September 2013

Structures in Real Theory Application: A Study in Feasible Epistemology

Robert H. C. Moir
The University of Western Ontario

Supervisor
Robert W. Batterman
The University of Western Ontario

Joint Supervisor
John L. Bell
The University of Western Ontario

Graduate Program in Philosophy

A thesis submitted in partial fulfillment of the requirements for the degree in Doctor of Philosophy

Structures in Real Theory Application: A Study in Feasible Epistemology

(Thesis Format: Monograph)

by

Robert H. C. Moir

Graduate Program in Philosophy

A thesis submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

The School of Graduate and Postdoctoral Studies
The University of Western Ontario
London, Ontario, Canada

Abstract

This thesis considers the following problem: What methods should the epistemology of science use to gain insight into the structure and behaviour of scientific knowledge and method in actual scientific practice? After arguing that the elucidation of epistemological and methodological phenomena in science requires a method that is rooted in formal methods, I consider two alternative methods for epistemology of science. One approach is the classical approaches of the syntactic and semantic views of theories. I show that typical approaches of this sort are inadequate and inaccurate in their representation of scientific knowledge by showing how they fail to account for and misrepresent important epistemological structure and behaviour in science. The other method for epistemology of science I consider is modeled on the methods used to construct valid models of natural phenomena in applied mathematics. This new epistemological method is itself a modeling method that is developed through the detailed consideration of two main examples of theory application in science: double pendulum systems and the modeling of near-Earth objects to compute probability of future Earth impact. I show that not only does this new method accurately represent actual methods used to apply theories in applied mathematics, it also reveals interesting structural and behavioural patterns in the application process and gives insight into what underlies the stability of methods of application. I therefore conclude that for epistemology of science to develop fully as a scientific discipline it must use methods from applied mathematics, not only methods from pure mathematics and metamathematics as traditional formal epistemology of science has done.

Keywords: mathematical modeling, computation, scientific data, epistemology, methodology, semantic view, syntactic view, feasible epistemology

Acknowledgements

The local origin of this project was work done during my eight month stay in Pittsburgh as a Visiting Scholar at the University of Pittsbirgh. The first four months I shared with Nic Fillion provided the inspiration, with many stimulating conversations along with our attending the fine foundations of mathematics lectures given by Wilfried Sieg at Carnegie Mellon University. During the entire period Bob Batterman played an important role in the shaping of the project and helping to develop clearer expression of my ideas. Conversations with Mark Wilson and Julia Bursten were also helpful for conceptualizing the project in this period. After my return to London, many conversations with Lori Kantymir and Emerson Doyle were essential for the clarification and development of ideas both before and during my writing the dissertation. Paul Wiegert was also a great help in getting oriented in contemporary computational approaches to modeling near-Earth objects.

The global origin of this project extends back to my time in Montreal, since which time my parents have been a constant source of support for my intellectual pursuits, no matter how wandering they have been. I simply would not have reached the point of completion of this project without their enormously generous care and support. During the same period a small group of people have been a constant source of support for the underlying ideas: Sona Ghosh, Emerson Doyle, Nic Fillion and Lori Kantymir. Nic's influence permeates this work. Our intellectual exchanges began shortly after we met when we quickly found that we shared a very similar perspective on the relationship between mathematics and phenomena, but deepened significantly following our encounter with backward error analysis in the graduate Introduction to Numerical Methods course taught by Rob Corless. Rob's influence on this work is similar in extent, from the idea of a "feasible epistemology" coming out of a conversation at the Grad Club one evening, to all I have learned from him about the epistemology of applied mathematics. The influence on my thinking coming from Bob Batterman has been great, as has his support and openness to allow me to find my own way. I am indebted to him for his scientific style of thinking about phenomena and his challenging my inclinations toward more formal representations. Possibly the deepest influence on my thought, however, has come from my long professional and personal relationship with John Bell. It is through his influence that I even came to

realize my case of the Platonistic measles, but more significant has been his constantly challenging me to walk that fine line between the ideal and the actual. Although it may not appear such on the surface, the ideas expressed and developed in this work are influenced by his work, thinking and conversations more than any other.

At a more practical level, a number of people have contributed to making the completion of the dissertation possible. Comments from and conversations with Bob Batterman, Chris Smeenk and Emerson Doyle over the last few months have been invaluable. As have comments on drafts of the dissertation itself from Bob Batterman, Bill Harper, Chris Smeenk, and Rob Corless. The editing work that Abby Egerter has skillfully and seemingly effortlessly contributed, even under very strict time constraints, has also been invaluable. Finally, Lori Kantymir has provided incredible support and encouragement through the arduous process of writing this summer. Although it is my parents, Anne and Frank Moir, that have made this work possible in the overall sense, in the most immediate and direct sense it is my dear Lori that has made it feasible.

Robert Moir
August 30, 2013
London

To Anne and Frank Moir
for their invariant love and support
and for making this work feasible

Contents

List of Tables

List of Figures

List of Abbreviations, Symbols and Terminology

Abbreviations

$\mathcal{A}$-$\mathcal{G}$	algebro-geometric	80
CCD	charge coupled device(s)	169
JPL	(the) Jet Propulsion Laboratory	115
$\mathcal{L}$-$\mathcal{S}$	logico-structural	78
NEO	near-Earth object(s)	35
ODE	ordinary differential equation(s)	4
PDE	partial differential equation(s)	4
$\mathfrak{S}$-view	semantic view	78
$\mathscr{L}$-view	syntactic view	78

Symbols

α	angle of weight on upper arm of a double pendulum	27
α	right ascension	125
β	angle of weight on lower arm of a double pendulum	27
β	ecliptic latitude	126
C	set of correspondence rules	47
C^n	a class of n times continuously differentiable functions	47
δ	declination	125
δ_q	backward error (defect) corresponding to constraint on quantity q	22
e	eccentricity	128
E	set of elementary sentences of a semi-interpreted language	61
ε	error tolerance	99
ϵ_q	forward error (error) corresponding to constraint on quantity q	21
$\dot{f}(t)$	first derivative of the function f with respect to time t	17
$\ddot{f}(t)$	second derivative of the function f with respect to time t	17
$\tilde{f}(t)$	numerically computed solution function to a constraint on $f(t)$	30
$\mathbf{f}$	vector field, force-to-mass function	15
F	constraint function	89

Terminology

Part I

Classical and Feasible Epistemology

Chapter 1

Applying Theories in the Real World: Feasible Epistemological Access to Natural Phenomena

1.1 Studying Scientific Method

1.1.1 The Phenomenon of Scientific Method

The objective of science, one might say, is to gain insight into the nature of certain phenomena by developing reliable descriptions, predictions, explanations and, perhaps, control of such phenomena. One of the distinctive features of scientific inquiry is the reliable character of its methods of investigation that provide these kinds of insight. Far from being monolithic, scientific method is rich, plural and variable. This is clearly seen from the enormous complexity and variability of methods of inference in contemporary science, methods that include various approaches to modeling, data handling, computation and error analysis. Indeed, the richness of methods in science shows that scientific method is a remarkable natural phenomenon in its own right, one generated from the process of doing science.

Rich variability in science is, of course, not a new phenomenon. Even in the earliest period of the scientific revolution there was no such thing as *the* scientific method. Although the scientific method of any scientist cannot be neatly summed up in a simple uniform account, there are nevertheless clear differences in approach among the thinkers of the scientific revolution. The approach Bacon advocated starts with empirical investigation, involving searching for patterns by taking stock of empirical

relationships pertaining to a particular phenomenon, followed by hypothesizing general empirical relationships that could then be tested by experiment (Bacon, 1620). Kepler's approach, *inter alia*, involved working with a range of mathematical models, aiming to find a precise fit of some model to available experimental data, in the interest of understanding and explaining patterns in observational phenomena. Galileo's approach varied depending on the subject matter (astronomy versus mechanics) and varied through the course of his life. In general, his approach involved a combination of empirically grounded mathematics, deduction and thought experiments, and direct experimental investigation (Segre, 1980). Mathematical symbols were used to represent measurable quantities, allowing the formulation of mathematical representations of particular phenomena that could be studied deductively and their consequences tested for concordance with experiment. His method involved, *inter alia*, a complex process of reasoning back and forth between theory and experiment in the search for a theory consistent with experimental results (Naylor, 1989). Later in the seventeenth century came the method devised by Newton, which used a general mathematical framework based on foundational physical and mathematical principles, which established systematic dependencies between mathematical relationships and natural phenomena, allowing the phenomena to accurately measure the parameters of the mathematical relationships that explain them (Harper, 2011).

Today, the variety of methods that fall under the umbrella of "science" is extraordinarily vast, demonstrates a wide range of methodological rigour, and is well beyond the capacity of any single individual to comprehend or employ. Among the more rigorous and robust of these methods are mathematical modeling techniques. These range from very abstract methods, such as the use of logical theories in computer science, to very concrete applications used throughout industry. Models are now used to gain insight into and control over very complex phenomena. An important example of this is *multiscale modeling*, which combines multiple models of dominant behaviour at different length and time scales, resulting in a considerably more accurate and robust representation of the system than any single model could provide. This approach finds use not only in theoretical contexts in physics, but has also been important in the design of large scale operations and in the study and design of materials Horstemeyer (2010). Since approaches like this combine a variety of different models, which can

involve different classes of mathematical problem, they are generally not feasible for purely analytic investigations. Indeed, it is the availability of high speed computing devices that has led to the recent rise in popularity of multiscale methods.

