and x_i must have a non-intrinsic property.

The 'principle of recombination' cited in NC, at least that adopted by Lewis and consistent with NC, is also inconsistent with non-local properties. To see this, recall Lewis's (1986) principle of recombination which states, roughly, that every little bit of any possible world can be glued to any other little bit or bits from other worlds to form a new world. In Lewis's (1986) words: "[A]nything can coexist with anything else, at least provided they occupy distinct spatiotemporal positions. Likewise anything can fail to coexist with anything else" (1986, p.88). Entanglement restricts this, since one of a pair of entangled particles cannot be randomly recombined with other particles without having some impact on its entangled pair.

Moreover, the problem of non-locality arises in both the collapse-theory and Everettian quantum mechanics variants of quantum mechanics above. In the case of hidden variable theories, where the hidden variable may be presumed to offer a local solution to non-local phenomena, the result of Bell's Inequalities (Bell 1987) demonstrates that at least some hidden variable theories (namely, the kind presumed by Einstein, Podolsky and Rosen (1935)) cannot offer such a solution without contradicting empirical results.

Thus, it appears that empirical data in the form of non-locality phenomena, and the possibility that quantum mechanics is at least one component of the fundamental laws of the actual world, contradicts two fundamental principles of Humean metaphysics: the

principle of recombination and the prohibition on necessary connections. In the next section it will be demonstrated that in fact there is no contradiction with the BSA theory of lawhood thus far developed.

5.6 The solution to non-locality

In this section, various proposed solutions to the problem of non-locality will be considered, but ultimately rejected, before considering an argument proposed by Esfeld (2014) that retains both the principle of recombination and the prohibition on necessary connections.

It is worth noting first that neither a revised principle of recombination to allow for non-locality, not Hume's original prohibition on necessary connections, conflicts with the foundational principle of empiricism that all metaphysical posits should be justified on empirical grounds alone. That is, the problem of non-locality does not require acceptance of the kind of metaphysical baggage proposed by non-Humeans such as Tooley (1993) and Armstrong (1985), typically in the form of (non-supervenient) laws of nature or primitive modality. At worse, it requires acceptance of some primitive relational properties in the stuff that makes up the actual world.

Nonetheless, one of the foundational assumptions in Lewis's variant of Humean supervenience as presented above, is of the subvenient base of laws of nature (and thereby chances and nomological modality) being a mosaic of permutable, interchangeable, discrete spacetime (or spacetime-like) points or point-sized objects. So long as the subvenient base is thus defined, the problem of non-locality may be thought unavoidable. This is because non-locality requires precisely the non-interchangeability of such points or point-sized objects, or at least that such objects be more than point-sized and bi-located (allowing two entangled particles to be redefined as a single object).

One response to the problem of non-locality that still retains the essence of Humean supervenience is to interpret the entangled two-particle system as a single object in configuration space $\mathcal{H}$. The latter proposal has been made by Loewer (1997). Loewer's proposal appears to be positing the configuration space of quantum mechanics, $\mathcal{H}$, as the fundamental space over which Humean supervenience is defined, rather than spacetime or something like it. Spacetime as we experience it then 'emerges' as the 'shadow' of a 'world particle' as it moves through $\mathcal{H}$. This, in effect, is Bohmian mechanics. As Loewer (1997) concludes:

> The lesson for a defender of [Humean supervenience] to take from quantum mechanics is to count a property as Humean in a world if and only if it is an intrinsic quality of points in the fundamental space of that world. If Bohm's theory (or any other version of quantum mechanics that construes the wave function realistically) is correct, then that space is configuration space. Given this account of Humean properties, quantum nonlocality poses no threat to [Humean supervenience]. (1997, p.105)

Loewer's advocacy of Bohmian mechanics with configuration space as the fundamental space of the world is attractive, particularly since it is consistent with various arguments from both physics (summarized in Seiberg, 2006, for example), as well as in the philosophy of science (Paul, 2012), that Einsteinian spacetime is best considered an emergent phenomenon, arising out of processes at the quantum level.

However, these are arguments from physics and philosophy of science, not metaphysics: for it to be a solution to the metaphysical problem of non-locality, configuration space must be generalizable to all possible worlds. The same can be said of Einsteinian spacetime, or any space-like 'thing' in any possible world. Thus, a more generalized notion of space may be needed. Moreover, this applies to Humean supervenience irrespective of whichever version of quantum mechanics is adopted.

In the specific case of Bohmian mechanics and configuration space, there is also another problem: what results is a very uninteresting Parmenidean unity – a single object, the 'world particle', with no inner structure, and all observations of things and processes in the world being otherwise inexplicable. This makes the theory very unattractive for an empiricist, since the link between observation and theory becomes tenuous, reliant on concepts such as 'emergent spacetime'.

Another solution is to instead define the subvenient base in looser terms that allow for extrinsic fundamental properties and relations, thus allowing for fundamental phenomena such as non-locality. This contradicts NNC but is otherwise compatible with the BSA. For instance, to say that $\mathcal{M}$ is composed entirely of all logically possible permutations of instantiations of each predicate, says nothing about the arity of the predicates or whether the properties they designate must be local or not.

As a demonstration of this, the phenomenon of entanglement defies only local supervenience (in the sense that the qualities of a supervenient object or property cannot be localized to a particular non- disjoint space-time region). But Humean supervenience is global in nature – it requires only that all modal properties and relations be reducible to the global instantiation of non-modal properties. Allow the set of non-modal properties to also include relations and anything else designated by n-ary predicates, and global supervenience is still retained.

Recasting this line of reasoning in an analogous format to Lewis's formulation of a principle of recombination above, a more quantum mechanics-friendly version may be rendered as follows:

> **PR***: For any (non-empty) list of mutually and internally consistent predicates whose conjunction describes a set of n spatiotemporally

related particulars, those n particulars form a maximal spatiotemporally related mereological sum.

In the above, it is assumed that the requirement of being instantiated by a set of particulars thereby makes those particulars into a maximal spatiotemporally related mereological sum, which for the modal realist equates to a possible world. Notice also that the language in which the predicates are formed need not be specified – any language with predicates capable of specifying an instantiated set of objects will do. The only requirements need be logical consistency, both between predicates (e.g. "object x is 2 foot tall" precludes "object x is 20 foot tall"), and within each predicate (e.g. "object x is to the left and not to the left of object y" is precluded"), and restrictions on predicates that imply logical contradictions such as "is red and not red all over".

According to a bundle theory of individuals, the singular terms in PR* above can be identified using some logically equivalent bundle of properties from the list of predicates. Therefore, no primitive notion of 'object' as a spacetime point or anything else needs to be specified. The predicates can be considered as properties or relations, with no restriction needed on arity.

This lack of any restriction on arity appears to be needed if PR* is to be a metaphysical thesis applicable to any possible world. However, Darby (2012) has questioned whether this is appropriate, since if binary relations are allowed, to accommodate for entanglement in the actual world, then why not allow tertiary relations in some other possible world, and so on, until the scenario where some worlds are defined by a single n-ary predicate. This is undesirable not only because it prevents a fully reductive account of metaphysical modality, but also for reasons of parsimony. PR* in fact allows for this such n-ary predicates in the absence of entanglement, since for any n predicates F1, F2, … Fn, the conjunction F1 & F2 & … Fn = G is logically implied. This scenario must be accepted. Yet, it

doesn't matter to the guiding principle of Humean supervenience since laws of nature and other such modal notions would still be supervenient, rather than primitive.

Although incorporating PR* into Humean supervenience in place of PR doesn't imply any need for primitive modality, it does so by reducing modality to logic alone. For, it defines what is (metaphysically) possible in terms of what is logically possible. Moreover, it can still be argued that NC is not preserved for it refers explicitly to intrinsic properties only: "all particulars are fully specified by their intrinsic fundamental properties". This – assuming something like quantum mechanics is a possible fundamental theory in some possible world, even if not the fundamental theory of the actual world – is clearly no longer consistent with PR.

It may be suggested that this is fixed by just dropping the second clause as follows:

> **NC***: There are no necessary connections between fundamental particulars, and all worlds are built up according to some principle of recombination (namely, PR*).

However, it is worth considering briefly what form of necessity is being invoked in NC*. Clearly it is not nomological necessity, for the very purpose of laws of nature is to define such necessary relationships between fundamental objects. Because such relationships are supervenient (at least according to the BSA theory of lawhood) they are not primitive components of the subvenient base of fundamental objects over which metaphysical no-go theorems such as NC* range. Rather, NC* refers to metaphysical necessity which, for the modal realist, refers to some common characteristic of all possible worlds. That is, it is a prohibition on their being any relationship between any objects in the subvenient base that is common to all possible worlds. Alternatively: all non-intrinsic properties of objects are metaphysically contingent. Thus, prohibition of necessary connections between

fundamental objects implicitly requires that all such objects are described exclusively by their intrinsic qualities, and thus NC = NC*.

Proponents of a non-Humean position involving the existence of a second-order necessitation relation, such as Tooley (1993) and Armstrong (1985), could claim that such a second-order necessitation relation is such a characteristic relation common to all possible worlds. But, it seems that many worlds would be Humean in the sense of NC or NC*, since it remains a logically possibility.

To summarize the options discussed thus far to deal with the problem of non-locality: Loewer's proposal to make configuration space the fundamental space of the actual world, apart from seeming like an ad hoc fix to the problem that also has the unfortunate requirement of being tied to a specific version of quantum mechanics, hardly seems to be generalizable to all possible worlds. Further, it requires re-interpreting Einsteinian spacetime as some kind of illusion or shadow of the world particle in configuration space. On the other hand, revising PR to allow for non-intrinsic properties in the principle of recombination serves no purpose, since NC specifically prohibits any non-intrinsic properties.

