Reduction as an *A Posteriori* Relation

Joshua Rosaler

Abstract

Reduction between theories in physics is often approached as an *a priori* relation in the sense that reduction is often taken to depend only on a comparison of the mathematical structures of two theories. I argue that such approaches fail to capture one crucial sense of "reduction," whereby one theory encompasses the set of real behaviors that are well-modeled by the other. Reduction in this sense depends not only on the mathematical structures of the theories, but also on empirical facts about where our theories succeed at describing real systems, and is therefore an *a posteriori* relation.

1 Introduction

One important sense of the term "reduction" requires that it be possible to model all real behaviors that are well-modeled in the reduced description at least as accurately in the reducing description. So, for example, the claim that Newtonian mechanics reduces to [1] special relativity is often interpreted to mean that any behavior that is modeled successfully (i.e., within some reasonably close margin of error) by Newtonian mechanics can be modeled at least as accurately in special relativity. Analogous claims are often made about the relationship between classical and quantum mechanics, Newtonian gravitation and general relativity, quantum mechanics and quantum field theory, thermodynamics and statistical mechanics, and many other theory pairs; a theory of quantum gravity, presumably, should reduce both general relativity and the Standard Model of particle physics in this sense. Unless explicitly stated otherwise, the reader should understand "reduction" here as the relationship whereby one description successfully models all real behaviors that are well-modeled on another - or, somewhat more concisely, the relationship whereby one description subsumes the domain of applicability of another.

While implicitly taking reduction to include this requirement of domain subsumption, much of the literature on reduction in physics also approaches reduction as an *a priori* relation in the sense that questions of reduction are assumed to depend only on the relationship between the mathematical structures of the two theories. Here, I argue that the tendency to view reduction as domain subsumption and the tendency to treat reduction as though it were

[1] Here, I will employ the "philosopher's" convention of referring to a less encompassing theory as being "reduced to," or "reducing to" a more encompassing one. As Nickles and others have noted, the opposite convention is usually employed in the physics literature [17].

an *a priori* relation in this sense are mutually incompatible, at least within the context of most realistic inter-theory relations in physics. More precisely, I argue that knowledge of the mathematical structures of two theories alone is not generally sufficient to determine whether one encompasses the domain of applicability of the other, and that further empirical input concerning the scope and precision of our theories in describing real physical systems is also needed. Whether the degree of dovetailing between two theories, as determined through mathematical analysis, suffices for domain subsumption depends on facts about the particular world we inhabit. Elaborating on this general line of argument, I show below that reduction, if understood as domain subsumption, is an *a posteriori* relation between theories or models.

Here, I distinguish between what I call "formal" and "empirical" approaches to reduction:

1. **"Formal" Approaches to Reduction** treat reduction as an *a priori* relation between theories/models - i.e., as a *two-place* relation between abstract descriptions. They assume that, given two theories or models, one can determine on the basis of logical or mathematical analysis alone whether one subsumes the domain of applicability of another. No further empirical input is required to determine whether reduction in this sense holds.

2. **"Empirical" Approaches to Reduction** treat reduction as an *a posteriori* relation between theories/models - i.e., as a *three-place* relation among some set of systems and two alternative ways of representing those systems. Domain subsumption is taken to rest not only on an abstract analysis of logical or mathematical relations between these representations, but also on further empirical input concerning where they succeed at describing real physical behaviors.

I argue that formal approaches operate at a high level of abstraction away from real physical systems and that as a result, the physical import of such analyses for questions of domain subsumption is often obscure. I then argue that a more direct and physically transparent line of attack on questions of domain subsumption is provided by an empirical approach, which keeps the real behaviors that are described by our theories, and the correspondence between our theories and real systems, constantly in view.

Understood as domain subsumption, reduction in the present context is not a direct relationship between high- and low-level [2] descriptions, but rather between the low-level description and those particular systems and *behaviors*

[2]Given a pair of theories or models, I use the term "high-level" to designate a theory that is purportedly reduced and "low-level" to designate the theory that purportedly does the reducing. By contrast with the reduced/reducing distinction, my use of the high/low terminology here is intended to allow for the possibility that reduction between the theories remains unproven and merely conjectural.

that are well-modeled on the high-level description. From this perspective, questions of reduction are questions about how broadly the domain of reality that is well-described by the low-level theory extends - specifically, whether it extends to include all behaviors that are well-modeled by the reduced theory. It is not inconceivable that one theory could subsume the domain of another solely *by virtue* of a direct relationship between the mathematical structures of the theories, and without regard to empirical facts about where our theories succeed at describing real systems. In a case where the high-level theory is strictly speaking a special case of the low-level theory, for example, any behavior that is well-modeled on the high-level theory is, *a fortiori,* well-modeled on the low-level theory, irrespective of empirical facts concerning where the theories succeed at describing real systems. However, this sort of scenario does not seem to be borne out in many, if any, of the inter-theoretic relations that tend to be of interest in the physics or philosophy of physics literature. In most realistic cases, domain subsumption occurs through some approximate dovetailing between the mathematical structures of the theories. One of the essential points that I emphasize here is that whether this dovetailing (which holds solely by virtue of the mathematical relationship between the two structures) is sufficiently precise or robust for the low-level theory to encompass the domain of successful applications of the high-level theory is unavoidably empirical in most, if not all, interesting cases.

The discussion is outlined as follows. In Section 2, I clarify further what is meant, and what is not meant, by the requirement that one theory subsume the domain of another. In Section 3, I describe a formal approach to reduction that seeks to extract one theory as a mathematical limit of another, and argue that the connection between results generated by this approach and questions of domain subsumption is often obscure. In Section 4, I argue on the basis of examples that reduction depends on both the mathematical relationship between theories and on further empirical facts about the correspondence between our theories and the world; I then describe an empirical strategy for reduction that connects more directly and transparently to questions of domain subsumption. This approach rests on a particular type of empirically constrained mathematical relationship between two models (typically, one from each theory) of the same physical system. Section 5 is the conclusion.

2 Reduction as Domain Subsumption

In this section, I attempt to characterize more precisely the sense of reduction with which we will be concerned here, which requires that one description of nature (in the context of physics, a theory or model) subsume the domain of applicability of another. As suggested above, subsumption of the domain of some high-level description by a low-level description requires that any real behavior that is accurately represented by the high-level description be represented at least as accurately and in at least as much detail by the low-level description.

So, for example, the reduction of classical to quantum mechanics requires that every classical system [3] be more precisely regarded as a quantum system, the reduction of Newtonian mechanics to special relativity that every Newtonian system be more precisely regarded as a relativistic system, the reduction of quantum mechanics to quantum field theory that every quantum mechanical system be more more precisely regarded as a quantum-field-theoretic system, and so on.

To further characterize this sense of reduction, it will serve to contrast it with some views of reduction that have been advanced in the philosophical and scientific literatures.