The methods used in numerical computation provide another example of the complexity and variety of scientific methods. Numerical computing has had an enormous impact on the kinds of problems that can be studied scientifically, which has led to the development of new scientific methods for solving problems efficiently by computer. In terms of specific applications, numerics have allowed modeling problems to be solved that were not feasibly solvable by any other means. Examples include weather forecasting, multiscale climate models, and protein folding. In general, the new methods that make these applications possible are methods for converting a mathematical model into a form that can be reliably solved by a computing device. There is enormous variety in the kinds of numerical methods that have been developed. Indeed, there is a wide variety of methods developed for solving any given class of mathematical problem, such as the range of single-step and multi-step methods for ordinary differential equations (ODE) and the range of finite-difference and finite element methods for partial differential equations (PDE). Combined with the use of high performance computers these methods allow applications of mathematics to problems that would otherwise be inaccessible in practice. Indeed, a major factor in the role of computers in expanding the range of accessible modeling problems is the fact that the computing capacity of microprocessors has increased exponentially since the invention of the integrated circuit in 1958. This trend that the number of transistors on integrated circuits doubles every two years, now known as Moore's Law, has remained accurate since the initial observation made by Moore *et al.* (1965).

In addition to expanding significantly the range of problems to which mathematics can be applied, numerical computation has also led to an ability to process and analyze large amounts of data. This has led to significant changes in pre-existing scientific fields, such as bioinformatics and econometrics, but it has also led to the development of new scientific disciplines that are only made possible by the ability to analyze large data sets. An example of this is the field of genomics, which gives us access to information about biological systems that would be inaccessible without the ability to process large amounts of data.

This gives a sense of the complexity and variability of methods used to gain knowledge about phenomena in contemporary science. This, together with the vast plurality of methods used in the intervening period between the scientific revolution and today, shows more clearly the sense in which scientific method is a remarkable natural phenomenon in its own right. And it is precisely this phenomenon that the epistemology, *qua* methodology, of science seeks to elucidate. Although epistemology of science tends to deal with quite general questions concerning the limitations of our knowledge and its reliability, these questions cannot always be answered in general terms. Understanding the limitations and reliability of the knowledge that is *actually accessible* to us requires understanding what makes the specific methods used in a given context both limited and reliable. Gaining this kind of epistemological insight requires detailed study and investigation of the methods used in real science. Consequently, one of the central issues considered throughout this study concerns the methods we should use as epistemologists of science in order to best gain clear insight into the complexity of method in scientific practice. I will argue that gaining clear insight into this complexity requires new epistemological methods, specifically adapted from those used in applied mathematics. To see how this need can be motivated, let us consider some of the aims and approaches of epistemology of science.

1.2 Epistemological Approaches to Scientific Knowledge

The classical approach to studying the epistemology of science has focused on addressing general issues of scientific knowledge and the relation between scientific theories and the world. In the early 20th century, through the work of the logical empiricists, this led to the discussion of general epistemology of science using the tools and concepts of mathematical logic. Although the approach favoured by the logical empiricists has largely fallen out of favour, the concepts of mathematical logic still permeate philosophical considerations of scientific knowledge and method.

A problem with the consideration of epistemology of science in terms of mathematical logic is that the methods used in actual scientific practice seem to bear little relation to the highly schematic deductive reconstructions of scientific theories and methods considered in philosophy. The incompatibility in form between the vast

range of methods used to actually gain knowledge about phenomena and the reconstructions or representations of science using logic has the effect that philosophical considerations about scientific knowledge often provide limited insight into how we actually gain knowledge about phenomena and what is actually involved in relating our theories to the world. Indeed, attention to the details of how theories are connected to the world in practice is quite often regarded to be a "merely pragmatic" concern, and hence not requiring philosophical consideration. On the contrary, I argue throughout this study that there is a great deal to be learned about the nature and limitations of our knowledge provided by scientific theories by studying the process of how theories are applied to phenomena in scientific practice.

The consideration of how knowledge is actually gained and the structure of the theory-world relation in scientific practice reflects different epistemological concerns than those of the classical concerns of epistemology of science. Rather than inquiring into the question of how theories in general are related to the world, or the structure and interpretation of particular theories, the concern is understanding the structure and behaviour of the methods used to gain actual knowledge about phenomena. One of the goals of this sort of understanding is to determine what kinds of knowledge we can be certain that scientific theories actually provide and to determine the limitations on knowledge that are imposed by the methods that are accessible for gaining knowledge. The ability to be certain about actually accessible knowledge and actually impossible knowledge has, aside from intrinsic philosophical relevance, the benefit of constraining philosophical accounts of epistemology of science. Thus, the details about how theories are actually related to the world in practice does have philosophical significance.

In general, the study of the actual process of gaining knowledge in science is a study of the structure and behaviour of scientific methods that make knowledge about phenomena accessible to us. An important part of this is understanding what information from theories is computable and how this information can be computed. But much stronger constraints are placed on knowledge that we can access in real science, since we do not simply need to be able to compute information, we need to be able to do so within a period of time such that the result is useful while also ensuring that the result is accurate. And how speed and accuracy can be optimized

depends very much on the scientific context and upon the technology available for the task at a given time. Since we cannot know *a priori* what this contextual and technological dependency will be, developing an understanding of the features that are common among different epistemic contexts requires studying the methods used to gain knowledge in science. This is a study, then, of what makes scientific knowledge *feasible*. To distinguish this pursuit from the classical concerns of epistemology of science I will refer to the study of the feasibility of scientific knowledge as *feasible epistemology*. In contrast, then, I will refer to the usual epistemological concerns with general questions concerning the relation of theories and the world as *classical epistemology*. These are not meant to be entirely distinct approaches, since actual epistemology of science typically involves some combination of classical and feasible concerns. The distinction is nevertheless helpful to distinguish between knowledge and methods, and studies of them, that are possible *in principle* and those that are possible *in actuality*.[1]

A question that arises here is: What methods should epistemologists of science use to study the feasibility of scientific knowledge? This is not a question with a ready answer, since the ability to feasibly gain knowledge about phenomena can be quite sensitive to the constraints imposed in particular epistemic contexts in real science. This means that epistemological methods to elucidate the feasibility of knowledge require the ability to reveal and explain patterns concerning what determines the limitations or reliability of certain kinds of knowledge in certain kinds of contexts. For instance, given that funding sources and limitations can influence the quality of research and the nature of the methods used, we might seek to understand patterns of such influence across different kinds of context. Thus, just to answer this question about methodology in epistemology of science, we will require detailed consideration of the methods used by scientists in real science. But let us consider some of the general features that adequate epistemological methods for studying feasible scientific knowledge will be required to have.

Since scientific knowledge is made feasible by the methods practicing scientists

[1]Note that the concept of a *feasible solution* from optimization and search algorithms, as a quantity that satisfies all of the imposed constraints, is not necessarily "feasible" in the sense intended here, since it may not actually be computable under the types of strong constraint that are imposed in practice.

use, a central concern for feasible epistemology is the methodology of science. Since science provides the model for our best kinds of methods for gaining knowledge about phenomena, we should attempt to emulate these methods in attempts to study scientific method. This is to say, we should study scientific method scientifically. Given that no scientific method appears out of a vacuum, any method of investigating a new phenomenon must attempt to adapt methods of investigation that have been useful in other contexts to use in a new context. And *scientific epistemology*, in the sense of an epistemology that studies knowledge and method using scientific methods, is no different in this regard. But the matter of a scientific methodology of science is peculiar in at least the sense that it is the study of scientific method itself. This reflexive nature of a scientific methodology of science has two important implications. From a constructive point of view, scientific methodology should use methods similar to those used in scientific practice in the interest of a scientific study of scientific method. And from an applications point of view, the general results of scientific methodology should apply to scientific methodology itself. Our concern in the present study is not this latter case of a general theory, since addressing this matter would require a broad survey of scientific practice, and our present focus is very narrow. Rather, we will be interested in addressing the first case, that of assessing the matter of what variety of scientific method the methodology of science should look to for methods that can be adapted most effectively to the purpose of studying scientific method.

The overwhelming trend in the epistemology of science has been to look to methods used in modern pure mathematics, particularly those used in its foundations, *i.e.*, formal logic, set theory and model theory, but also particular mathematical theories, such as probability theory and various theories in abstract algebra, *e.g.*, C^* algebras. Two formal approaches to the general treatment of scientific theories that proceed in this manner are the so-called syntactic and semantic views of theories, based, respectively, on logical $\mathscr{L}$anguages (*theories*) and sets of $\mathfrak{S}$tructures ($\mathfrak{models}$), each approach having a number of different variants. As was mentioned above, even outside of these particular approaches the use of the language and semantical tools of formal logic permeate the form of philosophical discourse about methods in science, *e.g.*, philosophical models of induction, confirmation and explanation. Many of these rely on a picture of the theory-world relation where application involves the logical deduction

of laws from theory, which, together with a set of conditions, allows the logical deduction of descriptions, predictions and explanations that are compared with data. Thus, the application of a theory is regarded, essentially, as a process of derivation of true consequences from theory using valid rules of inference.

Given that pure mathematics, in its modern form, is largely the study of pure abstract structure without any express connection to experience or the world, I wish to raise the question as to whether methods of investigation adapted from pure mathematics are sufficient for the task of methodology of science. Granted, one of the advantages of basing scientific methods on pure mathematics is that it allows for a kind of pristine clarity of expression and exactness of consequence, which are indeed desirable characteristics. And it is well suited to the construction of general theories, also desirable for any useful and powerful science. Furthermore, it is well motivated since mathematical methods have provided the conceptual tools necessary for our best and most widely successful scientific theories, dating back to the scientific revolution and even back to the ancients, a shining example being the remarkable scientific advances of Archimedes. What pure mathematics lacks, however, in terms of its status as an art or science of pure *abstract* structure, is the need to grapple with the practical problems involved in using mathematics to gain insight into the nature of *real* phenomena in experience and in the world. Since scientific method is an epistemological phenomenon constructed in the process of doing real science, it is not clear that the methods of pure mathematics alone are sufficient for its elucidation.

I suggest that rather than solely adapting methods of *pure* mathematics to the study scientific method, we should also look to *applied* mathematics, a suggestion that emphatically made by Wilson (2006). Given that applied mathematics provides the scientific methods needed in the actual practice of elucidating natural *phenomena*, and that scientific method is itself a phenomenon we wish to elucidate, applied mathematics seems to a natural place to look for methods that are suited to the *investigation* of scientific method. I argue that given the extraordinary complexity of method that we face as epistemologists of science, we should be looking to the tools that applied mathematics has developed for the purpose of overcoming complexity in its attempts to answer deeper and more varied questions about natural phenomena.