At this point, it might seem that the whole Humean perspective has to be erroneous. But there is another solution to the problem of non-locality, one that is simpler than Loewer's reification of configuration space and which does not tie itself to a particular interpretation of quantum mechanics nor contradict NC. On the downside, it requires the introduction of a primitive ontology in place of some kind of primitive modality.

The argument relies on the point made above that the necessity implicitly referred to in NC and NC* above is metaphysical necessity, for it is a metaphysical thesis regarding the nature of the Humean mosaic. In contrast, non-locality refers to a phenomenon that is

supervenient upon that Humean mosaic, and the entanglement phenomenon is of a nomological character, not a metaphysical one and thus remains unaffected by any prohibition on metaphysically necessary connections. That is, non-locality is expressed in terms of the BSA laws for a possible world and the properties and particulars that those BSA laws define. Non-local properties are not among the more primitive properties and objects – whatever they may be – of the subvenient base. So long as the subvenient base – the Humean mosaic – can be defined, in every possible world, in local terms Humean supervenience can be left intact. This argument, from a philosophy of physics perspective, has been made by Esfeld (2014).

Consider the example above of two entangled particles, each of which could flip between red and green, but not independently of each other. The resulting non-local property that we may call "Entangled red-greenness" is a supervenient property on the more fundamental properties of the subvenient base properties, such as "Localized redness" and "Localized greenness" and their instantiations. By induction, this demonstration of non-locality between two spacetime-separated objects could be generalized to cases of n spacetime-separated objects sharing a n-ary non-local property. In each case however, there is an empirically equivalent – though clumsier and less efficiently expressed – description in terms of local properties only.

In short, for any two entangled particles a and b in the actual world, the same particles exist in an unentangled state in some other world. This implies that entanglement is not an ontologically primitive property. It may be a fundamental property in the laws of nature of a world such as ours (although, as will be shown in Chapter 5, it will be argued that it would not be since it would not maximize the notion formally introduced later of modal entropy), but according to the BSA theory of lawhood, such fundamental laws are supervenient, and therefore not ontologically primitive.

Implicit in this picture is the presumption that there are such primitive properties and objects upon which non-local phenomena supervene. Moreover, it requires that such primitive properties and objects are themselves Humean in character, lest some similar problem of non-locality recur at this more primitive level. In Lewis's original formulation of Humean supervenience he spoke of points and point-sized objects and their properties to comprise the primitive ontology of the Humean mosaic. However, he did not specify what such primitive properties would be. Esfeld (2020) however has gone further in proposing a more fully developed primitive ontology of the Humean mosaic, claiming that the actual world can be fully described by – and the laws of fundamental physics ultimately translated into – point-sized material objects and their spatiotemporal relation to each other. That is to say, all physical processes can ultimately be described in terms of the spatiotemporal relations between fundamental (presumably indivisible or point-sized) material objects. He makes two claims:

(1) There are distance relations that individuate simple objects, namely matter points.

(2) The matter points are permanent, with the distances between them changing.

Esfeld's proposal seems reasonable when considering particle physics and general relativity. It does however seem to imply that things such as quantum wave functions or fields are not fundamental. That may not be considered important however, since many physicists themselves may consider such things to be instrumental rather than fundamental. There is a more serious problem with Esfeld's proposal however: it appears more to be a reductive account within physics, rather than metaphysics. That is, it is not clear that notions such as 'matter points' and spatiotemporal relatedness are sufficiently generalizable to describe the Humean base of all possible worlds in modal space.

The primitive ontology of both Lewis and Esfeld also seem to imply the need for something unattractive to the Humean – not primitive modality, but primitive properties of those points or point-sized object. Esfeld's proposal – even if workable in all possible worlds – implies the primitive property of spatiotemporal location and being made of matter. Lewis did not specify what such primitive properties would be, but did stipulate their existence, albeit to serve a different purpose than to retain metaphysical locality (see section 2.5 for a discussion of perfectly natural properties).

In the next Chapter, it will be shown how this additional ontological cost incurred by both Lewis and Esfeld can be reduced somewhat. It will be demonstrated that by harnessing some more concepts from physics and the concept of information theory from mathematics, a supervenient alternative to Lewis's PNPs can be defined. Thus, the 'primitive' properties comprising the 'primitive' ontology will be shown to be not primitive at all, but supervenient on the original points and point-sized objects of each world's Humean mosaic. Thus, the primitive ontology of the Humean need comprise nothing more than worlds with points, point-sized objects, the mathematical space in which they may be mapped, and then supervenience gets you everything else.

The purpose of this chapter has been to examine the consistency between the BSA theory of lawhood and Humean supervenience more generally, with our current best theory of fundamental laws – quantum mechanics, and its various interpretations. It was concluded that, while both the collapse theory and Everettian interpretations were consistent with the BSA theory of chance, the BSA criteria for simplicity would appear to prefer the Everettian interpretation over the collapse theory.

On the interpretation of chance and indeterminacy within the BSA theory of chance and Everettian quantum mechanics, both appear to share the same problem of being deterministic theories at the level of the T-world or Everettian multiverse, respectively.

However, it was shown that chance and indeterminisim could still be meaningful to world-bound epistemic agents in both theories. Moreover, it was argued that the branching nature of the T-world would clearly distinguish indeterministic worlds from deterministic ones, and thereby act as the truth maker for chance propositions, if one were needed beyond the epistemic features.

Finally, the problem of non-locality was introduced and discussed as a possible inconsistency between the NC and PR clauses of Humean supervenience. After examining solutions to the apparent inconsistency, it was concluded that a solution along the lines proposed by Esfeld (2014) was to be preferred which accepted non-locality as a supervenient feature of the world, but denied that such non-local relationships need to appear in the subvenient base of a world, thus retaining NC and PR in their original form.

6. Supervenient perfectly natural properties

6.1 Introduction

According to the BSA theory of lawhood and chance thus far developed, the best systematization of the Humean mosaic is defined in terms of simplicity, strength and fit. Each of these three parameters may be defined quite precisely. Strength may be defined in terms of the 'amount' of information contained in the Humean mosaic that is explained by the systematization; simplicity in terms of measures such as the total string length of the systematization measured in bits, for example, and fit may be defined using the statistic proposed in Chapter 3 based on the sum of mean-squared error.

Each of these measures is well understood mathematically, may be objectively and systematically quantified in all possible worlds and readily be used in the process of selecting the best systematization according to the BSA, assuming some exchange rate between them. However, Lewis (1999, p42) has noted a problem first raised by Armstrong regarding the issue of simplicity: the length of the string encoding the systematization must always be relative to some language or set of predicates denoting properties. Change the language, and the length of the string may change.

As the BSA is conceived of as a realist theory of lawhood, such language relativization is problematic: while a law statement may denote the same law in multiple languages, the relativization to language here refers to the selection of the laws themselves, and thus what counts as a law. Such relativization makes the resulting theory of fundamental laws, as in Cohen and Callender's (2009) account, an anti-realist account. On the other hand, for the Humean at least, the alternative of stipulating primitive natural kinds may be more unpalatable, although this was indeed the solution proposed by Lewis (1987, pp.ix-x).

Lewis's solution (1999 pp.8-55) to Armstrong's problem was to introduce a privileged class of fundamental properties. Those properties instantiating this second-order property Lewis called perfectly natural. The term "second order" is used here to denote any property of properties, with no implication regarding the ontological status of such a second order property being intended. Lewis himself does not use the term, and thus it is only used in this thesis in this unburdened sense. Requiring that any systematization of the Humean mosaic be made in a language composed of all and only those predicates denoting perfectly natural properties (plus logical connectives, mathematical functions and so on) would then solve Armstrong's problem by playing the role of the privileged universal language of lawhood.

Lewis did not expect that scientists would necessarily come to know all of the perfectly natural properties that appear in the best systematization of the actual world. But such epistemic access would not be impossible:

> And we may reasonably hope that physics — present-day physics, or anyway some not too-distant improvement thereof — will give us the inventory of all the perfectly natural properties and relations that ever appear in this world. (1999, p225).

Perfect naturalness was not assumed to be supervenient on the Humean mosaic of any particular world – it couldn't because it would be instantiated in multiple possible worlds – nor on the characteristics of modal space itself. To be sure, without any properties, the Humean mosaic of each world would be an undifferentiated mass of bare particularity. But it is important to note that this is not the alternative to there being no perfectly natural properties, since perfect naturalness is a second order property: the particulars comprising the Humean mosaic need not be bare, only that they need not instantiate any perfectly natural properties. This is, in effect, the position of the Humean anti-realist stance towards natural kinds adopted by, for example, van Fraassen (1980). Instead of adopting an anti-

realist position consistent with Humeanism, Lewis introduced perfect naturalness by stipulation. But for perfect naturalness to work in solving Armstrong's problem, it has to be a non-Humean addition to Lewis's overarching Humean metaphysical worldview. It must be a second order property, yet if the BSA is to be a realist theory of lawhood, it cannot be nominal second order property (as in Callender and Cohen, 2009), for it is the property that allows for language-independent comparisons between candidate best systems. But it cannot be supervenient on any one world's Humean mosaic, because it must be instantiated in all isonomic worlds. Moreover, any first order properties causal role seems independent of whether it is perfectly natural or not and thus results in quiddditicism, as argued below.