In doing so, it will be useful to draw a distinction here between *concepts of* reduction and *approaches to* reduction. A concept of reduction is a particular *meaning* that one assigns to the term "reduction," while an approach to reduction is a particular *strategy* for showing that some particular concept of reduction holds. Existing accounts of reduction, such as the "physicist's" notion that reduction is generally a matter of extracting one theory as a mathematical limit of another, or philosophical accounts such as Nagel/Schaffner reduction, may be considered either as concepts of or as approaches to reduction. [4] Here, we consider them as approaches to reduction, since reduction in the sense that we consider here is not *fundamentally* about recovering one theory as a mathematical limit of another, as would be the case in a limit-based concept of reduction, or about deriving the laws of one theory from those of another, as would be the case in the Nagel/Schaffner concept of reduction. *Fundamentally*, the concept of reduction that we investigate here is about showing that all real behaviors that can be accurately modeled in one theory can be modeled at least as accurately in another. Taking limits and deriving one set of laws from another may turn out to be useful *strategies* toward this goal, but neither requirement is regarded from the outset as a *sine qua non* of reduction.

In some important respects, the concept of reduction as domain subsumption is closely akin to the concept associated with Kemeny and Oppenheim's well-known account of reduction, which requires that the reducing theory explain all observational data that is explained by the reduced theory and that the reducing theory be at least as "well-systematized" [5] as the reduced theory

[3] By "classical system," I do not mean that the system conforms exactly to classical mechanics in all of its features - only that some of its features are described to good approximation by the regularities of classical mechanics. "Quantum system," "quantum-field-theoretic system," etc. should be interpreted analogously.

[4] For recent, up-to-date discussions of the Nagel/Schaffner account of reduction, see Dizadji-Bahmani et al's [9] and Schaffner's [23]. For a recent account of the relationship between the Nagel/Schaffner and limit-based approaches, see Butterfield's [6].

[5] Roughly, the degree of systematization of a theory corresponds to the "ratio" of number of phenomena that the theory explains to the number of fundamental assumptions that the theory makes. In practice, this is very difficult to quantify in any precise or systematic way. Nevertheless, there is nearly uniform agreement that, however degree of systematization is to be defined precisely, special relativity is "better systematized" than Newtonian mechanics,

[16]. One salient difference between the concept of reduction considered here and the Kemeny/Oppenheim concept is that our concept does not rely on the existence of a clear distinction between observational and theoretical realms. Instead, we require that all behaviors - whether directly observable or not - that are well-modeled in the high-level description also be well-modeled in the low-level description. The term "well-modeled" here presupposes that there is some fact about the closeness of fit between the possibly unobservable [6] features of a given physical system and their representations in our mathematical models. In this respect, our concept of reduction presupposes a broadly realist view of the correspondence between empirically successful theories and the physical world, in contrast to the logical empiricist view of theories that formed the background for Kemeny and Oppenheim's account.

Much of the skepticism toward reduction in the philosophical literature of the past several decades has been rooted in the famous critiques of reductionism based on multiple realizability given by Putnam and Fodor [12], [18]. It is important to emphasize that the concept of reduction that we consider here is clearly distinct from the concepts of reduction investigated by these authors, and is consistent with mutliple realization. The types of reduction that Putnam and Fodor consider impose requirements that go beyond what is necessary for one theory to encompass the domain of another in the sense that concerns us here, and it is precisely because of these additional requirements that these types of reduction fail in cases where multiple realization occurs.

To further clarify the concept of reduction as domain subsumption, it will prove useful to contrast it with the concept of reduction considered by Putnam. In Putnam's analysis, reduction requires a low-level description to "explain" all behaviors that are explained by the corresponding high-level description. The specific notion of explanation that Putnam adopts prohibits the inclusion of details that are extraneous in the sense that changes in these details do not affect the occurence of the high-level behavior that we wish to explain. In cases of multiple realization, a low-level description of some high-level behavior - to take Putnam's example, a detailed molecular description of a wooden peg that fits into one hole but not another - includes details that are extraneous in just this sense, so that the low-level molecular description fails to explain (in Putnam's sense) behaviors that the high-level description explains well. The sense of reduction that we consider here does not rely on anything like Putnam's sense of explanation. A behavior that lies in the domain of some high-level description is not excluded from the domain of a low-level description simply because the low-level description of this behavior is more detailed, or because the high-level behavior would remain unchanged under certain alterations in these details. What's important for domain subsumption, rather,

general relativity "better systematized" than Newtonian gravitation, and likewise in many other cases.

[6] By "unobservable" here, I do not mean that we have no empirical access to these features, only that they our empirical access to them is relatively indirect.

is that the quantities in the low-level description that represent the relevant high-level features of a system faithfully track (to within some suitably small margin of error) the behavior of those features in all cases where the high-level description does.

As another point of clarification, it will help to contrast reduction as domain subsumption with the sense of reduction considered in Fodor's discussion. Whereas Putnam focuses the question of whether a low-level description can explain all singular occurrences that a low-level description does, reduction on Fodor's analysis requires a low-level description to explain the *laws* of a high-level description. This, in turn, requires that it be possible to identify a given natural kind [7] in the high-level description with a single natural kind in the low-level description across all contexts where the high-level description applies. But in cases where multiple realization occurs, a natural kind according to the high-level description may be instantiated by many different low-level features across different contexts, where the disjunction of these low-level features cannot reasonably be regarded as a natural kind of the low-level description. As a result, a given natural kind of the high-level description cannot be identified with any single natural kind of the low-level description, so that reduction in Fodor's sense fails in cases of multiple realization. On the sense of reduction examined here, it is not necessary that a given natural kind of the high-level description be identified with some natural kind in the low-level description. As long as there is *some* low-level description of each instance of the high-level behavior in question that represents that behavior at least as accurately as the high-level description does, we may say that the low-level description has subsumed the domain of the high-level description. This condition may be satisfied even when a given natural kind in the high-level description cannot be identified with any single natural kind in the low-level description.

In response to Putnam and Fodor, Sober has argued that the notion of explanation adopted by both authors is unreasonably narrow - he argues that detailed microscopic descriptions of macroscopic phenomena, such as Putnam's peg, do not fail to be explanations of these phenomena simply because they are highly detailed or because certain changes in these microscopic details leave the macroscopic features of the system unchanged [24]. More recently, Batterman has argued that Sober's response misses the force of the multiple realizability argument because a highly detailed microscopic explanation of some multiply realized macroscopic phenomenon still leaves us without an explanation of why "systems that are heterogeneous at some micro-scale exhibit the same pattern of behavior at the macro-scale" - that is, it leaves us without

[7] A natural kind is generally understood to be any among a set of physical properties that "carves nature at its joints" - often, in the sense that these properties are those that occur in the laws of some physical theory (for example, mass and electric charge). Fodor himself does not use the term "natural kind," although his critique of reductionism if often assimilated in these terms - see, for example, Sober [24].

an explanation of why multiple realization occurs [3]. While explanations of multiple realization are certainly desirable - and, as Batterman has emphasized, can often be effected through renormalization group analysis - they are not required for reduction in the sense of domain subsumption. It may be that we possess many distinct low-level accounts of the same high-level regularity across different contexts, and therefore have domain subsumption, but lack an understanding of the salient commonality across all of these contexts that explains the occurence of the same regularity across all of them. While the issues of domain subsumption and multiple realization are closely related, they are distinct. In fact, questions about how to explain multiple realization seem to presuppose that reduction in our sense holds, in that they assume that high-level phenomena are encompassed by the low-level description in the sense that domain subsumption requires.