Some may question the idea of looking to the methods of applied mathematics

for guidance in the epistemology of science because applied mathematics has no clear foundation, except insofar as it, like the rest of mathematics, can be founded in set theory. I suggest that this thinking is wrongheaded. It is true that applied mathematics as a discipline lacks a foundation in the classical sense, but this is because it is organized according to different principles, as works like those of Holmes (2009) illustrate well. Pure mathematics relies heavily on the assurance of certainty of semantic reference and logical consequence. This is why set theory is such a successful classical foundation for mathematics and why the classical universe of sets is an essential tool in model theory, even for highly non-classical theories. Applied mathematics relies on this support from pure mathematics, but it also has a strong directed interest in the construction of empirically reliable and widely applicable (mathematical) models and methods. Generally speaking, the interest in applied mathematical practice is in tools that are effective in the description, understanding and control of real phenomena. General methods are sought, but so also are methods that are flexible, variable, stable and feasible.[2] And since understanding method in science requires the ability to answer specific questions about the form and function of methods, not only general ones, the epistemological intent of applied mathematics provides a much more powerful model to look to for insight into the structure and behaviour of scientific methods.

One of the aims of this study is to construct conceptual tools that are well suited for the investigation of real scientific method, tools that are modeled on methods used to apply mathematics to gain knowledge about natural phenomena. Given the complexity of scientific method as a phenomenon, I argue that these tools should provide a systematic means of abstraction from the full detail of scientific practice to clarify structural and behavioural features of scientific reasoning. This requires an approach that makes use of formal tools for conceptual clarity and precision. The dominant approaches in general epistemology of science that use formal tools are the syntactic and semantic approaches to theories. Consequently, we will find it useful to contrast some of these logical approaches to the methods developed in this study to show both the advantages of the conceptual tools developed here and the limitations

[2]Interestingly, when one looks at the *practice* of pure mathematics, rather than its content, one finds precisely the same sort of thing, *viz.*, flexible, variable, stable and feasible methods, in a quite different form (*cf.*, Mac Lane, 1986).

of methods based in mathematical logic. To this end, we will consider the so-called "received view" of the logical empiricists and Suppes's and van Fraassen's semantic approaches in the next chapter.

Since we seek methods that are suited to the investigation of scientific method, we must consider examples of the scientific method in order to either test extant formal tools or as a means of developing an approach that is capable of accurately capturing the structure and behaviour of scientific methods. We will begin this task in the next section through the consideration of a simple case of theory application in applied mathematics, *viz.*, mathematical modeling of simple and double physical pendula. We will see that just from the consideration of this simple case, the syntactic and semantic approaches to theories, as they are typically conceived, have significant problems and limitations from the point of view of accurately representing how theories are applied to phenomena in applied mathematics. The task of determining whether methods are suited to the investigation of real science cannot be accomplished only through the consideration of simple cases, however, since the real test of effective methods to study feasibility in science is to produce clarity from complexity. For this reason we also consider a realistic case from scientific practice, *viz.*, answering the question of whether a celestial object in an orbit close to Earth is likely to impact Earth in the future. This provides a detailed example of how applied mathematics is used to answer real-world questions.[3]

[3]These examples are not intended to be representative of the practice of applied mathematics in general. They have been chosen for the reasons just provided together with the requirement of relatively simple mathematics and relatively simple phenomena. For the purposes of initially exploring feasible epistemology, it is effective to follow the common scientific practice of starting with simple cases before proceeding to more complex ones. Future research projects, however, will either seek at the outset more representative cases or devise some method of choosing cases at random. The result of the present choices, however, is that only structural features common to all cases of mathematical modeling will be expected to generalize across applied mathematics. Other structural features will be more specific to peculiarities of the chosen cases. For this reason, the model of feasible theory application developed through the study is designed to be *modular*, in the sense that different structural features can be treated as independent from one another and adjusted or replaced to extend the model to describe cases that the model was not constructed to fit.

1.3 Understanding the Methods of Applied Mathematics

1.3.1 Using Applied Mathematics to Answer Real-World Questions

Scientific engagement with phenomena has produced extraordinary advances in our knowledge of the world. In the last one hundred or so years we have gone from thinking of the universe as fixed and filled uniformly with stars to realizing that it is a universe with an origin in a finite time in the past, which is now in accelerating expansion and full of dendritic filaments of groups of galaxies. We now see just how insignificant our tiny blue planet is, and we now recognize just how fragile it is as well. We now know with near certainty that a large asteroid or comet struck the Earth 66 million years ago, bringing an end to the Mesozoic era and three-quarters of the plant and animal species on Earth, including the dinosaurs. And though extinction events from the impacts of such large objects are thought to be very rare, impacts of the size of the Tunguska event in Russia in 1908, an impact carrying the energy of a 10-15 megaton nuclear device without the attendant radiation, are considerably more likely. Such an impact is in the range of 650-1,000 times more powerful than the bomb dropped on Hiroshima. And an even more recent reminder of our vulnerability came on February 15 of the present year, with a meteor impact, also in Russia, now called the Chelyabinsk meteor. From calculations performed by the Meteor Physics Group in the Department of Physics and Astronomy at Western University, the blast carried an energy equivalent of approximately 440 kilotons of TNT (about 30 times the energy of the Hiroshima bomb).

Recognizing the Earth's vulnerability to comets and asteroids in the solar system is one thing, but figuring out what we can do about it is another. These events provide strong motivation for an increased awareness of the massive objects in the vicinity of the Earth and an urgency to find ways to prevent future collisions that could threaten the existence of life on Earth. This, then, presents a very particular question about certain phenomena in the solar system and our means of knowing about them: How can we figure out whether a massive object is going to impact the Earth and how can we stop the impact from happening? This is an example of an urgent question that requires a quick and reliable answer. And it is just the sort of question that the methods of applied mathematics are suited to answer.

The strategy of applied mathematics in relation to a practical question such as this is to refine the question sufficiently such that it can be formulated as a set or sequence of mathematical *problems*, which can then be *solved* using mathematical techniques in order to provide a timely and reliable answer to the practical question. In this case, we first need to know what objects are out there in the vicinity of the Earth, what they are made of, and where they come from. This is because the origin of an object can tell us a great deal about its constitution, and its constitution can determine the nature of the threat that an impact poses. In addition to this we need to catalogue the objects that we know about. To do this, we need to be able to discern very precise orbits for these objects. And once we have determined precise orbits, we can then determine the probability that one of these objects will impact the Earth. Once we know that a collision is certain or highly likely at a certain point in the future, we can try to deflect the object from its current orbit onto one that will not intersect the path of the Earth. In this way, we can convert a broad question into a plurality of smaller questions that are capable of being addressed effectively. Once these questions have been precisely formulated, such as the question of how to determine the orbit of a near-Earth object, they can be formulated mathematically and then answered precisely and reliably.

The modeling of near-Earth objects provides a case of a complex logistical and mathematical problem on which the tools of applied mathematics can effectively be brought to bear. It is an example of a serious problem of concern for humanity and one that involves an enormously complex embedding in the practical life of scientists and the use of scientific methods. For this reason, it will provide an effective case study for the purposes of developing a model of theory application. This example illustrates the kind of complexity that is encountered in real life modeling, which will help to ensure that the model of theory application we construct on its basis will make a decent amount of contact with the actual experience of applied mathematicians engaged in modeling the world to answer real-world questions.

The specific parts of this problem that we will focus on are how the orbit determination problem for near-Earth objects is solved and how the impact probability for an object with a catalogued orbit is determined. Examining how this is done will take us through a thorough consideration of not only how theory and data are handled in

real science, but *computation* as well. Indeed, one of the most glaring omissions in the general philosophy of science is a full account of the theory-world relation that includes the role of computation. Accordingly, one of the central aims of this study is to determine how we can begin to understand the role that real computation plays in the entire process of application of theory. Not only will we find that computation is an essential feature of the feasible approach, we will find that it is *characteristic* of it, in the sense that feasible methods are computational in nature. More specifically, it is fast, efficient computation that provides solutions in real time that is characteristic of a feasible approach. Accordingly, we may begin to see how the lack of feasibility considerations in epistemology and the lack of detailed consideration of practical computation go hand in hand.

In order to construct a useful model of theory application it will be helpful not only to consider a complex real world example, but also to consider a simple theoretical example that still forces us to consider significant issues of feasibility. The complementary consideration of a simple model will help to clarify the nature of the computational procedures used in applied mathematics, consequently facilitating the elucidation of feasible epistemology. And we will find that the case of the ideal double pendulum evidences kinds of complexity not clearly present in the case of modeling near-Earth objects.

1.3.2 Modeling Simple Pendulums Using Newtonian Mechanics

Before considering the ideal double pendulum, we will begin by considering a simpler case. This serves both as a warm up to the double pendulum case, since we will see that even this simple case the detailed method of application is quite complex, and as a recurring example that we will use throughout the study. The simple case we begin with, then, is point-particle newtonian mechanics[4] and the construction of a model of the simple pendulum within it. We will use newtonian mechanics again when we turn to the consideration of near-Earth object modeling in chapter 3.

The basic equation we may take to define the theory of point-particle newtonian

[4]I am following the convention of Arnold (1989) in using lower case letters for names that have been converted into adjectives.

mechanics is Newton's second law,[5]

$$\mathbf{F} = m\frac{d^2\mathbf{x}}{dt^2} \tag{1.1}$$

which is taken to determine the motion of each particle in an n-particle system. The symbols in this equation have both a mathematical and a physical interpretation. Physically, $\mathbf{x}(t)$ is a function specifying the position of a given particle as a function of time t, m is the mass of that particle, and $\mathbf{F}(\mathbf{x}(t), t)$ is the total vector force acting on the particle at time t. The mathematical meaning of this equation is made more clear by representing it in a slightly different form.