The introduction of perfectly natural properties (PNPs) into the otherwise Humean metaphysical worldview of Lewis makes the BSA work, at least in terms of providing a common language in each possible world with fundamental laws of nature. But it no longer works as a Humean metaphysical thesis. At least, it doesn't work as a Humean thesis, narrowly construed: that all there is to every possible world is the Humean mosaic and properties that supervene upon it. This is true, it is worth noting, for the realist and anti-realist regarding lawhood. For the realist, it should therefore only be preferred over non-Humean stipulations of primitive modality in the absence of an alternative non-arbitrary Humean solution to Armstrong's problem.

Moreover, accepting PNPs results in quiddditicism – defined here as the claim that at least some properties (of whatever order/arity) have primitive identities independent of their causal role. Quiddditicism in turn results in a conflict with our intuitions that the predicates appearing in any fundamental law of nature should denote properties whose identity is wholly exhausted by their role. This conflict was highlighted by van Fraassen (1980) in a thought experiment, as follows: imagine a team of highly regarded fundamental physicists

announcing the discovery of the fundamental laws of physics. Van Fraassen claims this situation should elicit our belief that the physicists were correct, irrespective of any concerns as to whether they were framed in terms of PNPs. Epistemologically, Van Fraassen's thought experiment seems correct: educated humans would most likely believe them, and without verifiable epistemological access to PNPs and the BSA laws, this would have to be taken on faith.

Lewis's response to van Fraassen's thought experiment was Ramseyan Humility (Lewis, 2009). This thesis can be summarized as acceptance that epistemological access (in the sense that a scientist can know a property to be perfectly natural or not) to PNPs will forever be beyond any scientist. It's not a very satisfactory response. It could be improved a little by noting that if the BSA laws could be so easily confused for human laws of fundamental science, then all the better for the BSA as being a good description of how humans intuit laws should be.

A far better route than humility (Lewis) or resort to anti-realism (van Fraassen) would be to retain the Humean metaphysics of the BSA by not introducing non-Humean PNPs. In this chapter a conjecture will be proposed that does just this. This will be achieved by defining an entropy-like measure over modal space. The entropy-like measure, dubbed hereafter as 'modal entropy', plays the role of identifying some set of predicates in each possible world with fundamental laws of nature that each instantiate perfect naturalness. Because modal entropy is defined over modal space, the perfectly natural properties it identifies requires no additional stipulation or non-Humean entities, being instead supervenient on all of modal space in the same manner as logical or mathematical truth.

Before the conjecture is presented, Armstrong's problem will be more formally defined, and an alternative Humean proposal for its resolution proposed by Wheeler (2014, 2016) introduced and discussed. Wheeler's proposal, although wrong, is nevertheless insightful as

to how the problem may be correctly resolved, and for that reason is included in this chapter.

Following the detailed presentation of the conjecture and a more intuitive discussion of how it solves Armstrong's problem, it will subsequently be shown that it can also be shown to support the solution to the non-locality problem discussed in Chapter 5. It does so by prohibiting non-local properties from being PNPs, when perfect naturalness is defined using modal entropy.

6.2 Armstrong's problem

To see how Armstrong's problem works in the abstract, consider two languages, L1 and L2. The two languages are distinct in that at least one predicate appearing in one does not appear in the other. But the lawlike statements in the two languages are interdefinable, given some translation program. For the moment, count the simplicity of any set of laws framed within either language as the number of characters – including non-redundant white spaces, logical connectives and mathematical operators, etc. Thus, a simple law statement such as "if any x is an F then it is an H", rendered formally as $\forall x(Fx \rightarrow Hx)$ would have length 9, and so on.

Let L1 be composed only of the predicates F, G and H. And L2 be composed only of the predicates M and N. Further, let the best systematization, when expressed in L1 be the statement $\forall x(Fx \rightarrow (Gx \wedge Hx))$, and that the translation program between L1 and L2 establishes that i) $\forall x(Nx \leftrightarrow (Gx \wedge Hx))$, and ii) $\forall x(Mx \leftrightarrow Fx)$. It therefore follows that in L2 $\forall x(Mx \rightarrow Nx)$ expresses the same as $\forall x(Fx \rightarrow (Gx \wedge Hx))$, but with a string length of 9 characters versus 14 when expressed in L1. Thus, laws are language relative when simplicity – measured as string length – is included as one of the theoretical virtues to be optimized.

A real-world example is easy to provide, assuming the standard model of particle physics comprises part of the best systematization of the actual world: “behaves according to the Standard Model” is much shorter than listing out each individual force and particle and describing the interactions between them in precise mathematical detail.

In the most extreme case a self-referential predicate can be defined as “behaves according to the best systematization of laws”. Call this predicate F, then for any part x of the Humean mosaic of world w, Fx. Fx trivially satisfies all of the requirements of simplicity, strength and fit, and would almost certainly end up being the BSA fundamental law in all possible worlds with fundamental laws. Call this the trivialization problem.

It may be objected that the language relativization of laws as shown by Armstrong’s problem or their trivialization as shown by the trivialization problem is missing the point, since even the trivialization predicate F above still denotes the fundamental laws.

The stipulation that the language in which the BSA is formulated not contain trivializing predicates such as Fx can easily be made, and thus the trivialization problem would be solved. For example, it could be stipulated that the predicates not be self-referential: no predicate can be defined in terms of one or more other predicates within some privileged set of predicates. This set of predicates can be defined, for example, by van Fraassen’s hypothetical group of eminent physicists (van Fraassen, 1980). These physicists could hold a congress to determine the most fundamental physical properties from which all others could be derived and declare those to be exclusively used in formulating the best system laws. This is effectively the solution adopted by Cohen and Callender (2009). But such a solution would be relative to the particular language adopted by a particular sentient species (i.e. humans) in one of an infinite number of possible worlds and a particular stage of scientific progress. Thus, realism would be lost.

Armstrong's problem needs a non-arbitrary, non-relativistic solution if the BSA is to be a realist theory of lawhood. Lewis's (1994) solution to the problem was the introduction of perfectly natural properties (PNPs). PNPs were stipulated by Lewis to be all properties that instantiated the second-order property of perfect naturalness. Only PNPs could be used to formulate the best system laws, thus solving – according to Lewis – Armstrong's problem.

But PNPs, as defined by Lewis, instantiate an ontologically primitive second-order property: perfect naturalness. Thus, from an ontological perspective, little progress has been made towards a truly Humean theory of Lawhood: in terms of the ontological balance of account, ontologically primitive perfect naturalness is as costly to the Humean as ontologically primitive laws of nature.

One interesting non-arbitrary solution that has been proposed is the recasting of laws as algorithms that 'compress' the amount of information found in the Humean mosaic – so-called algorithmic compression theory (Wheeler, 2014; 2016). Here, the term 'information' is being used in its intuitive sense, not in its strict mathematical sense. Compression is therefore exactly akin to when we compress a large image file on our computer. Note that such compression requires there to be pre-existing properties, for example of *redness* or *brightness* when compressing an image; what it does not require as a pre-requisite is a privileged set of properties within those pre-existing properties.

In the next section a proposed example of algorithmic compression theory will be considered. Although it will be concluded that the proposal does not in fact solve Armstrong's problem, it will be noted that it does suggest an interesting approach to the problem's resolution, particularly in terms of it using information-theoretic concepts to address the problem.

6.3 Algorithmic compression theory

The algorithmic compression theory of lawhood (ACT) states that laws are information compression algorithms (Wheeler, 2014; 2016). Interpret "information" in this statement in the intuitive everyday sense for the moment; the concept will be made more precise below. What laws do, according to ACT, is to reduce the complexity of the world by showing how a description of it can be compressed in terms of the length of some string needed to reproduce it. This contrasts with the more conventional Humean view of laws as descriptions of patterns observed in nature (taken to be the Humean mosaic in the case of Lewis's Humean metaphysics), or the non-Humean view that they are rules that govern how the world 'works'.

Re-interpreting laws as information compression algorithms should be easier for the Humean than the non-Humean. The complexity of the world can only be reduced if there are recurring patterns in it; and any successful compression algorithm requires the presence of recurring patterns. Humeans and non-Humeans alike can accept this, for it merely states a mathematical fact coupled with an indisputable observation of the actual world. But the Humean can go further and claim the laws of nature are nothing above and beyond such compression algorithms. It could even be argued that an ACT theory of lawhood is likely to turn out to be effectively equivalent to a best systematization account, since both attempt to condense the information contained in the Humean mosaic down to some smaller set of statements (measured in terms of, for example, character length).

ACT can be fleshed out more formally using Kolmogorov complexity, K, (also called algorithmic complexity) as the measure of complexity used. K measures the combined size of the smallest input string y of a universal Turing machine (UTM) and the smallest program string running on a UTM p that together generate some output string x. Using bits

represented by 0s and 1s to measure the size of strings, each taken from the maximal set $\{0,1\}^*$. $K(x)$ is then some function such that $\{0,1\}^* \to \mathbb{N}$ and 1 below holds:

1. $K(x) := \min_{y,p:p(y)=x} l(p) + l(y)$

Where $l(i)$ is the length of some string i. It is important to note that, in this formulation, the minimization is of not just of the length of the input string y, but the sum of y and length of the UTM program p.

1 above is a generalized description of compression of Kolmogorov complexity. Replacing some terms with philosophical concepts will demonstrate its relevance to this discussion. It can be seen more clearly how the terms l, p and y in 1 above can be used to define ACT as a philosophical theory of lawhood. Namely, if the situation is considered when the output string x is a binary description, S_w, of the Humean mosaic of a world w, y is a corresponding set of laws qua compression algorithm L_w, and p remains a program running on a UTM that transforms the latter into the former, then the result summarizes the founding principle of ACT:

2. $K(S_w) := \min_{p,L_w:p(L_w)=S_w} l(L_w) + l(p)$

In other words: the laws of w, L_w, are those propositions that compress the (Komogorov) complexity of the Humean mosaic, S_w, when combined with some optimum UTM program p.