One might worry that in foregoing the problematic requirements imposed by Putnam and Fodor's construals of "reduction" and Batterman's requirement that a reduction furnish an explanation of multiple realizability, we have diluted the concept of reduction to a point where it is trivial or no longer interesting. Such worries are unfounded. It is precisely the sense of reduction that we adopt here that most directly concerns the question of how and whether it is possible, say, to accurately re-frame the world of everyday experience within the strange and powerful conceptual frameworks of quantum mechanics or general relativity. Such questions, I take it, are of profound interest both from a scientific and a metaphysical perspective. To see why the requirement of domain subsumption is non-trivial, one need only understand how it could fail to hold: for example, it would fail if the low-level description did not provide any representation of those features of a system that are well-modeled by the high-level description; it would also fail in cases where the low-level description did provide such a representation, but where the quantities that the low-level description took to represent these features of the system did not reflect the behavior of the system in all cases where the high-level description did. The non-triviality of this type of reduction is further supported by the fact that several prominent thinkers, including Nancy Cartwright and John Dupré, explicitly deny that this type of reduction holds in many of the cases where it is often presumed to hold - for example, between classical and quantum mechanics, or between thermodynamics and statistical mechanics [7], [11]. These authors advocate a more pluralistic metaphysics in which the domains of applicabilty of our theories are largely disjoint and in which domain subsumption therefore fails. Beyond being non-trivial, the sense of reduction as domain subsumption is relevant to a wide range of important scientific, metaphysical, and epistemological questions. Physicists' search for a theory of quantum gravity is the search for a theory that encompasses the domains of applicability of both general relativity and the Standard Model of particle physics, and that therefore reduces these theories in the sense we have been discussing. Within the foundations and philosophy of physics, questions about how far the metaphysical implications of a given physical theory extend depend critically on whether that theory subsumes the domains of applicability of theories that are purportedly less encompassing. Reduction in this sense also relates closely to the issue of what sorts of features of physical theories are preserved across theory changes, and to historical questions about how physicists construct new theories from existing ones.

3 A Formal Approach to Reduction in Physics

Having narrowed our concept of reduction, we now turn to the question of how one should determine whether this type of reduction holds between a given pair of physical theories. In this section, we will consider a certain type of formal approach to reduction in physics, which attempts to show that domain subsumption holds between theories solely by virtue of a certain limiting relation between their mathematical formalisms. I will argue that this approach, assuming that it is aimed at questions of domain subsumption and not simply at the elucidation of mathematically interesting correspondences, presents a misleading picture of the relationship between theories whereby one theory may encompass the domain of another.

An example of a formal approach to reduction is provided by the so-called "Bronstein cube" of physical theories, illustrated in Figure 1. The diagram indicates that classical mechanics, special relativity, Newtonian gravity, and general relativity should be the limits as Planck's constant $\hbar \to 0$, respectively, of quantum mechanics, relativistic quantum field theory, non-relativistic quantum gravity, and quantum gravity. Likewise, it indicates that classical mechanics, quantum mechanics, Newtonian gravity and non-relativistic quantum gravity should be the limits as $\frac{1}{c} \to 0$ (where c is the speed of light), respectively, of special relativity, relativistic quantum field theory, general relativity, and quantum gravity. Finally, it also indicates that classical mechanics, special relativity, quantum mechanics and relativistic quantum field theory should be recovered in the limit as the gravitational constant $G \to 0$, respectively, of Newtonian gravity, general relativity, non-relativistic quantum gravity, and quantum gravity. Although not always explicitly stated, it is often implicitly suggested that such limiting relations serve to ensure that all real behaviors that are well-modeled on the high-level theory can be modeled at least as accurately on the low-level theory. The idea of the Bronstein cube is thought to originate in a paper by Gamow, Ivanenko, and Landau, [14], and is further discussed in [25], [10], [8], and elsewhere.

The point that I wish to emphasize here is that the connection between such limits and issues of domain subsumptions is left obscure in most applications of this approach to reduction. An initial source of obscurity is that all of the quantities varied in these limits are constants of nature, and so are fixed for real systems. Notwithstanding our license to vary these quantities mathematically, it is initially unclear what relevance these limits have for the

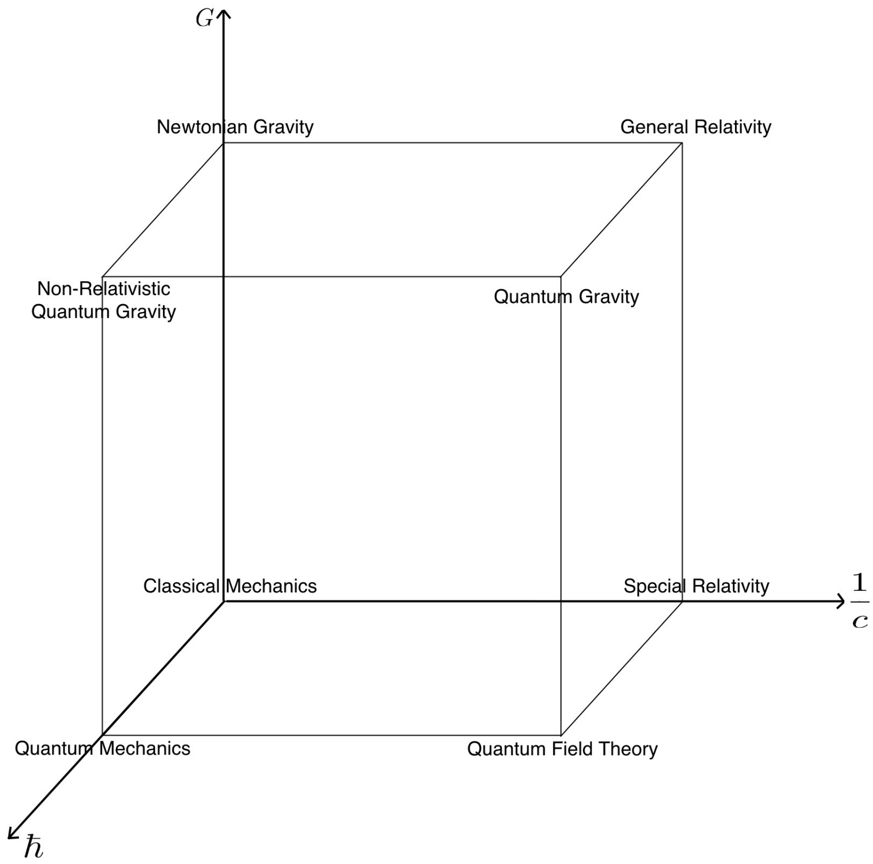

Figure 1: The Bronstein Cube of physical theories.

behavior of systems where the values of these constants never change. A common reply to this criticism is that the limits represented in the Bronstein cube are really shorthands for other related limits that involve a dimensionless ratio of the constant in question and some quantity of the same units that *can* be varied within or across real systems (see, for example, [4]). For example, it is often said that the limit $\hbar \to 0$ should be interpreted not literally, but rather as shorthand for the limit $\frac{\hbar}{S} \to 0$, where $\hbar$ is kept constant and S, a measure of the "typical classical action" of a system, is taken to values much larger than $\hbar$. Similarly, it is said that the limit $\frac{1}{c} \to 0$ should be interpreted as shorthand for the limit $\frac{v}{c} \to 0$, where c is kept constant and v, the velocity of a projectile or the relative velocity of two reference frames, is taken to values much smaller than c. Likewise, presumably, the limit $G \to 0$ should be interpreted as shorthand for the limit $\frac{G}{\Gamma} \to 0$, where G is kept constant and Γ, which has the same units as G, is taken to values much larger than G. Moreover, this approach tends to assume without proof that taking the limit as a dimensionful constant of nature ϵ goes to zero generally gives the same results as taking the limit of the associated dimensionless parameter $\tilde{\epsilon}$ approaches zero as ϵ is kept fixed.