A system of n bodies in three dimensional space is represented by a system of n copies of (1.2), one for each of n particles in a system. This determines a $3n$-dimensional configuration space, a point of which determines the state of the system, with the dynamics being governed by *Newton's second law*

$$\frac{d^2\mathbf{x}}{dt^2} = \mathbf{f}, \tag{1.2}$$

where $\mathbf{f}$ is a function on this $3n$ dimensional space that specifies the total force to mass ratio $\mathbf{F}_i/m_i$ for each of the n particles. We will refer to this as the *state space* of the system,[6] and we will use the same symbol $\mathbf{x}(t)$ to refer to the $3n$-dimensional position vector function. The $3n$-dimensional position vector $\mathbf{x}(t)$ then, away from collisions, lies in a space of twice differentiable functions from an interval I of the real line, picking out the time domain, into the state space. The mass m is simply some real number. And the total force-to-mass function $\mathbf{f}(\mathbf{x}(t), t)$ for the n-particle system lies in a space of continuous functions from the state space to itself, which may or may not vary with the value in I, the time.[7] Such a function from the state space to itself is called a *vector field*.[8] Physically the total force-to-mass on each

[5]There is also Newton's third law, *viz.*, $\mathbf{F}_{ij} = -\mathbf{F}_{ji}$, where $\mathbf{F}_{ij}$ is the force on particle i exerted by particle j, but we will not consider this explicitly here.

[6]This is distinguished from a *phase space*, which specifies $6n$ variables, the position state and the state of motion (velocity or momentum) of the n particles.

[7]The case where the forces do not depend explicitly on time is called the *autonomous case*, and where the forces do depend on time the *non-autonomous case*. We will be restricting attention to the autonomous case, but mathematically this implies no loss of generality, as will be explained below.

[8]This may be pictured as a field of arrows over a given space, *i.e.*, an arrow at each point

particle determines its acceleration at a position $\mathbf{x}_i$ and time t, which is what equation (1.2) says. Mathematically, the vector field (inversely scaled by the parameter m) determines the second order derivative of the position function $\mathbf{x}(t)$.

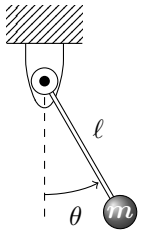

Figure 1.1: An ideal simple pendulum.

So much for the theory equation (1.2); what about methods used to construct models? Well, this requires the adding of constraints that determine the dimension of the state space (number of degrees of freedom). The simple case we consider here is the modeling a simple pendulum (see figure 1.1). We will suppose that the physical pendulum is such that the pivot has very little friction, the rod is very light and effectively rigid and the weight is small and dense. Under these conditions, we can construct an ideal model of the simple pendulum that has only one degree of freedom, specified by the angle θ that the rod makes with the vertical (see figure 1.1).[9] To do this we suppose that the pivot is frictionless, that the rod is perfectly rigid and massless, and that the weight is a point mass. We then represent these conditions as mathematical constraints.

determines the direction of the vector field at that point and its length determines the magnitude of the vector field at that point. For the circle that the ideal simple pendulum travels on, the vector field always points tangent to the circle and toward the ground. It reaches its maximum value when the pendulum rod is horizontal and the force is directed straight down and its minimum value of zero when the rod is vertical.

[9]Strictly speaking, θ is a *generalized position* variable and not a newtonian position variable. θ parameterizes a one dimensional circular manifold that defines the state space of the pendulum. The mathematical framework of newtonian mechanics uses only positions of particles in a three dimensional cartesian space. Accordingly, geometrically this model fits in the framework of *lagrangian mechanics*, which treats dynamics in terms of an n-dimensional configuration manifold parameterized by n generalized variables, together with their n time rates of change parameterizing the tangent bundle of configuration space. Nevertheless, by imposing the rigid body constraint on the rod in a cartesian newtonian setting, this still restricts the motion of the pendulum to a circle, which makes it natural to choose a parameterization of the circle that the motion is constrained to. Thus, in this way generalized variables arise naturally in newtonian mechanics as well. Furthermore, we have not used the apparatus of a lagrangian function and the euler-lagrange equations to construct the model, we have indeed used the framework of newtonian mechanics in the sense of computing the sum of vector forces on bodies and using Newton's second law to find the equation of motion. In this case, then, the angle θ can simply be regarded as a change of variable.

As an example, consider the length of the rod of the pendulum. A more detailed model would allow the length to vary depending on the applied forces, perhaps by modeling the rod as a very stiff spring. To explicitly formulate the rigid body assumption in a newtonian framework, then, we could add the constraint that the time rate of change of the length ℓ of the pendulum is zero, *viz.*, $\dot{\ell} = 0$, where the dot indicates a time derivative.

Under the applied constraints, and given that the force exerted by gravity on a mass is $m\mathbf{g}$, where $\mathbf{g}$ is a vector pointing toward the centre of the Earth and has magnitude $g = 9.81$ m/s, simple trigonometry allows one to show that the total force exerted on the weight is $\mathbf{f} = -\frac{mg}{\ell}\sin\theta$ (in a polar coordinate system (r, θ) centred at the pivot, see figure 1.2).[10] Thus, (1.2) becomes

$$\ddot{\theta} = -\omega^2 \sin\theta, \tag{1.3}$$

where the dots indicate time derivatives and $\omega = \sqrt{g/\ell}$, and the mass m has dropped out of the model, indicating the motion of the pendulum is independent of the mass of the weight.

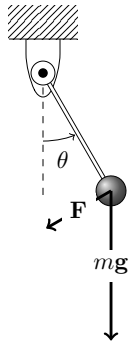

Figure 1.2: Forces on an ideal simple pendulum. The tension in the rod and the component of the gravitational force acting on the rod is not shown because these two forces cancel no matter what the mass m is because the rod is assumed to be rigid.

The equation (1.3) specifies the equation of motion of the ideal model pendulum, which when supplemented with initial conditions, will determine the form of the motion $\theta(t)$ of the model pendulum over time. Physically, this will describe the motion of the target physical pendulum for as long as the assumptions (no friction,

[10] r does not appear in the model because the constraint that the rod is rigid means that $\dot{r} = 0$, *i.e.*, that r is constant.

rigid bodies, point mass, *etc.*) are *effectively true* of the system, which occurs when the behaviour of the components of the system is not significantly distinguishable from the behaviour of their idealized counterparts.[11] The most significant factor in this case is the friction, which will quickly dissipate energy from the physical pendulum system, causing it to slow to a stop. But, for a short period of time the model pendulum will accurately describe the motion of the physical pendulum.

Mathematically, quite a lot has happened in the transition from the theory equation (1.2) to the model equation (1.3). The dimension of the state space, which could have been any finite dimension n, is fixed to be one. This implies that the functions $\mathbf{x}(t)$ and $\mathbf{f}(\mathbf{x})$ are in the function spaces $C^2(\mathbb{R} \to \mathbb{R})$ and $C(\mathbb{R} \to \mathbb{R})$, respectively. Before these constraints were applied, however, n was not determined and so they could have been in any of the spaces $C^2(\mathbb{R} \to \mathbb{R}^n)$ and $C(\mathbb{R}^n \to \mathbb{R}^n)$, respectively. Geometrically, however, the model is representing a pendulum that is located in three dimensional space $\mathbb{R}^3$. The model constraints imply that the state of the system is restricted to a circle of radius ℓ embedded in $\mathbb{R}^3$. Thus, the geometric presentation of the state of the model is as a (twice differentiable) function from an interval $I \subseteq \mathbb{R}$ into a one-dimensional circular manifold M embedded in $\mathbb{R}^3$.

There is also another significant shift from equation (1.2) to equation (1.3). In the theory equation (1.2) the symbols $\mathbf{x}$ and $\mathbf{f}$ have no distinct reference. Rather, they pick out a vast plurality of function spaces in which the interpretation of these symbols could lie ($C^2(\mathbb{R} \to \mathbb{R}^n)$ and $C(\mathbb{R}^n \to \mathbb{R}^n)$, respectively, for any n). Already, m is sufficiently meaningful because it is required to be a positive real number. Once the dimension of the state space was restricted to one, however, the symbols $\mathbf{x}$ and $\mathbf{f}$ are specified to lie in $C(\mathbb{R} \to \mathbb{R})$. This makes them sufficiently meaningful, since we know that they are functions of a specific type. And this also defines a distinct *space of models*, indexed by the possible values of $\mathbf{f}$ in $C(\mathbb{R}^n \to \mathbb{R}^n)$. The model construction process did not stop here, however, since we also specified the total force on the weight of the pendulum, which specified a particular value for $\mathbf{f}$, *viz.*, $-\frac{g}{\ell}\sin\theta$, which determines the dynamics on the state space. This now provides us with a *distinctly meaningful* equation, in the sense that it is talking about a particular

[11] Note that what counts as *significantly distinguishable* depends on tolerances for error in structure and behaviour in a given modeling context.

state space, and a specific dynamics on that state space. Moreover, we represented this mathematical form of the model in three dimensional space, in such a way that mirrors the motion of the physical pendulum. This sort of equation, which has a specific geometric interpretation (state space, vector field, with or without projection into 3D space), is what applied mathematicians often mean by a *model*. This model can then be studied subject to a variety of different initial and or boundary conditions in order to understand its behaviour.

Determining how the model will behave under particular circumstances requires imposing particular *conditions*: initial conditions or boundary conditions. This requires finding the *solution* to the model equation (1.3) under specific conditions. Suppose that here we are interested in starting the pendulum in motion from a state of rest. In this case we impose the initial conditions

$$\theta(t_0) = \theta_0, \qquad \dot{\theta}(t_0) = 0, \tag{1.4}$$

where t_0 is the initial time, the initial angle is θ_0 and the initial angular velocity is zero. These conditions together with the model equation (1.3) define a particular mathematical *problem*, called an *initial value problem*. The solution to this problem will tell us how the model behaves under the imposed conditions.