Before considering ACT further in terms of whether it can solve Armstrong's problem, it is worth noting immediately that ACT in its formulation above has problems with indeterministic worlds. This is because if the laws L_w contain some probability function, the output S_w will rarely be identical to the actual Humean mosaic of such a world. At this point however, two solutions already proposed in earlier chapters can come to the rescue.

First, from chapter 3, the difference between the two in terms of the sequences of chance events would not result in zero-fit, since fit is better measured in terms of X^2. Second, for differences in the relative frequencies of chance events, the T-world proposal discussed in Chapter 4 prevents such differences resulting in undermining. For, if it becomes clear during the step-wise systematization process that the Humean mosaic is not of a singleton T-world (i.e. that the best systematization includes probability functions), then S_w comes to denote the Humean mosaic not of a single B-world, but the entire T-world of which it is a component.

Let superscripts denote the language from which the string has been encoded into binary, thus: $K(S_w^i)$, L_w^i, and p^i. Armstrong's problem can then be simply stated as the claim that, for some i, j, and w, the following inequality holds:

3. $K(S_w^i) \neq K(S_w^j)$

While accepting this inequality, Wheeler (2014, 2016) maintains that any differences are negligible for ACT. His reasoning is based on Kolmogorov and Solomonoff's invariance theorem (Solomonoff, 1964), which states that the absolute difference in Kolmogorov complexity between any string encoded from different languages does not exceed some value c. In the case of strings that encode the Humean mosaic of some world w:

4. $|K(S_w^i) - K(S_w^j)| \leq c$

Wheeler further assumes, quite reasonably, that the output string of a UTM written in one language, S_w^i, is equivalent to the same information encoded from some other language, S_w^j, plus some additional translation string $p^{i\to j}$. He reasons therefore that the residual complexity c in 4 above is the string length of $p^{i\to j}$. Therefore:

5. $K(S_w^j) := \min_{p^i, L_w^i, p^{i\to j} : p^i(L_w^i, p^{i\to j}) = S_w^j} l(L_w^i) + l(p^i) + l(p^{i\to j})$

Critical to Wheeler's conclusion that ACT solves Armstrong's problem is that $c = l(p^{i \to j})$ is negligible in comparison to $l\left(L_w^i\right) + l\left(p^i\right)$.

But even if 'negligible' values were acceptable for a metaphysical thesis, the low complexity assumed of the translation program does not follow from the invariance theorem, or from the nature of $p^{i \to j}$ itself. To demonstrate this, consider the program length required for a translation program translating from a language i with a vocabulary length l^i to another language j with a vocabulary length $l^j = k^{l^i}$, for some $k \in \mathbb{N} > 0$. That is, where there is no straightforward linear translation between the two languages. The corresponding value $l(p^{i \to j})$ could not reasonably be considered negligible relative to $l\left(L_w^i\right) + l\left(p^i\right)$. This is the situation were i to be mapped onto literals that are the disjunctive normal form of j. For example, consider a language i with $2n$ predicates: $A_1, A_2, .., A_n, B_1, B, .., B_n$ and a systematization of the Humean mosaic that (as it happens) is in conjunctive normal form comprising n clauses, as follows:

$(A_1 \vee B_1) \wedge$
$(A_2 \vee B_2) \wedge$
…
$(A_n \vee B_n)$

Where $l = kn$, for some $k \in \mathbb{N} > 0$. But if these clauses were to be translated into language j, comprising 2^n predicates $C_1, C_2, .., C_n, \ldots, C_{2^n}$ there would need to be 2^n clauses. This is the case if j were the disjunctive normal form of i. Namely:

$(A_1 \wedge A_2 \wedge \ldots \wedge A_n) \vee$
$(B_1 \wedge A_2 \wedge \ldots \wedge A_n) \vee$
$(A_1 \wedge B_2 \wedge \ldots \wedge A_n) \vee$
$(B_1 \wedge B_2 \wedge \ldots \wedge A_n) \vee$
…
$(B_1 \wedge B_2 \wedge \ldots \wedge B)$

In other words: the translation into another language can involve an exponential amount of effort, and thus unlikely to be neglibigle.

There is a much greater problem with Wheeler's proposal however, which is that it guarantees no unique solution, irrespective of whether or not the translation program length is negligible. This is because Kolmogorov complexity, as with Shannon's information theory, attaches no semantic values – from within the measures themselves – to the binary digits (or whatever encoding alphabet is used) composing the strings or messages. That is, in the case of Kolmogorov complexity, the string $L_w^i = 0101010$ means nothing to the UTM; it is just a string of digits. Likewise for Shannon information theory or any other similar measure: the symbols are always uninterpreted. '0101010' could denote, within language i the Standard Model of quantum physics, or some other theory. Crucially, it could also denote a trivialization predicate of the type discussed in section 6.2 above.

Assume that it does denote a trivialization predicate. The founding principle of ACT given above does not exclude this as a best-systematization of the Humean mosaic. This is because there are two components to the right-hand-side of the founding principle: the string encoding the laws, $l(L_w)$, and the string encoding the UTM program that generates the string encoding the Humean mosaic using the laws as input, $l(p)$. It is reasonable to assume that, given a constant value of $K(S_w)$, that as $l(L_w)$ gets shorter, so $l(p)$ must counterbalance this by getting longer. That is, the more trivial or uninformative the laws get, the more complex the program has to become to interpret the laws in order to generate S_w. There is no reason, particularly given the redundancies of natural languages, that $l(L_w) + l(p)$ will be constant for some given $K(S_w)$. But likewise it should not be assumed that one can vary without the other varying also.

That the possible values of $l(L_w)$ and $l(p)$ are not independent given some $K(S_w)$ follows from the observation that the trivialization of $l(L_w)$ implies a larger computational burden

to 'unpack' all of the information contained within it. To demonstrate, consider UTM to be a super intelligent AI program. It is told that the best system of laws is whatever comes out of a best systematization process, such as the step-wise process depicted in Chapter 4. What would the shortest program, p, look like that is asked to output $K(S_w)$? It could not be L_w because this provides little information for p to use. Therefore the information necessary to output $K(S_w)$ has to be 'hard-wired' into p itself. That is, the shortest possible p would have to be a program that contains within it a description of the laws of nature in a manner that is useful – the information that would otherwise have been in L_w.

Although failing to address Armstrong's problem by not taking account the trade-off between the program length needed to decode input string length, and input string length, Wheeler's proposal is nonetheless informative regarding an alternative approach; one that is as equally applicable to the ACT interpretation of lawhood as the more conventional laws-as-generalizations interpretation. This proposal is based on the notion of Shannon information theory (Shannon, 1948), rather than Kolmogorov complexity theory, and is presented in the next section.

6.4 The modal entropy conjecture

The modal entropy conjecture is a proposed solution to Armstrong's problem that identifies a unique reference language for each set of isonomic possible worlds (i.e. each set of possible worlds all members of which are nomologically possible relative to each other). This unique reference language contains predicates that denote the perfectly natural properties of each possible world, just as in Lewis's original proposal. Crucially however, the PNPs according to the modal entropy conjecture will be defined non-arbitrarily as supervenient on the modal space of all possible worlds (as opposed to the Humean mosaic of any particular T-world or B-world), and thus need not be introduced as theoretical primitives.

Thus, in terms of the ontological balance of account, there is no need to introduce any ontologically primitive second-order properties beyond the particulars and their intrinsic properties that compose the Humean mosaic, and Humeanism is thus retained.

The modal entropy conjecture is based on the notion of Shannon information (Shannon C, 1948). In its conventional form, the Shannon information of some message x, $S(x)$, is given by

6. $S(x) = -\log P(x)$

Where $P(x)$ is the probability of message x being selected from all possible messages that could have been selected and the logarithm is base 2 (by convention). 'x' need not refer to a message in the conventional sense, but any probabilistic outcome value. For example, the toss of a fair coin, in which case the information 'contained' in the coin landing heads, $S(h) = -\log 0.5 = 1$. What the actual meaning of this value is remains a matter of some debate. Importantly though, it is worth emphasizing two points. First, Shannon information is not a semantic theory – it says nothing about what the 'x' denotes. Second, the values of Shannon information (measured in bits when the logarithm base is 2) do not correspond to the everyday meaning of 'information'. On the contrary in fact, for it seems natural to think of something that happens with certainty to convey more information than one that may not happen. But according to Shannon information theory, the former is associated with 0 bits of information, the latter with some non-zero positive amount.

To see the relevance of Shannon information theory to Armstrong's problem and language relativity, consider a language L that contains predicates denoting PNPs. If PNPs are to play any meaningful role in any theory of lawhood, whether Humean or not, it is reasonable to expect that a message can be constructed using only terms from L that uniquely identifies any possible world in modal space. Moreover, every possible world in modal space corresponds to some message in L. More explicitly: each possible world can be uniquely

described by some message in L. This, in effect, is the basis of ersatz theories of possible worlds.

Given this denotation relationship between messages in L and the points in $\mathcal{M}$, then just as it is possible to define the measure $S(x)$ for the possible messages, it would be possible to define the analogous measure for possible worlds. In brief, the modal entropy conjecture that will be fleshed out in more detail below states that this analogy holds, and that PNPs can be maximized from the set of all possible predicates as those which maximize the value of a measure related to $S(x)$ by this analogy.