Note that while adopting this "dimensionless" interpretation of the limit-based approach to reduction makes its connection to real physical systems more transparent than the version of the approach that varies constants of nature, the notion that reduction is a matter of simply extracting one theory from another in a dimensionless limit continues to reflect a formal rather than an empirical approach to reduction. The claim that one theory should be a dimensionless limit of another still concerns the mathematical formalisms of the theories exclusively. To see this, simply note that the variable quantities v, S, and Γ are all defined in the low-level theory's formalism; exploring the mathematical consequences of taking these quantities to zero or infinity does not require any reference to the circumstances under which the theories or models in question succeed at describing the behavior of real physical systems.

Taking the dimensionless form of the limit-based approach to be the most viable formulation of this strategy, I now discuss a number of problems that it faces when taken as an approach to domain subsumption.

The first concern is that this approach is underformulated. In practice, the variable quantity making up the dimensionless variable (e.g., S or Γ) is often left unspecified, and it is simply assumed that for some such quantity, one can recover the same results in a dimensionless limit that one recovers in the limits where the associated constant of nature is varied. Neglecting to specify these quantities results in a significant loss of physical insight - in cases where S or Γ is not specified, for example, we cannot be sure of the circumstances under which $\hbar$ or G, respectively, can properly be regarded as "small." Moreover, assuming that these quantities can be identified, it must be shown, for example, that taking $S \to \infty$ while leaving $\hbar$ fixed generally yields the same results as taking $\hbar \to 0$, or that taking $\Gamma \to \infty$ and leaving G fixed yields the same results as taking $G \to 0$. On a fully formulated version of the limit-taking approach, it should be possible to clearly identify the relevant variable quantities, such as S and Γ, and to derive elements of the reduced theory by taking the associated dimensionless limits, rather than having to rely on the more physically opaque procedure of varying constants of nature.

A second worry is that, even allowing that limits of constants can legitimately be interpreted as shorthands for limits of some corresponding dimensionless variables, the limits in Bronstein's cube often do not generate anything that resembles an equation or a quantity of the purportedly reduced theory. For example, the limit $\hbar \to 0$ of Schrodinger's equation, $i\hbar\frac{\partial|\psi\rangle}{\partial t} = \hat{H}|\psi\rangle$, yields the nonsensical result $0 = \hat{H}|\psi\rangle$, [8] rather than an element of classical mechanics. The limit $c \to \infty$ of the relativistic relation $E = pc$ likewise yields the nonsensical result $E = \infty$ rather than anything that resembles a structure of Newtonian mechanics. The limit $G \to 0$ of Einstein's field equations yields the

[8]Purely incidentally, this result happens to bear a close resemblance to the Wheeler-DeWitt equation. This has little significance for our present purposes, given that what we are *supposed* to recover in the limit $\hbar \to 0$ is some equation of classical mechanics.

vaccuum field equations, which allow for strongly curved spacetimes that do not resemble the flat spacetime of special relativity. One could interpret these results as signalling failures of reduction between the theories in question, or, alternatively, one can argue that further general restrictions must be placed on the limit-taking procedure, which rule out the kinds of naive applications of this approach discussed here. However, it is not clear what these restrictions might be, once again raising the concern that the limit-taking approach is excessively vague and underformulated.

A third difficulty, following on this last remark, is that even allowing for the sake of argument that the equations of one theory can be smoothly and unproblematically recovered from another in some dimensionless limit, this may not suffice to ensure that all behaviors that are well-modeled by the high-level theory are also well-modeled in the low-level theory. One can imagine worlds, perhaps physically different from our own, in which the convergence of the low-level theory to the high-level theory in some limit does not occur rapidly enough for the low-level theory to successfully model all real behaviors that are well-modeled in the high-level theory. As I discuss further in the next section, whether a given instance of mathematical dovetailing between two theories or models is sufficient for the low-level description to encompass all real behaviors that are well-modeled in the high-level description depends on empirical facts about where our theories are successful at describing real systems. One theory does not subsume the domain of another solely by virtue of recovering its equations in some limit.

There are still other concerns that one might raise about the notion that reduction is simply a matter of extracting one theory from another in some formal mathematical limit, which relate to the fact that this limiting relationship has been only extremely vaguely defined. Is it necessary that every equation of the low-level theory return some equation of the high-level theory in the specified limit? Or that each equation of the high-level theory be recovered as the limit of some equation of the low-level theory? Or does this limiting relation require only that certain specific elements of the high-level theory be limits of certain corresponding elements of the low-level theory? If so, how are these specific elements to be selected? The more vagueness we tolerate on questions such as these, the less insight we can hope to glean from this approach regarding the general features of the relationship that enables on theory or model to encompass the domain of another. As it currently stands, the limit-taking approach is perhaps best understood as a vague but sometimes useful heuristic, which tells us that taking limits can sometimes be an illuminating exercise in understanding the mechanisms whereby one theory encompasses the domain of another, and also about other features of inter-theory relations not directly related to domain subsumption.

It is not my intention here to suggest that limits do not play an important role in the study of inter-theoretic relations in physics, or even in questions of domain subsumption in certain cases. I only wish to question the notion that subsumption of a high-level theory's domain by a low-level theory is generally a matter of simply taking some limit, as the Bronstein cube and related analyses seem to suggest. This way of thinking about reduction often generates results that are mathematically interesting but physically opaque, particularly on the question of how a low-level theory can be used to model a given phenomenon in the domain of a high-level theory. In the case of quantum-classical relations, this purely formal manner of thinking about reduction is illustrated in Berry's [5] and Batterman's [2], and criticized in [20]. In Berry's work, it seems that "reduction" is often understood to require subsumption of one theory's domain by another, [9] while Batterman takes the term to be defined either by Nagel and Schaffner's account of reduction or by the notion that one theory should be a limit of another. Of course, in the case where reduction is defined by the limit-based account, failure of one theory to smoothly recover another in the relevant limit implies failure of reduction only trivially, and purely as a matter of definition. This kind of definitional failure of the limit-based *concept* of reduction need not - and does not appear to - have any bearing on the question of whether reduction in the sense of domain subsumption occurs in a given case.

Beyond critiquing the sort of naive limit-taking approach to reduction that is exemplified in the Bronstein cube, I also wish to question the notion that it is generally possible to show that one theory subsumes the domain of another exclusively through an analysis of the relationship between their mathematical structures (whether this analysis involves the taking of limits or some other type of procedure). As I argue further in the next section, questions of domain subsumption between theories depend *both* on the relationship between the theories' mathematical structures and on further empirical facts about where the theories are successful at describing real systems.