This model is not entirely straightforward to solve as it stands, however, since the vector field $\mathbf{f}$ is a nonlinear function of θ.[12] If we are content to consider only the behaviour for small angles, then we can use the small angle approximation for $\sin\theta$, *viz.*, θ, which simplifies (1.3) to

$$\ddot{\theta} = -\omega^2\theta, \tag{1.5}$$

which is just the equation for simple harmonic motion, which has under the conditions (1.4) the solution

$$\theta(t) = \theta_0 \cos(\omega t).$$

This approximative technique, called *linearization*, is not strictly deductive. It preserves the structure of the equation, however, changing only the form of the vector

[12]It does have a general solution, however, and one that can be fairly straightforwardly found using Jacobi elliptic functions (see, *e.g.*, Lawden, 1989). More typical cases of nonlinear models, however, are not so amenable to closed-form analytic solution, if they can be solved analytically at all.

field, buying us an equation that is easily solved analytically at the expense of its only being valid under conditions that ensure small angles.

Models are used for a variety of different purposes in applied mathematics, including description, prediction, explanation and control. Let us suppose that in this case we are interested in the classic case of prediction. This simple model makes a number of predictions. It tells us that for small angles and over a short time interval the motion should be described by a cosine function. It also tells us that the frequency of oscillation $\omega = \sqrt{g/\ell}$ is inversely proportional to the length ℓ of the pendulum rod in a particular way, *viz.*, that quadrupling the length of the pendulum would reduce the frequency by a factor of $\frac{1}{2}$. Or, equivalently, that the period $T = 2\pi/\omega$, *i.e.*, the time of one oscillation, of the pendulum is directly proportional to the the square root of the length. To test these consequences requires observation and measurement.

The square root dependency of the period on the rod length could be studied simply by constructing a series of pendulums with a low friction pivot, rods of different lengths made of a light, rigid material, and a small dense weight, much heavier than the rods. Using a measuring tape and a stop watch to measure the lengths of the rods and the time of one period, respectively, it would be possible to determine if the $\ell^{1/2}$ dependency is replicated by real pendulums.[13] A strategy for measurement here could be to run several trials of period measurement for each rod length, each trial being a time average of several oscillations, and then averaging the results of the trials for each rod. The result would be a measurement of the (short time) period of each of the pendulums. To see if the square root dependency obtains, several methods could be used. An example would be to compute the logarithms of the period and length measurements, plotting these on a log-log plot of the period versus the length, performing a linear regression analysis and then seeing if the slope of the regression line is equal to $\frac{1}{2}$ within experimental error. Determining the experimental error is another matter, which would require estimating the uncertainty of the devices

[13]The $\ell^{1/2}$ dependency could actually be tested without the use of a time measuring device, by using a measuring tape to measure the lengths of the rods and using rod lengths that differ by a factor of 4. It would then be possible to determine if the $\ell^{1/2}$ dependency is replicated by real pendulums by determining whether there are two oscillations of a given pendulum for every one of a pendulum with a rod four times as long. This method avoids the potential circularity of using a time measuring device that itself uses vibrations of a quartz crystal, an essential equivalent of a pendulum. For our purposes here, however, we will require some time measurement device, so it is convenient to work with the simpler case of a common time measuring device.

themselves, calculating the standard error of the trial measurements, and then tracing the propagation of error through the calculation done on these quantities. At the end of this process we would be in a position to determine whether the square root dependency prediction of the model is correct.

The measurement devices for the procedure just described (a tape measure and a stop watch) would not be sufficient to determine whether the form of the function describing the evolution of the angle θ over time is, to a high degree of approximation, a cosine function for small angles. First of all, this means of measuring time will probably not be sufficiently accurate to precisely determine the form of the function. But we also need a device to measure the angle accurately at particular times. For example, this could be done with a digital camera and a strobe light or with a video camera with a precise time stamp. However this is accomplished, the result will be a number of pairs of measurements of angle and time. These could be compared point-wise with the prediction of the models by calculating $\theta(t)$ for each measurement time t_i. If the difference is small, then this is one measure of whether the pendulum and model agree. Or the measurements could be interpolated to obtain a twice continuously differentiable function that minimizes the deviation from the measured angles.[14] This would provide us with a differentiable function $\tilde{\theta}(t)$, which is a *model* of the raw angle-time data.

There are two basic ways that we could assess whether the model and the measurements agree. One approach is to compute the difference between the two functions $\theta(t)$ and $\tilde{\theta}(t)$. This difference

$$\epsilon_\theta(t) = \theta(t) - \tilde{\theta}(t)$$

is called the *forward error*. If for any given time, $\epsilon_\theta(t)$ is small, then we know that the form of the motion of the physical pendulum is approximately a cosine function. This is a continuous version of the discrete calculation mentioned above where we compute $\theta(t)$ from the model solution at each measurement time t_i and calculate the

[14] This can be done, for example, using a combination of forward differences and Hermite interpolation (for an example of this approach to interpolation see Moir, 2010). Forward differences could be used to estimate the angular velocity at each point in time, and Hermite interpolation could then give a C^2 interpolant that matches the measured angles and the estimated angular velocities at each time.

difference from the measured value θ_i. There is another way, however, that we could assess agreement between the model and the data.

A second approach is to compute the amount by which the data model function $\tilde{\theta}(t)$ fails to satisfy the linearized model equation (1.5), or for that matter the original model equation (1.3). Since $\tilde{\theta}(t)$ is twice differentiable by construction, we can plug it in for θ in either of these model equations. If we did this in equation (1.5) then we would find that it does not satisfy the equation exactly, but leaves a residual or *defect* $\delta_\theta(t)$, *viz.*,

$$\frac{d^2\tilde{\theta}}{dt^2} = -\omega\tilde{\theta} + \delta_\theta(t). \tag{1.6}$$

This defect $\delta_\theta(t)$ is called the *backward error*. It is an assessment of how closely the data function $\tilde{\theta}(t)$ satisfies the constraints and conditions of the model. This allows us to see that the backward error can be interpreted as a *perturbation* (small variation) of the vector field, a perturbation which causes the vector field to vary as a function of time. Since the vector field specifies the dynamical behaviour of the target physical system in the context of the modeling framework, *a small backward error means that the dynamical behaviour predicted by the model and that evidenced by the data are very close*.

Rather than being useful for assessing agreement of the model and data functions, this approach is useful for assessing whether we have good evidence that the model as a whole is valid. Since we modified the vector field of the original model by linearizing it, and since we used strong idealizing assumptions to construct the original model, there is no reason to expect the dynamics of a physical pendulum to be exactly specified by either (1.3) or (1.5). Accordingly, the calculation of the backward error allows one to *measure* the dynamics of the physical pendulum *within the modeling framework of a one-dimensional newtonian model*. This is to say that perturbed vector field in equation (1.6), which has been measured, picks out a model in the 2D $(\theta, \dot{\theta})$ framework for which the data model $\tilde{\theta}(t)$ is the exact solution. We may thus regard the perturbed dynamics as the *effective dynamics* of the physical pendulum in the two-dimensional modeling framework. And if the backward error is small, then the effective 2D dynamics is modeled well by the ideal model we have constructed. This provides a considerable epistemic advantage over the forward error approach, since this allows us to have a quantity that is easily calculated from the data inform

us directly as to how well we have modeled the dynamics of the physical pendulum.[15]

There is one other epistemic advantage to considering the backward error that should be mentioned here. We see that the data function $\tilde{\theta}(t)$ does not satisfy either of the model equations (1.3) or (1.5) exactly, making it an approximate solution to these models. Regarding the solution as an approximate solution to one of these models would make it natural to judge the quality of the solution in terms of the size of the forward error, *i.e.*, how close an approximation to the exact model solution the data solution is. But we may notice that the data function satisfies the modified model equation (1.6) *exactly*. Thus, we can also regard the data model function as the *exact solution to a modified model*. Regarding an approximate solution to a problem as the exact solution of a modified problem is the essence of the approach taken by *backward error analysis*, which we will discuss in more detail in chapter 5. Thus, rather than regarding the constructed model as exact and the measured data solution as approximate, we can regard the modeling framework as enabling us to measure the closest model for which the data solution is the exact solution.[16]

This now gives us a basic picture of the application process in an exceedingly simple case. There is little in this case that goes outside of what is typically considered in philosophical models of theory application. Other than perhaps the non-deductive linearization step, the only epistemically significant feature that is not considered is the analysis of backward error. Indeed, I will argue in the next chapter that typical approaches on the syntactic and semantic views of theories are unable to accommodate

[15]It is worth noting that this approach allows the defect to *carry information* about how well we have modeled the physical system. It is a subject of future work to determine whether there are deeper epistemic gains to this approach than are evidenced here. For evidence that such gains may be possible, consider that in the context of the analysis of computational error introduced by numerical methods, specific computable perturbations of the model dynamics can explain the structure of the defect, as is shown by Corless (1994*a*), work extended by Moir (2010). The explanation is accomplished by showing that certain higher order perturbations of the original model result in modified models with much smaller defects. If, in the case of data handling, the defect can be explained in terms of computable perturbations that have a specific physical interpretation, then it could be possible to detect physical sources of error that explain the deviations from the constructed model. If this can be accomplished, then the epistemic gains of using backward error analysis in data handling would appear to be strongly analogous to the epistemic gains in Newton's methodology that are afforded by treating the model as exact and allowing the deviations from it to carry information about the phenomenon, as discussed by Harper (2011).

[16]In this case, this requires increasing the dimension of the state space by one, since the vector field in (1.6) is now explicitly dependent on time. Thus, the modeling framework must be modified. But, in the case of non-autonomous models, the data function will be the exact solution of a model in the same modeling framework.

methods of this kind, because they rely on both specific syntactic manipulations and their corresponding semantic variations. Both model reduction methods like linearization and backward error analysis will figure largely in the model of theory application we obtain from this study. It is interesting, however, that even the detailed consideration of one of the simplest models in applied mathematics already goes outside of what extant formal classical approaches to the theory-world relation can accommodate. There is another important feature, however, that is not adequately represented on classical views in general. This feature is the use of numerical methods, which are often *required* for successful theory application; for example, in cases where nonlinear models force their use to compute solutions to a model.