Consider modal space $\mathcal{M}$. Every point in $\mathcal{M}$ is some possible world. Partitions can be made on $\mathcal{M}$ that splits it into the nomologically possible and the nomologically impossible sets of possible worlds, relative to some set of fundamental laws of nature. Assume for the moment some partition on $\mathcal{M}$ between nomologically possible and nomologically impossible worlds, with an equal number of worlds in each partitioned subspace. Then, we can calculate the Shannon information associated with a possible world being nomologically possible. As before, $S(w) = -\log 0.5 = 1$.

An interpretation of this application of Shannon information to $\mathcal{M}$ should now be evident: the number of bits represents the number of binary partitions on $\mathcal{M}$ needed to identify the set of possible worlds of interest.

This can equally be expressed in terms of property instantiation: Shannon information defined over $\mathcal{M}$, for any predicate f, is a measure of the number of partitions on $\mathcal{M}$ necessary to uniquely identify all possible worlds that instantiate f-ness (that is, all worlds where there is at least one point or point-sized object that is an f). When f denotes some trivialization predicate, then $S(w) = -\log 0.5 = 1$ as above. However, in other cases the

situation is more difficult because some f may be instantiated in both nomologically possible and impossible worlds.

To generalize this approach, suppose that some candidate best system has partitioned $\mathcal{M}$ into the subspace of nomologically impossible worlds, $W^{\neg L} = \{w_1^{\neg L}, w_2^{\neg L}, \dots, w_k^{\neg L}, \dots, w_\infty^{\neg L}\}$, and the subspace of nomologically possible worlds, $W^{L} = \{w_1^{L}, w_2^{L}, \dots, w_k^{L}, \dots, w_\infty^{L}\}$. It can now be asked, for any predicate f that appears in the candidate best system, what is the probability of some possible world in $\mathcal{M}$ that instantiates f being a nomologically possible world? If the size of W^L and $W^{\neg L}$ are both assumed to be equivalent, then for some predicate f_j appearing in the fundamental laws of nature, the relative frequency of nomologically possible worlds among all worlds that instantiate f_j, $r(f_j)$ is given by:

7. $$r(f_j) = \left[\lim_{k\to\infty}\left(\frac{|f_j(w_k^{\neg L})|}{k}\right) + \lim_{k\to\infty}\left(\frac{|f_j(w_k^{L})|}{k}\right)\right]^{-1}$$

Where $f_j(w_k^{\neg L})$ denotes all nomologically impossible worlds that instantiate f_j in some k-sized set of nomologically impossible worlds, and $f_j(w_k^{L})$ all nomologically possible worlds in some k-sized set of nomologically possible worlds, all of which will instantiate f_j by definition. Thus, as k tends towards ∞, $0.5 \geq r(f_j) \geq 1$. More specifically, $r(f_j) = 0.5$ when f_j is ubiquitous in both nomologically possible and impossible worlds, and $r(f_j) = 1$ when f_j is ubiquitous in nomologically possible worlds and uninstantiated in all nomologically impossible worlds.

Note that the terms “nomologically impossible” and “nomologically possible” in this context refers to where the partition lays as specified by the candidate best system.

It is reasonable to assume that any predicate f_j which appears as a term in the fundamental laws of nature of a world must be instantiated in all nomologically possible

worlds relative to that world and its fundamental laws. This assumption is not the result of a logical inference: a law of the form $\forall x(f_j x \rightarrow Gx)$ does not imply $\exists x(f_j x)$. But a world where the fundamental predicates were not instantiated would be strange indeed. Consider, for example, a world with the fundamental law $E = Mc^2$, but with no matter in it! Making this assumption means that $\left|f_j\left(w_k^L\right)\right| = k$. Therefore 7) simplifies to:

8. $r(f_j) = \left[\lim_{k\to\infty}\left(\frac{|f_j(w_k^{\neg L})|}{k}\right) + 1\right]^{-1}$

Which implies that $0.5 \leq r(f_j) \leq 1$.

Now, call the modal Shannon information of a predicate f_j, $S_{\mathrm{mod}}(f_j)$, what is obtained by using the relative frequency $r(f_j)$ in place of $\mathrm{P}(x)$ in the equation of Shannon information above:

9. $S_{\mathrm{mod}}(f_j) = -\log r(f_j)$

Again, the measure $S_{\mathrm{mod}}(f_j)$ is a measure of the difficulty in partitioning $\mathcal{M}$ into those worlds that instantiate f_j and those that do not, given that the existing partition into the nomologically possible and nomologically impossible goes some way to achieving this.

The relevance of modal Shannon information is this: the trivialization predicate F discussed above has $S_{\mathrm{mod}}(F) = 0$; mathematical and logical predicates are all such that $S_{\mathrm{mod}}(-) = 1$. All other predicates lie in the range $0 < S_{\mathrm{mod}}(-) < 1$. This measure defined over predicates lies at the heart of the modal entropy conjecture and will be fleshed out in detail below.

As a first step towards fleshing out the conjecture, notice that the notion of Shannon information, and by extension that of modal Shannon information, can be extended to entire sequences of messages or predicates. In the case of conventional Shannon

information, the extension is to the concept of Shannon communication entropy H of a sequence of messages, $X = x_i(i = 1, \dots, n)$, and is calculated as the weighted mean of the Shannon information of each x_i:

10. $H(X) = -\sum_{i=1,\dots,n} P(x_i) \log P(x_i)$

For a set of predicates, $F_i = f_j(i = 1, \dots, n; j = 1, \dots, m)$, the analogous extension from modal Shannon information to Shannon modal entropy is straightforward:

11. $H_{\text{mod}}(F_i) = -\sum_{j=1,\dots,m} r(f_j) \log r(f_j)$

Let each set $i = 1, \dots, n$ of predicates be called a language. The modal entropy conjecture is then: the perfectly natural properties are those properties denoted by the predicates in the language, F^L, that is used to express the best systematization of a world (as defined by the BSA account of lawhood thus far developed) such that:

12. $F^L := \max_{i=1,\dots,n} H_{\text{mod}}(F_i)$

Put more simply: the PNPs are denoted by that set of predicates used in a candidate best system that maximizes Shannon modal entropy. Or alternatively: the best systematization of the Humean mosaic is that which is expressible using predicates that are most ubiquitously instantiated in both W^L and $W^{\neg L}$.

It may be objected that the predicates that are most ubiquitous are the logical and mathematical predicates. True enough. But that's the point: because they are ubiquitous, they cannot be used alone to encode contingent patterns in the Humean mosaic. There must be *some* empirically defined predicates in the language F^L. The modal entropy conjecture says that this empirical component should be minimized.

The issue of where this minimization fits in with the maximization of strength and simplicity in the selection process of the best system will be discussed in the next section. First, it will

be demonstrated intuitively why the maximization of modal entropy is an attractive and non-arbitrary requirement for any language used to encode the best system of a world. This is best demonstrated by returning to algorithmic compression theory and Kolmogorov complexity.

Consider again the foundational principle of ACT:

2. $K(S_w) := \min_{p,L_w:p(L_w)=S_w} l(L_w) + l(p)$

It was argued above that decreasing the length of one term on the right-hand-side of 2 must lead to a concomitant increase in the other term on the right-hand-side if S_w is to be fully expressed. That is, there is a trade-off between the amount of empirical information encoded in the laws and that encoded in the program that translates the laws into the Humean mosaic. 'empirical information' in this context refers to contingent information that identifies a particular world or group of isonomic worlds from other worlds; 'information' is being used in the colloquial sense. Analogously, it is a consequence of the modal entropy conjecture that as $l(L_w)$ is maximized so $l(p)$ will be minimized and vice-versa. This is because if all of the predicates used in L_w contain the maximal amount of Shannon modal entropy, then any examination of the meaning of the predicates used in L_w provides the least amount of information (in the colloquial sense) of whether the predicates are most frequently instantiated in W^L worlds or $W^{\neg L}$ worlds. That is, they contain the least amount of empirical information as in the colloquial sense described above.

What does this 'least amount of empirical information' imply about the predicates in L_w? Mathematically, the least amount of empirical information implies that all of the predicates in L_w would be logical and mathematical properties that are ubiquitous. But this would not be sufficient to define any best systematization, since it would contain no information

particular to any possible world. Thus, there must be some empirical information, just the bare minimum to express the best system and thereby partition $\mathcal{M}$ into W^L worlds and $W^{\neg L}$ worlds.

At this point two objections may be made. The first is a circularity objection: the best system partitions $\mathcal{M}$ into W^L worlds and $W^{\neg L}$ worlds, yet this presumes identification of the PNPs using the measure of modal entropy, which itself presumes a partition of $\mathcal{M}$ into W^L worlds and $W^{\neg L}$ worlds. This objection is valid but can be overcome through the incorporation of the modal entropy conjecture into the selection process of the candidate best system, which will be detailed in the following section.

Second, it may be objected that incorporating the maximization of Shannon modal entropy into the optimization parameters of the best system requires adoption of the ACT theory of lawhood. This objection is not valid; the modal entropy conjecture establishes the language in which the best system is to be expressed, not the nature of what a law comprising the best system should be. The purpose of the discussion of ACT above was to describe why the modal conjecture works in identifying PNPs, it was not an argument for ACT itself.

6.5 Modal entropy and the BSA

The preceding discussion introduces the notion of Shannon modal entropy and uses the notion to define PNPs as those properties denoted by the set of predicates used in a candidate best system that maximizes Shannon modal entropy. Consequently, that the best systematization of the Humean mosaic is that which is expressed using predicates that are most ubiquitously instantiated in both W^L and $W^{\neg L}$. The question remains of how Shannon modal entropy and the maximization criteria should be incorporated into the best systematization selection process.