4 An Empirical Approach to Reduction in Physics

In this section, I begin by arguing that reduction in physics is a partly empirical affair. Then, in Section 4.2, I describe an approach to showing that one theory or model encompasses the domain of another that is more direct, more transparent, and more empirically oriented than the purely formal type of approach discussed in the previous section.

While the question of whether and how one physical theory subsumes the domain of another depends crucially on how the mathematical formalisms of these theories relate, knowledge of the mathematical structures of the theories alone is not generally sufficient to determine whether one theory reduces to the other. Reduction also depends on empirical facts about where and how well our theories succeed at describing real physical systems. While this point is perhaps most clearly illustrated through examples, the general reasoning be-

[9] See, for example, [4].

hind this claim can be summarized as follows. In realistic cases of reduction in physics, the manner in which one theory encompasses the domain of another is typically through some approximate dovetailing between their mathematical formalisms, rather than through one theory's being a mathematically special case of the other. This dovetailing ensures that the low-level theory generates approximately the same predictions as the high-level theory in those cases where the dovetailing holds, but also allows the predictions of the theories to diverge substantially outside of the domain where the dovetailing occurs. The question of whether this dovetailing - which holds solely by virtue of mathematical facts about the relationship between the two theories' formalisms - is sufficiently precise and robust to ensure that the low-level theory succeeds in all cases where the high-level theory does depends on empirical facts about where the high-level theory is successful and where it isn't. From a somewhat different perspective, questions of reduction are, at the most basic level, questions about how far the domain of the low-level theory extends - and, in particular, about whether it extends to include the domain of applicability of the high-level theory. And, of course, questions about where the low-level model succeeds at describing real physical behavior are ultimately empirical.

As an example, consider the relationship between Newtonian mechanics and special relativity, which may seem on its face to be a clear success of a formal approach to reduction. As nearly every undergraduate student of relativity learns, for velocities much less than that of light, relativistic equations of motion closely approximate Newtonian equations, and likewise for solutions of these equations. This result, and the degree of precision within which the approximation holds, follow purely from a mathematical analysis of the relationship between the relativistic and Newtonian equations. However, we can imagine other possible worlds with different physical laws from our own, in which the mathematical relationship between Newtonian and relativistic models necessarily continues to hold, but in which the relativistic model does not encompass all of the empirical successes of the Newtonian model. Perhaps the simplest example of such a world is one in which Newtonian laws hold universally and exactly. In such a world, it would *not* be true that special relativity subsumes the domain of validity of Newtonian mechanics, for special relativity would fail at modeling behaviors that Newtonian mechanics describes accurately. This example shows simply that whether a given low-level theory encompasses the domain of a given high-level theory depends on empirical facts that go beyond the mathematical relationship between the formalisms of the two theories. We can also imagine other worlds in which Newtonian mechanics and relativity apply in domains such that neither theory's domain of success is wholly contained in that of the other, although they overlap in the realm of small velocities.

To take a second, somewhat more complicated example, consider the relationship between classical and quantum models of an alpha particle, whose trajectory in a background magnetic field can be accurately described using classical equations of motion. The classical model represents the trajectory of the alpha particle with a point in phase space whose time evolution is given by Hamilton's equations. The quantum model's description of the alpha particle, which relies heavily on results from decoherence theory (at least, on certain formulations), is more involved [10] However, we can summarize the quantum model's description of the alpha particle as follows. Through environmental decoherence, the combined pure state of the alpha particle and its environment acquires a branching structure. At each branching, one branch is selected probabilistically as the effective state of the system (through a mechanism that depends on one's interpretation of quantum theory). This effective state will be narrowly localized in both position and momentum, so that the quantum description of the alpha particle assigns it well-defined values for these properties. [11] From Ehrenfest's Theorem for open quantum systems, it follows that this trajectory will almost surely approximate the corresponding deterministic trajectory prescribed by the classical model over timescales where spreading of the alpha particle's ensemble wave packet can be ignored. This analysis, which follows purely from an examination of the mathematical relationship between the classical and quantum models, illustrates a specific sort of dovetailing between classical and quantum models of the alpha particle that holds over timescales where ensemble spreading in the quantum model can be ignored. However, even given this analysis, we cannot know whether this dovetailing is sufficient to ensure that the quantum model tracks the particle's behavior over at least those timescales where the classical model does without further empirical input concerning the timescales over which the classical trajectory succeeds at tracking the alpha particle's behavior. There is no logical or mathematical inconsistency in supposing that the timescale over which the classical model tracks the behavior of the alpha particle might be longer than the timescale over which the classical and quantum trajectories are approximately equal according to the above analysis; in such a case, reduction would fail in spite of the mathematical dovetailing just described. [12] Reduction between the models, and by extension between the theories of classical and quantum mechanics, therefore depends on empirical facts that go beyond a purely mathematical analysis of relations between the models.

4.1 A Potential Concern

At this stage of our discussion, it is worth taking a moment to address one likely objection to our argument that reduction is a partly empirical affair. The objection proceeds as follows: it is the *low-level theory*, and not empirical

[10] See [1], [20],[22] and references therein for further discussion of this case.

[11] The sense of narrowness here is defined relative to macroscopic scales of position and momentum, and is in keeping with the uncertainty principle.

[12] Note that there *would* be an inconsistency if we had assumed from the outset that the quantum model applies in all cases where the classical model does - but, of course, this would beg the question, since we are taking this as something that needs to be *shown*.

observation, that determines the domain of validity of the high-level theory, and it does so in a completely *a priori* and mathematical way; in the first example above, it is special relativity that tells us where Newtonian mechanics is and is not successful, solely by virtue of a mathematical analysis of relations between the two theories' equations; in the second example, it is the quantum model of the alpha particle that tells us over what timescales classical equations of motion provide a good approximation to the alpha particle's trajectory, again, solely by virtue of the mathematical relationship between the models; in both cases, we have no need of empirical observation to tell us where the high-level theory is successful at describing real systems since the low-level theory does this for us; reduction is therefore a *a priori* relation.

As I have already suggested above, this objection is question-begging in that by deferring to the low-level theory, rather than empirical observation, to tell us where the high-level theory is and is not empirically successful, it presupposes that the low-level theory furnishes a strictly more accurate description of reality in all cases where the high-level theory is successful - which is precisely what we aim to *show*. Without this assumption, there would be no ground for insisting with such confidence that the low-level theory must be the main authority on where the high-level theory is successful at describing real systems and where it is not. Certainly, we are entitled to pose as a *hypothesis* the view that the low-level theory encompasses the domain of the high-level theory; but the manner in which this hypothesis must be tested is by determining *empirically* whether the high-level theory does in fact succeed only in those cases where the low-level theory entails that it should, and not in any others. While the expectation of a certain unity in nature may reasonably lead one to expect that the low-level theory *should* determine where the high-level theory succeeds at describing real systems and where it does not, in the final analysis this is an empirical hypothesis that depends on facts about the relationship between our theories and the world.