1.3.3 Modeling Double Pendulums Using Hamiltonian Mechanics

The double pendulum is a simple extension of the simple pendulum, obtained by adding another rod and weight to a simple pendulum (see figure 3.3). This very simple change to the physical change results in wildly different behaviour than that exhibited by the simple pendulum, since the system goes from very regular behaviour to being *chaotic*. The proper description of chaotic behaviour requires representing the system in a higher-dimensional system space than the state space of newtonian mechanics. Rather than simply representing the configuration of n bodies, this sort of system space represents the configuration and the *state of motion* of the n bodies.

The two general ways of doing this both use *generalized coordinates* to describe the motion of the bodies. Both start with a generalized position coordinate $\mathbf{q}$ for a body, an example of which is the angle θ used to parameterize the motion of the simple pendulum. The first approach represents the motion in terms of the time rate of change $\dot{\mathbf{q}}$, or the *generalized velocity*, of the body. The coordinates $(\mathbf{q}, \dot{\mathbf{q}})$ then specify the spatial configuration and state of motion of the body. This is the mathematical framework used by *lagrangian mechanics*. The configuration of a system with n degrees of freedom is represented by a point in an n-dimensional *configuration space*, specified by an n-vector function $\mathbf{q}(t)$. The state of motion is then represented by the time derivative $\dot{\mathbf{q}}(t)$ of the configuration function. The total space representing the configuration and motion of the system is the $2n$-dimensional tangent bundle of configuration space. The *system space* of lagrangian mechanics is the n-dimensional

configuration space, but the system space determines a distinct larger space, the tangent bundle, that specifies the configuration and motion of the system.

The other way to represent the configuration and motion of a system involves augmenting the generalized coordinate $\mathbf{q}$ not with a generalized velocity but with a *generalized momentum* $\mathbf{p}$. The coordinates $(\mathbf{q}, \mathbf{p})$ then specify the spatial configuration and state of motion of the body. This is the mathematical framework used by *hamiltonian mechanics*. The entire state of motion of a system with n spatial degrees of freedom is represented by a point in a $2n$-dimensional *phase space*, specified by an n-vector function $\mathbf{q}(t)$ for the generalized position and an n-vector function $\mathbf{p}(t)$ for the generalized momentum conjugate to the generalized position $\mathbf{q}$. The phase space of a system has an important geometrical property that arises from a kind of symmetry between the generalized position and momentum coordinates, a symmetry imposed by hamiltonian dynamics.

To see how this arises let us first consider how the dynamics of a system is specified in lagrangian mechanics. This is done by specifying the *lagrangian* function $L(\mathbf{q}, \dot{\mathbf{q}})$,[17] which is equal to the difference in kinetic energy T and potential energy V of the system. In the framework of lagrangian mechanics this function L determines the equation of motion of the system from a principle that the physical action of a system evolving between two points in configuration space is locally minimal. This principle determines the *Euler-Lagrange equations* of lagrangian mechanics, into which the lagrangian function L for a particular system can be substituted to find the (differential) equation of motion of the system. This approach is useful, particularly for the modeling of real world physical systems.

Another approach, which is more useful for theoretical purposes, is the hamiltonian approach, a special case of which is obtained from lagrangian mechanics. Using a mathematical transformation, a legendre transformation,[18] the lagrangian function can be used to define another function $H(\mathbf{q}, \mathbf{p})$, the *hamiltonian*. The combination of this legendre transformation and the Euler-Lagrange equations *defines* the generalized momenta $\mathbf{p}$ in terms of partial derivatives of the lagrangian $L(\mathbf{q}, \dot{\mathbf{q}})$. The

[17] We are considering here the autonomous case where the lagrangian does not have an explicit dependency on time.

[18] This same mathematical transformation, which essentially involves representing the change in a quantity in terms of changes in the other quantities it depends on, is very important in the definition of various quantities in thermodynamics.

dynamics implied by the Euler-Lagrange equations in terms of the hamiltonian is then specified by *Hamilton's equations*:

$$\dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}}, \qquad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}}. \tag{1.7}$$

This transformation from the lagrangian framework to the hamiltonian framework replaces a system of n *second order* differential equations, *i.e.*, equations with second order time derivatives, on an n-dimensional configuration space with a system of $2n$ *first order* differential equations on a $2n$-dimensional phase space. And the transformation also exposes a distinct kind of symmetry between the position and momentum coordinates, reflected in Hamilton's equations in terms of the presence or absence of a minus sign in the two equations above.

This symmetry between position and momentum variables in hamiltonian mechanics implies a distinct kind of geometric structure on phase space. Mathematically this is a nondegenerate differential 2-form, and it characterizes the kind of symmetry between $\mathbf{q}$ and $\mathbf{p}$ represented in Hamilton's equations (1.7). This kind of structure on a manifold is called a *symplectic structure* and a manifold with this kind of structure is called a *symplectic manifold*.

The significance of symplectic structure then is that it *encodes* the structure of hamiltonian dynamics. To clarify the sense in which this is the case, consider transformations of the phase space that *preserve* the symplectic structure, which we may think of as preserving the symmetry between $\mathbf{q}$ and $\mathbf{p}$ encoded in Hamilton's equations. Since such transformations respect hamiltonian dynamics, they describe the potential *dynamical evolution* of a hamiltonian system, *i.e.*, how the state of a hamiltonian dynamical system could evolve over time. Mathematically such transformations are called *symplectic diffeomorphisms*, or *symplectomorphisms*. The set of all such transformations forms a mathematical group, which acts on phase space. Each one-parameter subgroup of symplectomorphisms, picked out by a symplectomorphism φ_t, determines a particular hamiltonian evolution on phase space. The connection with Hamilton's equations (1.7), then, is that the specification of a hamiltonian H in (1.7) *fixes* a particular symplectomorphism φ_t on phase space, determining the specific dynamics of the system being modeled.

The importance of this geometric picture for understanding a hamiltonian model of a system as a mathematical *problem* is the following. For newtonian mechanics, specifying the mass m of a body and the forces on it determines a *vector field* on the n dimensional state space of the newtonian model. Similarly, specifying the hamiltonian H of a system determines a vector field on the $2n$-dimensional phase space of a hamiltonian model. Thus, the form of the hamiltonian equations can be cast in a very similar form to that of Newton's second law (1.2), where the "position vector" is $\mathbf{x} = (\mathbf{q}, \mathbf{p})$ and the "vector field" is $\mathbf{h} = (\partial H/\partial \mathbf{p}, -\partial H/\partial \mathbf{q})$, so that Hamilton's equations (1.7) can be written as

$$\frac{d\mathbf{x}}{dt} = \mathbf{h}, \tag{1.8}$$

where we have replaced a system of n *second* order equations (1.2) with a vector field $\mathbf{f}$ on n-dimensional state space with a system of $2n$ *first* order equations with a vector field $\mathbf{h}$ on a $2n$-dimensional phase space. Now, if we impose conditions that determine a unique *particular* solution to this problem, such as the position $\mathbf{q}_0$ and momentum $\mathbf{p}_0$ of the system at some time t_0, then the solution is a particular trajectory through phase space. But, if we only fix the *dynamics* without specifying particular initial conditions, then the *general* solution to this problem is a particular *symplectomorphism* φ_t acting on the $2n$-dimensional phase space, called the *phase flow* of the dynamical system. The particular trajectory obtained before is then recovered simply as $\varphi_t(\mathbf{q}_0, \mathbf{p}_0)$. Thus, the geometry of hamiltonian mechanics makes a clear connection between the vector field $\mathbf{h}$ that determines the model from the specified hamiltonian H and the phase flow φ_t that constitutes its general solution.

Now that we have clearly defined the geometric framework of hamiltonian mechanics and its physical interpetation, we may now briefly consider the treatment of the double pendulum in this framework. Using similar kinds of simplifying assumptions as we used to construct the ideal simple pendulum model (1.3) we can construct an ideal model of the double pendulum. In this case, rather than the angle θ and angular velocity $\dot{\theta}$ for the simple pendulum, we use the two angles α and β and their angular velocities $\dot{\alpha}$ and $\dot{\beta}$ to describe the motion (see figure 3.3). We will also work with the simplifying assumption that both rods are the same length ℓ and that both weights have the same mass m. It is then a straightforward matter to compute the lagrangian $L(\alpha, \beta, \dot{\alpha}, \dot{\beta})$ for the double pendulum in these variables. The legendre

transformation of this then allows us to compute the hamiltonian and the generalized angular momenta l_α and l_β for the two weights. The specification of the hamiltonian of the system then fixes the hamiltonian vector field $\mathbf{h}$, a hamiltonian dynamics, on the four dimensional phase space $(\alpha, \beta, l_\alpha, l_\beta)$ (*cf.*, equations (1.7) and (1.8))

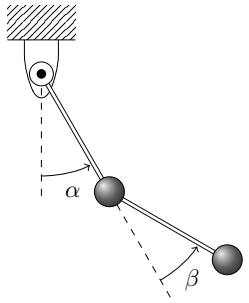

Figure 1.3: An ideal four-dimensional double pendulum system.

This specifies the hamiltonian model for the double pendulum, which more explicitly is (*cf.*, equations (1.7) and (1.8))

$$\left(\dot{\alpha}\dot{\beta}\dot{l_\alpha}\dot{l_\beta}\right) = \left(\partial H/\partial l_\alpha \partial H/\partial l_\beta - \partial H/\partial \alpha - \partial H/\partial \beta\right) = \left(h_\alpha h_\beta h_{l_\alpha} h_{l_\beta}\right), \tag{1.9}$$

where the $h_i(\alpha, \beta, l_\alpha, l_\beta)$ are the components of the hamiltonian vector field on the four dimensional phase space.