What is clear is that a best systematization of the Humean mosaic cannot presume identification of PNPs using Shannon modal entropy, for the latter measure depends on the former systematization identifying the partition between W^L and $W^{\neg L}$ worlds. Thus, the following circularity objection could be made: the modal entropy conjecture establishes the language in which the best system is to be expressed. And the best system establishes the fundamental laws of nature, which in turn distinguishes between W^L worlds and $W^{\neg L}$ worlds. But the modal entropy conjecture requires that the distinction between W^L worlds and $W^{\neg L}$ worlds has already been made.

Similarly, derivation of the set of PNPs using Shannon modal entropy cannot presume identification of the set of fundamental laws of nature using the BSA, since the latter requires the former set of PNPs to have been determined.

The following method avoids this circular reasoning by first identifying all possible systematizations of the Humean mosaic – each one relative to some logically possible set of predicates; each one corresponding to some partition over $\mathcal{M}$ (which may be trivial – it merely identifies some empty subspace of worlds), and then comparing the resulting values for modal entropy, strength and simplicity. The procedure is as follows:

Consider the set of all n possible sets of instantiated predicates in a world w, $\Pi = \{F_1, F_2, \ldots, F_i \ldots, F_n\}$. This set can be defined as the power set of the set of all instantiated predicates, closed under Boolean algebra and excluding the null set. It follows that each $F_i \in \Pi$ represents some possible language for which a system of candidate laws could be formulated, when allowing for the likely possibility that most such systems will partition some empty subspace of $\mathcal{M}$, as mentioned above.

The following statements immediately follow from this setup, where F^L denotes that set of instantiated predicates used in formulating a best systematization:

- For some $F_i \in \Pi$, $F_i = F^L$, and thus maximizes the Shannon modal entropy H_{mod} and forms the language in which the best-systematization is to be written.
- For every $F_i \neq F^L$, it will be the case that the following inequality holds: $H_{\mathrm{mod}}(F_i) \leq H_{\mathrm{mod}}(F^L)$, where '$H_{\mathrm{mod}}(F_i)$' denotes the value of Shannon modal entropy measured for the set of predicates comprising language F_i.
- For every $F_i \in \Pi$, there exists some Lewisian (in the sense of not including Shannon modal entropy as a parameter in selecting the best-systematization) best-systematization of w, S_i, according to the parameters of simplicity, strength and fit, and therefore for every F_i there will correspond some subspace of $\mathcal{M}$, M_i, corresponding to the modal notions of nomological necessity and possibility.

These considerations indicate that every possible systematization of a world w corresponds to a 3-tuple $< F_i, S_i, M_i >$. Moreover, neither the language, F_i that is used in the best-systematization relative to that language, S_i, nor S_i itself, nor the resulting partition M_i requires any measure of Shannon modal entropy to have been made a priori. However, for every $< F_i, S_i, M_i >$ there is some value of H_{mod}. Therefore when considering all m possible $< F_i, S_i, M_i >$ systematizations, where $m \geq n$ because of possible non-unique best-systematizations for each F_i, they collectively represent a search space over which the maximum value of H_{mod} can be identified. That possible best-systematization that results in the maximum value of H_{mod} is then the best-systematization coded using the language F^L, and so the predicates comprising F^L denote the perfectly natural properties of w, and the best-systematization S_i denote the fundamental laws of w and the partition M_i which distinguishes between W^L worlds and $W^{\neg L}$ worlds

Because the above account does not presume any particular partition on $\mathcal{M}$ in the process of best-systematization, it is non-circular. However, it does introduce the additional variable of modal entropy into the systematization process. Most importantly though – it gets the job done, without recourse to either primitive perfectly natural properties nor circular reasoning.

One issue that may be raised at this point is to claim that in fact there may be multiple, non-identical, best-systematizations resulting from the above process since two or more systematizations may be Lewisian (in the sense above) best. This non-uniqueness problem is frequently noted as an objection to Lewis's original formulation. Further, because for every $F_i \in \Pi$ there is one unique value of H_{mod}, the introduction of Shannon modal entropy as an additional optimization parameter in selecting the best system does not solve this problem. This is unfortunate, but so be it: there seems no a priori requirement on possible worlds that they should have a unique set of fundamental laws. The minimal requirement that seems reasonable is that whatever sets of laws prevail, that they be consistent in terms of their predictions in the overwhelming number of instances (given that different trade-offs between strength and simplicity may be made).

6.6 Supporting the conjecture

The term 'conjecture' is intended to denote a theory whose truth has not been proven but which seems reasonable to assume on intuitive grounds. The term is reasonably applied to the claim that Shannon modal entropy can be deployed as an additional optimization parameter in a best-system account of lawhood to identify an objective set of supervenient PNPs. This is because the preceding introduction of this claim presumes that, if it is granted that fundamental laws are those identified by a best systematization of the Humean mosaic, the predicates that appear in those laws should be as simple as possible. This presumption is not proven in the preceding discussion, and the Shannon modal entropy conjecture therefore remains a conjecture.

Nonetheless, some more support can be given to the requirement that fundamental laws should be framed using only the simplest predicates. Clearly this support cannot come from any regular, linguistic, sense of simplicity, since this notion is immediately dependent on the language used to frame the laws. However, another line of reasoning is to define

simplicity in terms of presuming the least amount of empirical information about the rest of the Humean mosaic; anything in addition would mean that there was some more elementary predicate. This inference can be fleshed out with an example in the case of a world, where it is assumed that the subatomic particles – electrons, quarks, etc. - are the sole elementary building blocks of the Humean mosaic. Assume that the set of predicates "is a quark", and "is an electron", and so on are thus defined according to the above maxim as the simplest in terms of presuming the least amount of empirical information about the rest of the Humean mosaic. These predicates contain the least amount of empirical information about the rest of the Humean mosaic because any other predicate that describes the composition of the mosaic must describe a multiplicity of these elementary building blocks, and therefore – assuming the world in question is finite – restricts what can be said about how the remainder of the Humean mosaic is configured. Namely, the instantiation of "is a quark" is implied by "is a bundle of quarks in the formation of a proton" and with "is a bundle of quarks in the formation of a neutron", but not vice-versa. This is pertinent to the Shannon modal entropy conjecture, since Shannon modal entropy is precisely that measure of empirical content that formalizes the above line of reasoning: there are more possible worlds with quarks than with neutrons or protons, since all worlds with the latter contain the former, but not vice-versa. The Shannon modal entropy of a language containing "is a quark" is therefore greater than one containing "is a bundle of quarks in the formation of a proton" and "is a bundle of quarks in the formation of a neutron".

While the above line of reasoning is not a proof, it does lend significant intuitive support to the notion that fundamental laws should be written in fundamental predicates, and that 'fundamental' in this respect correlates well with some concept of the amount of empirical content of the predicate that may be formally measured using the tools of modal metaphysics and Shannon modal entropy.

6.7 Objections and conclusion

For the strict Humean, non-Humean PNPs and the quiddditicism that results are contrary to the parsimony of Humean metaphysics. Although Lewis's introduction of a privileged set of properties that were perfectly natural did not explicitly invoke any second order property, if the BSA is to be a realist theory of lawhood, then it is difficult to see what else perfect naturalness could be if it is to be a founding notion of such a theory. Perfectly natural properties that are defined instead in terms of human science (such as Cohen and Callender's "Better Best System Account", 2009) contradict the very purpose of developing a Humean metaphysics that does not slip into anti-realism.

To achieve a truly Humean metaphysics it has been argued here that a concept of perfect naturalness can be devised where PNPs are instead supervenient on modal space $\mathcal{M}$ or some part of it; the only other primitive Humean entity in Lewis's metaphysics.

The modal entropy conjecture effectively identifies the most fundamental empirical predicates, and therefore the language in which the best system must be expressed. Incorporating the maximization of Shannon modal entropy into the set of optimization parameters of the BSA allows these predicates to be identified as PNPs without circularity. Because the modal entropy conjecture identifies PNPs without presuming anything regarding their metaphysical status, they can remain supervenient entities that supervene on patterns of instantiation in worlds comprising modal space $\mathcal{M}$.

The modal entropy conjecture presumes Lewisian modal realism. That is, because Shannon modal entropy is a measure taken over modal space $\mathcal{M}$, and since every point in $\mathcal{M}$ is a logically possible world, then Lewisian modal realism is implicitly presumed. However, Shannon modal entropy may in theory at least be measured over ersatz possible worlds, or any concept of non-concrete possible world, just so long as the resulting modal space does not contain holes, for example. That is, the language used to construct the space of ersatz

possible worlds is formally complete with respect to modal space $\mathcal{M}$ of the modal realist. Lewis (1986) gives some reasons for thinking non-concrete formulations of possible worlds would be incomplete, but a full discussion of whether this implies their wholesale rejection is beyond the scope of this thesis.

It may also be objected that a Humean best systematization of the laws of nature using supervenient PNPs cannot get off the ground, since in the Humean ontology, properties are not reified, but properties must be presupposed in order to measure the corresponding entropy values across $\mathcal{M}$. This objection is incorrect on two counts: first, the Humean ontology can accept reified properties, just so long as they are supervenient on the Humean mosaic; second, there's no reason why modal entropy could not be calculated using nominal properties, since all that matters is the pattern of instantiation across points in $\mathcal{M}$. While further discussion of Humean theories of properties is also beyond the scope of this property, it is noteworthy that Esfeld (2020) has proposed a Humean theory of PNPs that presupposed no reified properties. He argues that all physical laws can ultimately be expressed in terms of relative spacetime co-ordinates and trajectories. Whether this is true or not is beside the point: in worlds like ours, predicates denoting spacetime co-ordinates or trajectories, seem like they would have high modal entropy, and thus provide excellent candidates to be PNPs.