Nothwithstanding these philosophical arguments, in practice, physicists often do rely on a low-level theory to give them some sense of where a high-level theory applies and where it breaks down. To give a few examples: Newton's theory of gravity is often used to determine where we can expect Galilean theories of gravity to succeed and where we can expect them to break down; general relativity is used for this same purpose with regard to Newton' s theory of gravity, as is special relativity with regard to Newtonian mechanics. However, the primary reason, arguably, that we feel confident that this practice will give correct results regarding the high-level theory's domain of validity is that in each case there exists a large body of empirical data confirming that the high-level theory does in fact break down in the circumstances and manner predicted by the low-level theory. It is these empirical data, rather than the low-level theory itself, that most directly constrain the domain of validity of the high-level theory. Without the benefit of such empirical data, it is not clear that we would, or should, be as confident in the practice of using a low-level theory to determine the high-level theory's domain of validity.

4.2 A Local, Empirical, Model-Based Strategy for Reduction in Physics

We turn now to the matter of how, precisely, one should go about determining whether one theory encompasses the domain of another. We have argued that the naive strategy of simply taking limits often yields results whose relevance to questions of domain subsumption is tenuous or obscure. Here, I describe a mathematical relationship between models of the same system that connects more directly and transparently to questions of domain subsumption.

Subsumption of the domain of a high-level theory T_h by a low-level theory T_l requires that every real behavior that is well-modeled by T_h is also well-modeled by T_l. More precisely, it requires that for any real system K in the domain of applicability of T_h, the specific features of K that are successfully described by T_h are represented at least as precisely by T_l. Because the manner in which a physical theory T typically represents the behavior of a system K is through some specific *model* M of that system, reduction between theories requires that T_l provide a model of K's behavior in all cases where T_h does. Moreover, T_l's model of K, M_l, must represent the behavior of those features of K that are accurately described by T_h's model of K, M_h, at least as precisely as M_h does. Thus, one may understand reduction between theories, $reduction_T$, as resting on a more fundamental concept of reduction between two models of a single system, $reduction_M$. By definition, $reduction_M$ of M_h to M_l with respect to physical system K requires that M_l successfully track all features of K that are accurately described by M_h at least as closely as M_h does. This ensures that any behavior of K that is well-modeled by M_h is modeled at least as accurately by M_l. If $reduction_M$ holds for every system in the domain of T_h, T_h $reduces_T$ to T_l. This local, model-based approach to reduction also allows for a natural account of partial reduction, insofar as $reduction_M$ may hold for some systems in the domain of T_h but not others. [13]

For our purposes in this discussion, it suffices to understand a model M of a system K as consisting of some state space S that represents the range of possible values for some subset of K's properties, together with some set of mathematical constraints (often in the form of dynamical equations) that restrict the allowed values and behaviors of the state, which we designate schematically by the relation $F(x) = 0$, where $x \in S$. Here, the expression

[13] Concerning the relationship between thermodynamics and statistical mechanics, it is possible that one may have to settle for partial reduction in this case since certain systems - namely, black holes - seem to exhibit thermodynamic regularities without having any underlying statistical mechanical description. Depending on precisely how one sets the bounds of a theory, black holes may count as an example of a system that is in the domain of thermodynamics but not in the domain of statistical mechanics. Thanks to Jos Uffink and Patricia Palacios for drawing my attention to this example.

$F(x)$ is not restricted to being simply a function of x, but may also include derivatives of x with respect to some parameter such as time. Most commonly, the constraint $F(x) = 0$ will specify the model's equations of motion; however, in some cases, like the Ideal Gas Law model, the model's constraints may be non-dynamical in nature. Below, we will designate a theory's model of a particular system by an ordered pair of some state space S and some constraints [14] F over that state space, $M = (S, F)$. [15]

Having cast the matter of reduction between theories in terms of reduction between two models of a single system, we now pose the question: by virtue of what sort of relation between models is a low-level model M_l able to describe those features of a system K that are well-described by a high-level model M_h, in all cases where M_h is successful? That is, how should one go about checking whether $reduction_M$ holds between two models M_l and M_h of a given system K? Here, I describe a general type of mathematical relationship between models that relates more directly to the problem of domain subsumption than does the procedure of simply taking limits, and that takes account of the partly empirical nature of reduction. This type of relationship generalizes a pattern that occurs across many successful instances of inter-model reduction in physics. While certain finer details of the mathematical relationship that underlies $reduction_M$ will depend on the specific mathematical form of the models (e.g., on whether they are deterministic or stochastic, dynamical or non-dynamical), it is possible to identify a certain general, but less detailed, pattern that extends across virtually all classes of model pair. I have described this type of relationship in detail within the specific context of reduction between dynamical systems models in [21], and in the context of reduction of deterministic to stochastic models in [22]. In the present dis-

[14]The term "constraint" here is used generically, rather than in the more specialized sense that is sometimes employed in gauge theories. A constraint over some state space S is to be understood here as any mathematical restriction on the physically allowed values or transformations of points in S.

[15]The exact nature of the relation between theories and models, and the precise manner in which one should characterize the content of a scientific theory continues to be a matter of significant controversy in the philosophy of science literature, particularly between proponents of the "syntactic" and "semantic" views of theories. Whereas proponents of the syntactic view identify a theory with a body of axioms and theorems in some particular formal language, proponents of the semantic view identify a theory with a collection of models (defined in a more formal logical sense than we have done here) obeying certain constraints of the theory. Although the emphasis on reduction between models here may seem to strongly presuppose a semantic view of theories, it is not clear that proponents of the syntactic view would deny that a given system in the domain of a given physical theory is described by some model in the specific sense of "model" we have adopted here. Inasmuch as proponents of the syntactic and semantic views can agree on the specific mathematical description that a given theory provides of a given physical system in its domain (e.g., the Hamiltonian description of the moon's center-of-mass orbit around the earth), it is not clear that the view of reduction presented here necessarily conflicts with the syntactic view of theories. For recent contributions to the debate between the syntactic and semantic views, see, for example, [13] and [15].

cussion, the primary emphasis will be on the empirical aspects of the general type of relationship that underlies *reduction*$_M$. For purposes of illustration, I will begin by briefly discussing the intuitive and common case of reduction between deterministic dynamical systems models; I will then go on to explain how this strategy can be generalized to inter-model reduction in cases where one or both models is not of this sort.

4.2.1 An Empirical Approach to *Reduction*$_M$ in the Case of Dynamical Systems

Let us assume, then, that we have two dynamical systems models $M_h = (S_h, F_h)$ and $M_l = (S_l, F_l)$ that describe the same physical system K, where the constraints $F_h(x_h) = 0$ and $F_l(x_l) = 0$ represent the equations of motion of these models. The state spaces S_h and S_l may describe the same features of the system K (as occurs in cases where both Newtonian and relativistic models can be used to describe a slow-moving particle) or they may describe different features (as occurs in cases where models of both classical mechanics and quantum mechanics, or statistical mechanics and thermodynamics, can be used to describe the same system). In cases where the state spaces of the models represent different features of the system, how is it possible for M_l to describe those features of K that are represented by M_h, as *reduction*$_M$ requires? In many such cases, heuristic considerations will suggest some determination relation between the features of K represented by points in S_h and the features of K represented by points in S_l; where the relationship between the models is concerned, this determination relation is usually captured by some function $B : S_l \rightarrow S_h$ from the low-level to the high-level state space. The features of K that are represented in the high-level model by points in S_h are represented in the low-level model by the quantity $B(x_l)$, where x_l is a point in the low-level space (as we will see, typically belonging to some restricted subset of this space).