To specify particular dynamical behaviour for this model we need to impose *conditions* on it, *viz.*, initial or boundary conditions. The case we will consider is where the upper arm of the pendulum is lifted to an initial angle α_0 at t_0 and the second arm is left free to point downward, and the whole system is started at rest. These conditions are specified by the initial conditions

$$\alpha(t_0) = \alpha_0, \qquad \beta(t_0) = -\alpha_0, \tag{1.10}$$

on the angles of the two weights and the initial conditions

$$l_\alpha(t_0) = 0, \qquad l_\beta(t_0) = 0, \tag{1.11}$$

on their angular momenta. Unlike the simple pendulum model, the double pendulum does not so easily admit of a general solution. And unlike the simple pendulum

model, the double pendulum model exhibits *chaotic behaviour* in the sense of (generic) exponential divergence of nearby trajectories in phase space.[19] This means that a very small change in the initial angle α_0 can lead to very different behaviour over a short period of time. Although there are general analytic techniques for the study of chaotic behaviour (*cf.*, Lichtenberg & Lieberman (1992)), applied mathematicians in a typical modeling context want to directly study the behaviour of the model under different conditions. Without the availability of analytical tools to compute closed-form analytic solutions, applied mathematicians turn instead to *numerical methods*.

As we will describe in more detail in chapter 5, this process involves replacing the continuous phase flow φ_t of a dynamical system with a *discrete phase map* φ_i, which transforms the system forward in discrete time from t_i to t_{i+1}. It is possible to define numerical methods that provide a formula for this map φ_i. The equations defining this map replace differential equations (1.7) with *difference equations*. In the case of the Störmer-Verlet method, the difference equations take the form [20]

$$\mathbf{q}_{n+1} = \mathbf{q}_n + \tfrac{h}{2}(\mathbf{k}_1 + \mathbf{k}_2), \quad \mathbf{p}_{n+1} = \mathbf{p}_n + \tfrac{h}{2}(\mathbf{m}_1 + \mathbf{m}_2), \tag{1.12}$$

The trouble is, however, that these difference equations cannot be solved analytically either. For this reason, the ability to feasibly compute solutions to nonlinear models typically requires solving the difference equations of a numerical method using *machine arithmetic*. This requires writing computer code to implement the numerical method in a form that can be performed on a computer.

The result of this process is typically a discrete set of system space points $\mathbf{x}_n$, approximate values of the solution to the model under the imposed conditions.[21] This discrete set of points mirrors the discrete data obtained from measurements that we considered for the simple pendulum above. And in a similar way, this discrete set can be *interpolated* to obtain a continuously differentiable numerical solution function

[19] Although the mathematical notion of *genericity* applies in the much broader context of function spaces, 'generic' in this context can be understood to mean "except on a set conditions of measure zero".

[20] The quantities on the right hand sides are defined by $\mathbf{k}_1 = \mathbf{f}(\mathbf{q}_n, \mathbf{p}_n + \frac{h}{2}\mathbf{m}_1)$, $\mathbf{m}_1 = \mathbf{g}(\mathbf{q}_n, \mathbf{p}_n + \frac{h}{2}\mathbf{m}_1)$, $\mathbf{k}_2 = \mathbf{f}(\mathbf{q}_n + \frac{h}{2}(\mathbf{k}_1 + \mathbf{k}_2), \mathbf{p}_n + \frac{h}{2}\mathbf{m}_1)$, $\mathbf{m}_2 = \mathbf{g}(\mathbf{q}_n + \frac{h}{2}(\mathbf{k}_1 + \mathbf{k}_2), \mathbf{p}_n + \frac{h}{2}\mathbf{m}_1)$, where h is the time-step.

[21] There are classes of methods, however, called *continuous numerical methods*, that use interpolation to produce a continuous and usually differentiable solution to the numerical method evaluated using machine arithmetic.

$\tilde{\mathbf{x}}(t)$. And, similarly, this numerical solution function can be substituted for $\mathbf{x}(t)$ in Newton's second law (1.2) or Hamilton's equations (1.8),[22] yielding a residual or *defect* $\boldsymbol{\delta}_{\mathbf{x}}(t)$:

$$\frac{d^2\tilde{\mathbf{x}}}{dt^2} = \mathbf{f}(\tilde{\mathbf{x}}) + \boldsymbol{\delta}_{\mathbf{x}}(t), \quad \text{or} \quad \frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{h}(\tilde{\mathbf{x}}) + \boldsymbol{\delta}_{\mathbf{x}}(t), \tag{1.13}$$

showing that the error introduced by the use of numerical methods and machine arithmetic can also be regarded as a (time-dependent) *perturbation* of the vector field of the model.

Again, since the vector field specifies the dynamics of the target physical system in the context of the (newtonian, lagrangian, hamiltonian) modeling framework, a small defect (backward error) means that there is only a small difference between the dynamics that gives rise to the behaviour *computed by the numerical method* and the dynamics specified by the model. And since the construction of the model typically involves strong idealizing assumptions and mathematical reduction methods, there is rarely reason to expect the dynamics of the actual physical system being modeled to be specified exactly by (1.2) or (1.8). Accordingly, obtaining a small defect (backward error) allows one to show that the change caused by numerical computation to the dynamics of the specified model is small. And if this defect is small compared to other sources of (modeling, measurement, *etc.*) error in the problem, then the numerical perturbation changes the dynamics less than other sources of error in the model.

This backward error approach to analyzing numerical error has a distinct epistemic advantage over forward error analysis from the point of view of feasible epistemology. For nonlinear models, like the double pendulum, where solutions generally cannot be obtained analytically, we cannot *know*, in the sense of feasible access, the exact solution to the model we constructed because computation must be used and computation introduces error. The exact solution exists in a space of solutions in a classical sense of existence, but it is not feasibly accessible, *i.e.*, we cannot obtain a closed-form

[22] It is important to note that the interpolant of the machine-computed solution will generally not preserve the constraints on a hamiltonian model, including Hamilton's equations, and so the interpolant will generally not even be a solution to a hamiltonian model. Skeel (1999) has shown how to ensure that there exists an interpolant that is a solution to a modified hamiltonian model, even in the presence of rounding and other computing errors, but this approach does not provide a means of efficiently *computing* the interpolant. Moreover, Skeel's method results in modified hamiltonians with large derivatives and does not guarantee numerical stability. Thus, given currently available methods, the feasibly computable interpolants to machine-computed solutions to hamiltonian models will generally not be solutions of a modified hamiltonian model.

solution in terms of known functions. Accordingly, there is no way to *compute* the forward error, *viz.*,

$$\boldsymbol{\epsilon}_{\mathbf{x}}(t) = \mathbf{x}(t) - \tilde{\mathbf{x}}(t),$$

because there is no way to *compute* the exact solution $\mathbf{x}(t)$ to the constructed model ((1.2) or (1.8)). Provided that we can interpolate the raw numerical solution $\mathbf{x}_n$ so that it is suitably differentiable, then the backward error, $\boldsymbol{\delta}_{\mathbf{x}}$, *can* be computed, simply by substituting the interpolant $\tilde{\mathbf{x}}(t)$ into the differential equation and calculating the amount by which the interpolant fails to satisfy it.

We have seen how the defect makes it possible to interpret numerical error in terms of a perturbation of the model. In any case of modeling of real-world phenomena, other sources of error, such as *modeling error* introduced by abstractions or approximations introduced in model construction and *physical error* resulting from interactions between that part of the world being modeled and its environment, a perturbation analysis of a constructed model must be performed to ensure that conclusions drawn on the basis of the model actually apply to the phenomenon. Thus, a perturbation analysis is a standard part of error analysis in modeling. Thus, we see that the ability to compute the defect for a numerical solution provides us not only with a general feasible means of assessing numerical error, it also allows us to feasibly assess numerical error in the *same manner* as we assess other forms of error introduced into a model, *viz.*, in terms of perturbations of models.

A further epistemic advantage that the feasible backward analysis of numerical error is that, just as for the analysis of measurement error considered above, it allows the numerical solution function $\tilde{\mathbf{x}}(t)$ to be regarded as the *exact solution to a modified model*. This is because, as already mentioned, the numerical solution function can be used to compute the defect $\boldsymbol{\delta}_{\mathbf{x}}(t)$ thereby showing that $\tilde{\mathbf{x}}(t)$ is the exact solution to a modified differential equation (1.13). Thus, this feasible error analysis allows us to find a model nearby to the specified one (nearby vector field/dynamics) for which we have numerically computed the *exact* solution. Accordingly, it is not necessary to regard numerical solutions to mathematical models as approximate either. This reveals the manner in which numerical methods are, like idealizing assumptions and model reduction methods, *just another kind of modification of a model used in the search for an exact solution to a model problem.*

Returning again briefly to consider the double pendulum model, recall that its primary mathematical interest is that it provides a very simple example of a chaotic dynamical system. Given this interest, it is important to be able to study the behaviour of the system over long time scales, long after a real double pendulum would have slowed to a resting state. And this is the reason that hamiltonian mechanics is used for a numerical study of the double pendulum. As we will discuss in more detail in chapter 5, it is possible to exploit the special geometric, *i.e.*, symplectic, structure of hamiltonian models in order to construct numerical methods that produce accurate results over very long time scales, *even using machine arithmetic*. This is accomplished by so-called *symplectic numerical methods*. For these methods, the map that computes a time step on the basis of the previous value of the solution, *i.e.*, the map that transforms the system forward in time, is a *symplectic map*. This means that, just like the phase flow of a hamiltonian dynamical system, it *preserves* the symplectic structure on phase space. This is to say that the map is also a symplectomorphism, only a discrete rather than a continuous one.

The upshot of this for a backward analysis of numerical error is that the phase space points $\mathbf{x}_n$ computed by a symplectic method lie exactly (or to a high degree of approximation) on the trajectories of a nearby *hamiltonian model*, *i.e.*, it also exactly (or very closely) satisfies an equation of the form (1.8) but with a slightly modified hamiltonian function H (see Sanz-Serna, 1992, 275). The result of this is that the discrete (numerical) model has essentially the same energy conservation properties that the specified (continuous) model had. The importance of this is that symplectic methods applied to hamiltonian models have an error that grows only linearly in time, rather than exponentially for general methods (Stuart & Humphries, 1996). In the case of periodic solutions, which are common in celestial mechanics, the error grows linearly in the number of periods. Accordingly, symplectic methods with very small error can be used to accurately compute orbits over long periods of time, making these methods very useful for various cases of modeling the motion of objects in the solar system. Numerical methods, including symplectic methods, that respect the special geometric structure of a mathematical model problem are called *geometric numerical methods*, and the backward analysis of error for such problems that takes the geometric structure into account is called *structured backward error analysis*. For

reasons that will become clear, it is significant that a proper understanding of the nature of these methods requires an appreciation for not only the complex geometry but also the algebraic form of the equations used to formulate them.