7. Conclusions

The purpose of this thesis has been to introduce the version of the BSA developed by Lewis, and to develop the account further by proposing solutions to four outstanding problems. These problems are the zero-fit problem, the problem of undermining, the problem of non-locality as it occurs in quantum mechanics and Armstrong's problem. In chapters 3 through 6 of this thesis, each of these have been shown to have non-arbitrary solutions that retain both the realism of Lewis's original formulation of the BSA, as well as its Humean character. Indeed, the introduction of the concept of supervenient perfectly natural properties developed in chapter 6 in place of the primitive analogues proposed by Lewis, makes the BSA theory developed here more Humean than Lewis's original formulation.

To recap the problems and the solutions proposed: the zero-fit problem results from the possibility that there are infinite permutations of chance events in some possible worlds for a given chance outcome. For such worlds, the notion of fit used in Lewis's formulation as the likelihood of a particular world or worlds to be most similar to our own in terms of the pattern of chance events then becomes a division by an infinite quantity, meaning all worlds have zero fit. It was argued in chapter 3 that, in the first place, fit should not be made in terms of the actual sequence of chance events, since this contradicts a reasonable assumption of statistical independence, but rather in the resulting relative frequencies. It was then argued that, given this assumption, a measure commonly used in discrete statistics, the X^2 statistic, could be used instead of likelihood to measure fit.

The second problem, addressed in chapter 4, is that of undermining futures. In this problem, the possibility of worlds that are presumed to be nomologically possible in terms of their non-chancy character relative to the actual world have relative frequencies of chance events would contradict the relative frequencies and resulting indeterministic laws

of the actual world. The solution proposed in chapter 4 to solve this contradiction was the derivation of a group of worlds that each shared the same deterministic laws and could be ordered in some branching structure defined by the probability spaces and resulting sequences of all possible chance outcomes. It was shown how this branching structure, referred to as a T-world, could be established a priori to the formulation of indeterministic laws, and would serve as the Humean mosaic from which those laws would be derived according to a more complex non-circular systematization process. It was shown that performing the systematization process over the T-world Humean mosaic would satisfy the Basic Chance Principle, and therefore avoid the problem of undermining.

The third problem of non-locality addressed in chapter 5 arose from a phenomenon observed in the quantum mechanics of the actual world whereby fundamental properties could involve 'spooky' connections between spatially separated particles. This has been thought to contradict a foundational principal of Humean metaphysics: that such necessary connection could not exist in the Humean mosaic. The solution proposed, after Esfeld (2014), was to note that the notion of necessity appealed to in the Humean doctrine of no necessary connections, was that of metaphysical necessity, not nomological necessity. Yet, the necessity involved in non-local phenomena and the laws of physics that describe such phenomena is nomological necessity. Thus, the Humean doctrine is not in contradiction with the laws of physics since such laws according to the BSA, even were they to be fundamental, would be supervenient on the underlying Humean mosaic. Because the Humean mosaic can be described exclusively in local terms, albeit with a loss of expediency and elegance than would be achieved by inferring non-local properties, ultimately there is no contradiction.

Finally in chapter 6, Armstrong's problem, as manifested in its most extreme form by the problem of trivialization, was addressed. This problem states that requirements of trade-

offs between strength and simplicity of competing systems of laws can only be meaningful and non-trivial if a privileged language was used that did not allow trivializing predicates such as “the best systematization of the Humean mosaic into laws for world w”. For, such predicates would invariably be the simpler and stronger than any more detailed formulation of laws. Lewis’s solution of requiring the candidate laws of nature to be expressed using only properties that had the second-order property of perfect naturalness was rejected as non-Humean. Instead, it was shown how a concept of entropy measured over all of modal space and the distribution of instantiated properties at each point in that space could be calculated. This measure of entropy, analogous to entropic measures in information theory and thermodynamics, was referred to as modal entropy. It was shown how the additional criterion of a system of laws that maximized this measure by the choice of predicates the laws were expressed in, would result in a notion of perfect naturalness that would be supervenient on modal space, rather than stipulated as an ontological primitive.

With solutions proposed for all of these four problems, some of the most serious objections to the BSA theory of lawhood as originally formulated by Lewis are overcome. As a result, the revised BSA theory of lawhood presented in this thesis means that it can be more widely accepted as a non-reductive realist theory of lawhood.

8. Bibliography

Armstrong, DM. 1978. *Universals and Scientific Realism: Nominalism and Realism.* 1 ed. Cambridge: Cambridge University Press.

Armstrong, DM. 1985. *What is a Law of Nature?* Cambridge: Cambridge University Press.

Arntzenius, Frank. 2008. Gunk, Topology, and Measure. In *Oxford Studies in Metaphysics vol. 4.* Edited by Dean Zimmerman. Oxford: Oxford University Press.

Baker, David J., "Measurement outcomes and probability in Everettian quantum mechanics," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38 (1): 153-169 (2007).

Barker, Stephen, "Can Counterfactuals Really Be about Possible Worlds?," *Nous* 45 (3): 557-572 (2011).

Beebee, Helen, "The Non-Governing Conception of Laws of Nature," *Philosophy and Phenomenological Research* 61 (3): 571-594 (2000).

Bell, John. 1987. *Speakable and Unspeakable in Quantum Mechanics.* Cambridge: Cambridge University Press.

Beth, E. W., "On Padoa's Method in the Theory of Definition," *Journal of Symbolic Logic* 21 (2): 194-195 (1956).

Bigelow, John, John Collins, and Robert Pargetter, "The Big Bad Bug: What Are the Humean's Chances?," *The British Journal for the Philosophy of Science* 44 (3): 443-462 (1993).

Briggs, Rachael, "The Big Bad Bug Bites Anti-Realists about Chance," *Synthese* 167 (1): 81-92 (2009).

Briggs, Rachael, "The Anatomy of the Big Bad Bug," *Nous* 43 (3): 428-449 (2009).

Butterfield, Jeremy, "Emergence, Reduction and Supervenience: A Varied Landscape," *Foundations of Physics* 41 (6): 920-959 (2011).

Carrol, JW. 1994. *Laws of Nature.* Cambridge: Cambridge University Press.

Cohen, Jonathan and Craig Callender, "A Better Best System Account of Lawhood," *Philosophical Studies* 145 (1): 1-34 (2009).

Collins, John, "Preemptive Prevention," *The Journal of Philosophy* 97 (4): 223-234 (2000).

Darby, George and Duncan Watson, "Lewis's Principle of Recombination: Reply to Efird and Stoneham," *Dialectica* 64 (3): 435-445 (2010).

Deutsch, David, "Quantum theory of probability and decisions," *Proceedings of the Royal Society of London.Series A: Mathematical, Physical and Engineering Sciences* 455 (1988): 3129-3137 (1999).

Dewar, Neil, "Ramsey Equivalence," *Erkenntnis* 84: 77-99 (2019).

Dziurosz-Serafinowicz, Patryk, "Chance, Resiliency, and Humean Supervenience," *Erkenntnis* 84: 1-19 (2019).

Earman, John. 1986. *A Primer on Determinism.* Vol. 32. Dordrecht: D. Reidel.

Earman, John and John T. Roberts, "Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience," *Philosophy and Phenomenological Research* 71 (1): 1-22 (2005).

Earman, John and John T. Roberts, "Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature. Part II: The Epistemological Argument for Humean Supervenience," *Philosophy and Phenomenological Research* 71 (2): 253-286 (2005).

Einstein, A, B Podolsky, and N Rosen, "Can the Quantum Mechanical Description of Reality be Considered Complete?," *Physical Review* 47: 777-780 (1935).

Elga, Adam, "Infinitesimal Chances and the Laws of Nature," *Australasian Journal of Philosophy* 82 (1): 67-76 (2004).

Emery, Nina, "Chance, Possibility, and Explanation," *The British Journal for the Philosophy of Science* 66 (1): 95-120 (2015).

Esfeld, Michael. 2010. Humean metaphysics versus a metaphysics of powers. In *Time, chance and reduction. Philosophical aspects of statistical mechanics.* Edited by Gerhard Ernst and Andreas Hüttemann. Cambridge: Cambridge University Press.

Esfeld, Michael et al., "The Ontology of Bohmian Mechanics," *The British Journal for the Philosophy of Science* 65 (4): 773-796 (2014).

Esfeld, Michael, "A proposal for a minimalist ontology," *Synthese* 197: 1889-1905 (2020).

Everett, Hugh, ""Relative State" Formulation of Quantum Mechanics," *Reviews of Modern Physics* 29 (3): 454-462 (1957).

Everett, Hugh. 1973. The Theory of the Universal Wavefunction. In *The Many-Worlds Interpretation of Quantum Mechanics.* Edited by B. DeWitt and N. Graham.Princeton University Press.

Frigg, Roman and Carl Hoefer, "The Best Humean System for Statistical Mechanics," *Erkenntnis* 80: 551-574 (2015).

Gao, Shan. Three possible implications of spacetime discreteness. 2010.
Ref Type: Unpublished Work

Greaves, Hilary and Wayne Myrvold. 2010. Everett and Evidence. In *Many Worlds? Everett, Quantum Theory, and Reality.* Edited by Simon Saunders et al.Oxford University Press.

Hall, Ned, "Correcting The Guide to Objective Chance," *Mind* 103 (412): 505-517 (1994).

Hall, Ned, "Causation and the Price of Transitivity," *The Journal of Philosophy* 97 (4): 198-222 (2000).