For M_h to *reduce*$_M$ to M_l, it is necessary that in all cases where the high-level state x_h successfully tracks the associated features of the system K, the low-level model's representation of these same features, $B(x_l)$, track their behavior at least as precisely. The behavior of $B(x_l)$ will be determined entirely by the behavior of x_l, which is determined by the constraints F_l of the low-level model and a choice of initial low-level state. To be more precise, reduction requires for any solution $x_h(t)$ of M_h that represents a physically realistic evolution [16] of K, there exist some solution $x_l(t)$ of M_l such that $B(x_l(t))$

[16] In general, not all solutions of an empirically successful model will represent physically realistic evolutions of a system. For example, while the Newtonian model of a given projectile (here, our system K) may be empirically successful within a certain domain, there exist solutions to the equations of motion of this model that do not represent physically realistic evolutions of the projectile - for example, solutions in which the projectile achieves speeds greater than the speed of light. In general, a model of a given system provides a realistic representation of that system only for some restricted subset of states in the state space

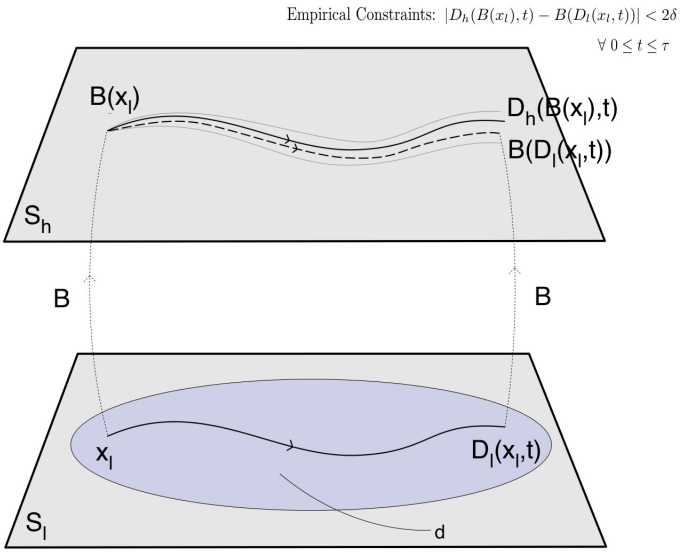

Figure 2: The local, empirical, model-based (L.E.M.) strategy for reduction rests on a certain direct but empirically constrained mathematical relationship between two models of the same physical system K. This relationship is illustrated here for the specific case where both models are deterministic dynamical systems. Given a "bridge function" B postulated on the basis of heuristic considerations, it requires that for any physically realistic high-level solution $D_h(x_h, t)$, $x_h = B(x_l)$ for some $x_l \in S_l$, and $D_h(B(x_l), t) \approx B(D_l(x_l, t))$ on at least the timescale τ for which $D_h(B(x_l), t)$ continues to represent a physically realistic evolution of K. In many cases, it is possible to show that this condition holds by showing that $B(x_l)$ approximately satisfies the high-level equations of motion for x_l in some domain $d \subset S_l$, and that the domain d is preserved by the low-level dynamics over the timescale τ. An upper bound of 2δ on the deviation between the trajectories $D_h(B(x_l), t)$ and $B(D_l(x_l, t))$ (designated by two faint outer lines in S_h) is determined by the least known upper bound δ on the margin of error within which the high-level model tracks the relevant features of the system in question. The factor 2 multiplying δ arises from the requirement that the induced trajectory $B(D_l(x_l, t))$ describe the system's behavior at least as accurately as the high-level solution $D_h(B(x_l), t)$. Since each can deviate from the system trajectory by at most δ, they can differ from each other by at most 2δ.

approximates $x_h(t)$ for all times where $x_h(t)$ continues to represent a realistic evolution of K. Writing the time-t evolution of a given high-level state x_h as $D_h(x_h, t)$ and the time-t evolution of a given low-level state x_l as $D_l(x_l, t)$, this entails that for any physically realistic high-level solution $D_h(x_h, t)$, $x_h = B(x_l)$ for some $x_l \in S_l$, and $D_h(B(x_l), t) \approx B(D_l(x_l, t))$ over intervals of t for which the high-level solution $D_h(B(x_l), t)$ continues to represent a physically realistic evolution of K.

---

and only over limited intervals of parameters such as time. Acknowledging that the phrase "physically realistic" could do with still further clarification, the example given here serves at least to illustrate the distinction between physically realistic and unrealistic solutions of models, as well as the salience of this distinction to questions of domain subsumption between models.

One widely applicable strategy for showing that this condition holds is to show that the quantity $B(x_l(t))$ approximately satisfies the high-level model's equations of motion- i.e., that $F_h(B(x_l)) \approx 0$ - for x_l in an appropriately defined subset d of S_l. "Appropriately defined" here typically includes the requirement that solutions $x_l(t)$ of M_l remain in d over the timescales where the corresponding high-level solution tracks the system's behavior; this serves to ensure that $B(x_l)$ continues to approximately satisfy the high-level equations of motion over these timescales, and, in turn, that $B(x_l(t))$ continues to approximate the corresponding high-level solution $D_h(B(x_l(0)), t)$. In a nutshell, our strategy for $reduction_M$ is to identify the appropriate bridge function B and a domain $d \subset S_l$ such that trajectories induced on S_h through B by solutions of M_l in d approximate all physically realistic solutions of M_h. Many successful applications of this strategy are detailed in [21]. This strategy is further extended to $reduction_M$ of deterministic to stochastic models in [22].

Note that the trajectory $B(x_l(t))$ need not be exactly equal to $x_h(t)$, nor does it need to approximate $x_h(t)$ for all times, in order for M_h to $reduce_M$ to M_l relative to a given system K. How closely, over what timescales, and for what initial states $x_h(0) \in S_l$ must $B(x_l(t))$ approximate $x_h(t)$ in order for $reduction_M$ to hold? Ultimately, this depends on how precisely, over what timescales, and for what initial states the model M_h succeeds at tracking the relevant features of the system K. The more coarsely and the more transiently M_h tracks these features, the more weakly is $B(x_l(t))$ required to approximate $x_h(t)$ in order to uphold reduction between the models. Thus, we can understand reduction between models as resting on a certain type of direct, formal mathematical relationship between the models that is parameterized by empirically determined bounds on the timescale, margin of error, and set of initial states for which M_h succeeds at tracking the relevant features of K. Note that the precise bounds within which the high- and low-level models are required to dovetail in this sense will vary depending on the specific system K; thus, in a certain important sense, this type of reduction is local. For a given pair of theories, the general results and patterns of proof that underwrite $reduction_M$ across different systems will tend to be quite similar, while the specific empirical margins within which dovetailing is required to hold may vary substantially across these systems - as, say between the parameters that constrain reduction between classical and quantum models of an alpha particle and those that constrain reduction between the classical and quantum models of a baseball. The implications of this locality for more general philosophical problems relating to reduction, such as multiple realization, are further discussed in [21].