Hall, Ned. 2015. Humean Reductionism About Laws of Nature. In *A companion to David Lewis.* Edited by Barry Loewer and Jonathan Schaffer. Chichester: John Wiley & Sons, Inc.

Halpin, John F, "Briggs on antirealist accounts of scientific law," *Synthese* 190: 3439-3449 (2013).

Halpin, John F., "Legitimizing chance: The best-system approach to probabilistic laws in physical theory," *Australasian Journal of Philosophy* 72 (3): 317-338 (1994).

Halpin, John F., "Lewis, Thau, and Hall on Chance and the Best-System Account of Law," *Philosophy of Science* 65 (2): 349-360 (1998).

Halpin, John F., "Scientific Law: A Perspectival Account," *Erkenntnis* 58 (2): 137-168 (2003).

Hankerson, Darrel. 2003. *Introduction to information theory and data compression.* Boca Raton: Routledge/CRC Press.

Hempel, Carl. 1965. *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science.* New York: The Free Press.

Hitchcock, Christopher, "Of Humean Bondage," *The British Journal for the Philosophy of Science* 54 (1): 1-25 (2003).

Hoefer, Carl, "On Lewis's Objective Chance: "Humean Supervenience Debugged"," *Mind* 106 (422): 321-334 (1997).

Hoefer, Carl, "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance," *Mind* 116 (463): 549-596 (2007).

Hume, David. 1955. *An Inquiry Concerning Human Understanding.* Edited by Charles W Hendel. Indianapolis: Library of Liberal Arts.

Hume, David. 1985. *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects.*Penguin Classics.

Humphreys, Paul, "How Properties Emerge," *Philosophy of Science* 64 (1): 1-17 (1997).

Jansson, Lina, "Everettian quantum mechanics and physical probability: Against the principle of State Supervenience," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 53: 45-53 (2016).

Karakostas, Vassilios, "Humean Supervenience in the Light of Contemporary Science," *Metaphysica* 10 (1): 1-26 (2009).

Kincaid, H. 2004. There are laws in the social sciences. In *Contemporary Debates in Philosophy of Science.* Edited by Christopher Hitchcock.Blackwell Pub.

Kuhn, Thomas S. 2012. *The Structure of Scientific Revolutions.* Chicago: The University of Chicago Press.

Ladyman, James. 2007. *Every Thing Must Go: Metaphysics Naturalized.* New York: Oxford University Press.

Laplace, P. S. 1951. *A philosophical essay on probabilities.* New York: Dover Publications.

Lewis, David, "Causation," *Journal of Philosophy* 70 (17): 556-567 (1973).

Lewis, David. 1973. *Counterfactuals.* Cambridge, MA: Harvard University Press.

Lewis, David. 1980. A Subjectivist's Guide to Objective Chance. In *The University of Western Ontario Series in Philosophy of Science (A Series of Books in Philosophy of Science, Methodology, Epistemology, Logic, History of Science, and Related Fields).* Vol. 15. Edited by Stalnaker R. Pearce G. Harper W.L. Dordrecht: Springer.

Lewis, David, "New work for a theory of universals," *Australasian Journal of Philosophy* 61 (4): 343-377 (1983).

Lewis, David. 1986. *On the plurality of worlds.* Oxford: Blackwell.

Lewis, David. 1987. A Subjectivist's Guide to Objective Chance. In *Philosophical Papers Volume II.* New York: Oxford University Press.

Lewis, David. 1987. Causal Explanation. In *Philosophical Papers Volume II.* New York: Oxford University Press.

Lewis, David. 1987. Postscripts to ''Causation''. In *Philosophical Papers Volume II.* New York: Oxford University Press.

Lewis, David. 1987. *Philosophical Papers Volume II.* New York: Oxford University Press.

Lewis, David, "Humean Supervenience Debugged," *Mind* 103 (412): 473-490 (1994).

Lewis, David. 1999. *Papers in metaphysics and epistemology.* Vol. II. Cambridge: Cambridge University Press.

Lewis, David, "Causation as Influence," *The Journal of Philosophy* 97 (4): 182-197 (2000).

Lewis, David, "How many lives has schrodinger's cat?," *Australasian Journal of Philosophy* 82 (1): 3 (2004).

Lewis, David. 2009. Ramseyan humility. In *Conceptual Analysis and Philosophical Naturalism.* Edited by David Braddon-Mitchell and Robert Nola. Cambridge, MA: MIT Press.

Loewer, Barry, "Humean Supervenience," *Philosophical Topics* 24: 101-127 (1997).

Loewer, Barry, "Laws and Natural Properties," *Philosophical Topics* 35 (1&2): 313-328 (2007).

McDermott, Michael, "Redundant causation," *British Journal for the Philosophy of Science* 46 (4): 523-544 (1995).

Mill, John Stuart. 1904. *A system of logic: Ratiocinative and induction.* 8th ed. New York: Harper & Brothers.

Nagel, Ernest. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation.*Harcourt, Brace & World.

Niiniluoto, Ilkka, "Scientific progress as increasing verisimilitude," *Studies in History and Philosophy of Science Part A* 46: 73-77 (2014).

Paul, L. A., "Keeping Track of Time: Emending the Counterfactual Account of Causation," *Analysis* LVIII: 191-198 (1998).

Paul, L. A., "Building the world from its fundamental constituents," *Philosophical Studies* 158 (2): 221-256 (2012).

Placek, Tomasz, "Comparative Similarity in Branching Space-Times," *International Journal of Theoretical Physics* 49 (12): 3233-3242 (2010).

Popper, Karl, "A note on verisimilitude," *British Journal for the Philosophy of Science* 27 (2): 147-159 (1976).

Putnam, Hilary. 1975. Scientific realism in mathematics and elsewhere. In *Philosophical Papers (vol I).* Cambridge: Cambridge University Press.

Ramsey, F. P. 1990. *Philosophical papers.* Edited by D. H. Mellor. Cambridge: Cambridge University Press.

Reynolds, Andrew. 2002. *Peirce's Scientific Metaphysics: The Philosophy of Chance, Law, and Evolution.* Nashville: Vanderbilt University Press.

Russell, Bertrand, "On the Notion of Cause," *Proceedings of the Aristotelian Society* 13: 1-26 (1912).

Schaffer, Jonathan, "Is There a Fundamental Level?," *Nous* 37 (3): 498-517 (2003).

Schrenk, Markus. 2008. A Theory for Special Science Laws. In *Selected Papers Contributed to the Sections of GAP.6, 6th International Congress of the Society for Analytical Philosophy, Berlin, September 2006.* Edited by S Walter. Paderborn: Mentis.

Schrenk, Markus, "Can physics ever be complete if there is no fundamental level in nature?," *Dialectica* 63 (2): 205-208 (2009).

Schwarz, Wolfgang. 2014. Proving the Principal Principle. In *Chance and Temporal Asymmetry.* Edited by Alistair Wilson.Oxford University Press.

Seiberg, Nathan. Emergent Spacetime. arXiv . 2006. Cornell Uniersity.
Ref Type: Electronic Citation

Shannon, C E, "A Mathematical Theory of Communication," *The Bell System Technical Journal* 27: 379-423 (1948).

Solomonoff, R. J., "A formal theory of inductive inference. Part II," *Information and Control* 7 (2): 224-254 (1964).

Thau, Michael, "Undermining and Admissibility," *Mind* 103 (412): 491-503 (1994).

Tooley, Michael. 1993. Causation: reductionism versus realism. In *Causation.* Edited by E Sosa and M Tooley.Oxford University Press.

Tosh, Nick, "Finite frequentism in a big world," *The British Journal for the Philosophy of Science* 67: 169-213 (2014).

Van Fraassen, Bas C. 1980. *The Scientific Image.*Oxford University Press.

Van Fraassen, Bas C. 1989. *Laws and Symmetry.*Oxford University Press.

Van Fraassen, Bas C. 2002. *The empirical stance.* New Haven, CT: Yale University Press.

van Strien, Marij, "On the origins and foundations of Laplacian determinism," *Studies in History and Philosophy of Science Part A* 45: 24-31 (2014).

Wallace, David, "Worlds in the Everett interpretation," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 33 (4): 637-661 (2002).

Wallace, David, "Everettian rationality: defending Deutsch's approach to probability in the Everett interpretation," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 34 (3): 415-439 (2003).

Wallace, David, "Quantum probability from subjective likelihood: Improving on Deutsch's proof of the probability rule," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38 (2): 311-332 (2007).

Wallace, David. 2009. A formal proof of the born rule from decision-theoretic assumptions [aka: How to Prove the Born Rule]. In *Many Worlds?: Everett, Quantum Theory & Reality.* Edited by Simon Saunders et al.Oxford University Press.

Wallace, David. 2012. *The Emergent Multiverse.* Oxford: Oxford University Press.

Ward, Barry, "Humeanism without Humean Supervenience: A Projectivist Account of Laws and Possibilities," *Philosophical Studies* 107 (3): 191-218 (2002).

Ward, Barry, "Projecting Chances: A Humean Vindication and Justification of the Principal Principle," *Philosophy of Science* 72 (1): 241-261 (2005).

Wheeler, Billy. Laws, Natural Properties and Algorithmic Compression. 23-6-2014.
Ref Type: Unpublished Work

Wheeler, Billy, "Simplicity, Language-Dependency and the Best System Account of Laws," *Theoria* 31 (2): 189-206 (2016).

Williams, JR, "Chances, Counterfactuals, and Similarity," *Philosophy and Phenomenological Research* 77 (2): 385-420 (2008).

Wilson, Alastair, "Everettian quantum mechanics without branching time," *Synthese* 188 (1): 67-84 (2012).