4.2.2 A More General Strategy for $Reduction_M$

It is possible to generalize the strategy for $reduction_M$ described in the previous section well beyond the set of cases in which both models are deterministic dynamical systems. The strategy is to identify the correct "bridge function" $B : S_l \to S_h$ from the low-level to the high-level state space and a domain

$d \subset S_l$ such that $B(x_l)$ approximately satisfies the constraints of the high-level model - that is, such that

$$F_h(B(x_l)) \approx 0 \text{ for all } x_l \in d. \tag{1}$$

Moreover, it must be the case that low-level solutions $x_l(s)$ with $x_l(0) \in d$ remain in d over intervals of s for which physically realistic solutions $x_h(s)$ of M_h continue to track the behavior of the relevant features of K. [17] As above, the precise strength of the dovetailing between the high- and low-level model required for reduction will depend on the empirical fact of how precisely and how robustly the high-level model describes these features of the system itself. In cases where one of the models is stochastic, it will be necessary to add certain further qualifications, such as the requirement that the above condition hold with very high likelihood. Furthermore, it is likely that this general strategy can also be applied to cases in which one or both of the models in question is not a dynamical system - for example, those involving the Ideal Gas Law model or models of spacetime in general relativity.

It is worth highlighting that this local, empirical, model-based (L.E.M.) strategy for reduction bears a certain kinship to the non-mathematical requirements for Nagel-Schaffner reduction in its use bridge functions, which serve an analogous role to the bridge principles of N.S. reduction. Like bridge principles, bridge functions serve to provide a sort of translation between distinct, and possibly alien, theoretical frameworks. However, the L.E.M. approach is much more mathematically precise than N.S. reduction, and far more explicit on points where N.S. reduction is silent or non-committal. While the feature of locality has been emphasized in some recent formulations of N.S. reduction (for example, [9]), the empirical and model-based aspects of the L.E.M. approach are seldom mentioned in the context of the N.S. approach. N.S. reduction is formulated fundamentally as a relationship between theories rather than between models. Moreover, in its emphasis on logical, derivational relations between theories, N.S. reduction can easily be construed as a formal rather than an empirical approach to reduction.

5 Conclusion

I have argued that if we interpret reduction as domain subsumption, reduction between two theories in physics depends not only on the abstract formal relationship between the mathematical structures of these theories, but also on the relationship that these theories bear to the actual systems that they describe. I have criticized a tendency to approach questions of reduction - which is often

[17] In many cases, it is necessary to supplement the bridge function, which maps between the state spaces of the models, with certain auxiliary "bridge rules" that relate certain fixed, non-state-related parameters of the high-level model (e.g., parameters in the model's Hamiltonian) to corresponding quantities in the low-level model. See [19] for more detailed discussion of this point.

taken to include the requirement of domain subsumption - in a purely formal manner, as though it were solely a question of extracting the mathematical structures of one theory from those of another in some formal limit, without taking sufficient care to give the derived results a clear interpretation in the context of actual physical behaviors. Instead, I have described an alternative type of relationship between theories in physics that more transparently underwrites the subsumption of one theory's domain by another and that clearly reflects the empirical nature of reduction. This relationship, which is based on a more fundamental type of relation between two models of a single fixed system, captures and extends a pattern that arises in many known, successful inter-model reductions in physics. The more moderate but highly non-trivial concept of reduction considered here remains viable in the context of many inter-theory relations in physics, even if it has not yet been rigorously proven to hold in some of these cases. It is also possible that the distinction between formal and empirical approaches to reduction may prove relevant to reductions involving sciences outside of physics, which are typically far less mathematical. However, that is a topic for another discussion.

Acknowledgements

Thanks to Jos Uffink, Michel Janssen, Alan Love, David Wallace, and Simon Saunders for their helpful feedback on various stages of this work.

References

[1] Guido Bacciagaluppi. The role of decoherence in quantum mechanics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy.* Winter 2012 edition, 2012.

[2] R. Batterman. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence.* Oxford University Press, 2002.

[3] Robert Batterman. Reduction and multiple realizability. *http://philsci-archive.pitt.edu/11190/*, 2014.

[4] M. Berry. Asymptotics, singularities and the reduction of theories. In Brian Skyrms Dag Prawitz and Dag Westerståhl, editors, *Logic, Methodology and Philosophy of Science, IX: Proceedings of the Ninth International Congress of Logic, Methodology and Philosophy of Science, Uppsala, Sweden, August 7–14, 1991 (Studies in Logic and Foundations of Mathematics: Volume 134)*, volume 134, pages 597–607, 1994.

[5] MV Berry. Semiclassical mechanics of regular and irregular motion. *Les Houches lecture series*, 36:171–271, 1983.

[6] Jeremy Butterfield. Less is different: emergence and reduction reconciled. *Foundations of Physics*, 41(6):1065–1135, 2011.

[7] N. Cartwright. *The dappled world: A study of the boundaries of science.* Cambridge University Press, 1999.

[8] Joy Christian. Exactly soluble sector of quantum gravity. *Physical Review D*, 56(8):4844, 1997.

[9] F. Dizadji-Bahmani, R. Frigg, and S. Hartmann. Who's afraid of Nagelian reduction? *Erkenntnis*, 73(3):393–412, 2010.

[10] Michael J Duff, Lev B Okun, and Gabriele Veneziano. Trialogue on the number of fundamental constants. *Journal of High Energy Physics*, 2002(03):023, 2002.

[11] John Dupre. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science.* Harvard University Press, 1993.

[12] Jerry Fodor. Special science: Or the disunity of science as a working hypothesis. *Synthese*, 28:97–115, 1974.

[13] R. Frigg. Scientific representation and the semantic view of theories. *Theoria. Revista de Teoría, Historia y Fundamentos de la Ciencia*, 21(1):49–65, 2009.

[14] George Gamow, Dmitrii Ivanenko, and Lev Landau. World constants and limiting transition. *Physics of Atomic Nuclei*, 65(7):1373–1375, 2002.

[15] H. Halvorson. The semantic view, if plausible, is syntactic. *Philosophy of Science*, 80(3), 2013.

[16] John G Kemeny and Paul Oppenheim. On reduction. *Philosophical Studies*, 7(1):6–19, 1956.

[17] T. Nickles. Two concepts of intertheoretic reduction. *The Journal of Philosophy*, pages 181–201, 1973.

[18] Hilary Putnam. Language and reality. *Mind, language and reality*, 2:272–290, 1975.

[19] Joshua Rosaler. *Inter-theory Relations in Physics: Case Studies from Quantum Mechanics and Quantum Field Theory.* PhD thesis, University of Oxford, 2013.

[20] Joshua Rosaler. Formal vs empirical approaches to quantum-classical reduction. *Topoi*, 34(2):325–338, 2015.

[21] Joshua Rosaler. Local reduction in physics. *Studies in History and Philosophy of Modern Physics*, 50:54–69, 2015.

[22] Joshua Rosaler. Interpretation neutrality in the classical domain of quantum theory. *Studies in History and Philosophy of Modern Physics*, forthcoming, 2015.

[23] K.F. Schaffner. Ernest Nagel and reduction. *Journal of Philosophy*, 109(8/9), 2012.

[24] Elliott Sober. The multiple realizability argument against reductionism. *Philosophy of science*, pages 542–564, 1999.

[25] John Stachel. Development of the concepts of space, time and space-time from newton to einstein. *One Hundred Years of Relativity*, page 1, 2005.