These two scientifically *very* simple examples of theory application now provide us with a glimpse of the depth of complexity of method involved in the application of real theories in applied mathematics. And though a thorough examination of the ideal simple pendulum already causes problems for classical accounts of theory application, the next more complex version of it requires the use of methods that are not in any adequate way accounted for in classical approaches, given the need for numerical methods, machine arithmetic *and* the feasible analysis of the error that they introduce. Through the consideration of details of the full application process in the cases of the double pendulum and the problem of estimating an impact probability for near-Earth objects in part II, I will argue that the syntactic and semantic approaches to theories and their relation to the world are unable to adequately account for the actual process of theory application in science. Along with this I will argue that the picture of the process of relating a theory to data or phenomena as a logical derivation is an false representation of the process. The central purpose of the detailed investigation in part II is to develop methods that do accurately account for the process of theory application and simultaneously elucidate significant aspects of the structure and behaviour of the process. The results of this development are interesting new perspectives on mathematical modeling in science and its relationship to the handling of data and numerical computation.

1.4 New Perspectives on the Relation of Theory to Data and Phenomena

Given the complexity of scientific method as a phenomenon, methods that are capable of generating insight into the feasibility of scientific knowledge must be capable of elucidating complex phenomena. This requires the screening out of details that are not relevant to a given epistemological task. The way that this is handled in applied mathematics is to use systematic means of abstraction from the full detail of the phenomenon being examined. Batterman (2002*b*) and Wilson (2006) have shown how

mathematical modeling techniques used in applied mathematics accomplish this kind of abstraction to gain insight into the dominant behaviour of a phenomenon. This abstraction process, *inter alia*, serves to make information about a phenomenon more feasibly accessible. The results of this study show that the processes of data handling and numerical computation play a formally identical role in theory application. The argument for this proceeds by developing a technique of epistemological abstraction from the details of the case examples, the use of which reveals interesting structural patterns throughout the application process.

The epistemological abstraction method just described is modeled on the mathematical model reduction methods used in applied mathematics. A characteristic feature of these reduction methods is that they begin with a more complex description and then generate a simpler description that still contains the information required to describe the dominant behaviour of the phenomenon. In a similar way, I treat the detailed consideration of the case examples as the complex description that requires simplification. I then use a systematic means of abstracting out detail from the examples in a way that preserves dominant features of the inferential process involved in the application process. This systematic abstraction method is based on a model of theory application that I develop through the study in part II.

In chapter 2, I argue that an adequate formal method for studying feasible theory application requires the representation of both the syntax and the semantics of a theory or model together with how they are *correlated*; in particular how they *covary* in the sense that syntactic modification of a model implies certain semantic modifications and *vice versa*. The ability to represent syntactic and semantic covariation allows a faithful representation of the different methods of variation or transformation of models used in the process of feasibly applying a theory, including an account of the error such methods introduce. In chapter 3, I present the basic structure of a model of theory application that is capable of representing application processes in terms of the syntax and covarying semantics of a theory, and in such a way that can model the effects of the kinds of error that are introduced in the process of application. The details of this model are extraordinarily complex, which requires the development of a way of presenting the model that brings out the epistemologically most significant, or *dominant*, behaviour. For this purpose I develop a conceptual language that can,

in a qualified sense, be regarded as generalizing concepts from mathematical logic by introducing imprecision and descriptions that are local to a given context. I present the details of this conceptual language as they are needed to clarify structural and behavioural features that are found in different parts of the application process.

Chapter 2 has a dual purpose for the study as a whole. In part, it provides clear examples of logical approaches to scientific epistemology, facilitating demonstration of the limitations of logical approaches by showing those features of the case examples that are omitted or misrepresented. But it also functions to set up the epistemological modeling approach I develop in this study, since there are respects in which the methods I use evolve out of the logical approaches. In particular, this helps to motivate and specify the details of the model of theory application in terms of covarying syntax and semantics in mathematical modeling. A major difference in my approach, however, is that it is designed to result in a way of clarifying the structure and behaviour of the application process by using an abstraction process to eliminate the details of the full model of applying a theory that are irrelevant to the dominant inferential processes involved in application.

The consideration of the case examples is spread throughout the three chapters in part II. In chapter 3, I consider in detail the process of constructing a mathematical model in the cases of the double pendulum and in near-Earth object (NEO) modeling. I show that this process relies heavily both on abstraction and approximation, and explain the ways in which logical approaches fail to accurately or adequately represent this structure. I also consider the main kinds of error that are introduced in the model construction process and I show how my epistemological model clarifies the ways that error affects the structure of a model.

In chapter 4, I consider in detail the process of handling data in astronomy, particularly in the case of solving the orbit determination problem for NEO tracking. I show that this process relies heavily on a complex variety of mathematical models in a way that is not adequately accounted for in typical presentations of the theory-world relation in terms of a "hierarchy of models" connecting theory with data or phenomena. In particular, I show that, at least in the case of astronomy, the *process* of using data in scientific inference is connected to the structure of the models used in data handling, a structure which is not hierarchical. I also consider what

is involved in ensuring that inference based on data is reliable and I show how my epistemological model accounts for this reliability in terms of the structural stability of transformations between mathematical models.

In chapter 5, I consider in detail the process of computing numerical solutions to mathematical models. I show that this process is not a "merely pragmatic" concern since it is required for us to actually know the predictions of a theory in many cases and I show how my epistemological model, unlike the usual logical approaches, is able to faithfully represent methods used in numerical computing. I consider the main kinds of error introduced in numerical computing and how they affect the structure of a mathematical model. And I also consider what is involved in ensuring that inference based on numerical computations is reliable and show how my epistemological model accounts for this reliability in terms of a kind of structural stability that is identical in essence to the structural stability of inference using data.

Finally, in chapter 6, I summarize the results of the study and show how my epistemological model represents the entire application process. The model depicts theory application in terms of recursive methods of inference that, after a small number of steps, terminate in either accessible data or accessible solutions to a model. The steps in inferences of this kind, whether used in modeling, data handling and computing, all share a common structure of translations between modeling frameworks that allow certain information to be more easily accessed. The recursive inference process terminates when the information sought is feasibly accessed in some, possibly encoded, form. My epistemological model also provides a way of clarifying what is involved in the stability of inference under the various kinds of error introduced in the process and can be a useful tool in the analysis of inferential stability.

Therefore, I conclude that the logical approaches are ill-suited to the investigation of feasible scientific inference and method. A suitable approach must be able to faithfully represent actual scientific methods in detail and provide a means of simplifying this description in order to clarify epistemologically important structure. I therefore conclude that for an epistemological theory to adequately account for actual scientific inference it must be modeled on the methods of applied mathematics and not pure mathematics. Given the insight into large-scale patterns in feasible scientific inference seen to come out of the epistemological tools that I develop through this

study, I conclude that these methods are effective for feasible epistemology, at least in scientific contexts that use methods similar to those used in the case examples. Given that accurate description of feasible scientific methods requires logical concepts that generalize those of mathematical logic, I conclude that the picture of the theory-world relation in terms of a derivation from theory provides a badly distorted and false representation of how we actually gain knowledge from scientific theories, obscuring the actual extent and limitations of our knowledge.

Chapter 2

Applying Theories in an Ideal World: Classical Accounts of Epistemological Access to the World

A central aim of this study is to determine the characteristics that an epistemological method must have to be able to gain insight into the complexity of method in real science. This relates to the syntactic and semantic views of theories in two main ways. The first relation concerns my claim that that there are numerous ways in which typical syntactic and semantic approaches to theories and their relation to data or phenomena are either inadequate or inaccurate in terms of how they represent how theories are applied to gain actual knowledge in real science. To argue effectively for this claim requires clearly demonstrating those features of the syntactic and semantic approaches that result in inadequate or inaccurate representations of scientific method. This, in turn, requires a clear specification of how syntactic and semantic approaches represent theories and their application.

The second relation of this central aim to the syntactic and semantic approaches concerns the approach to modeling the epistemology of feasible scientific inference that I develop through the study of the two central case examples in part II. This epistemological modeling approach that I develop can be seen, in certain respects, to evolve out of syntactic and semantic approaches to theories and their application. Showing how certain advantageous features of the syntactic and semantic approaches are preserved in a modeling approach suited to feasible scientific inference will provide partial motivation for the epistemological modeling method.

Another central aim of this study is to argue that for epistemology of science to

develop fully as a scientific discipline it must adapt methods from applied mathematics to the purposes of investigating feasible scientific inference, rather than relying solely on methods from pure mathematics and metamathematics. One of the ways that I argue for this thesis is by showing how an epistemological modeling method that adapts methods from applied mathematics, developed throughout part II, provides insight into feasible scientific inference that the syntactic and semantic approaches typically used in epistemology of science do not provide and are not naturally capable of providing. Indeed, this will be seen to provide an argument that in order to address these limitations, syntactic and semantic approaches must adopt an approach that is, in certain fundamental respects, essentially the same as the basic modeling framework I develop. This line of argument also relies on having a clear exposition of certain important features of the syntactic and semantic approaches to the theory-world relation.

It is with these two central aims in mind, then, that we consider the syntactic and semantic approaches in this chapter. We will begin with the consideration of the "received view" of the logical empiricists. I will argue that certain basic features of this approach to reconstructing scientific theories are incompatible with the aim of investigating feasible scientific inference. Since this view is often used as the representative of syntactic approaches and it has not been amended since the 1960s, I also show how a more up-to-date syntactic approach is capable of usefully modeling important structural and behavioural features of feasible scientific inference. This is followed by a consideration of the semantic approaches to theories developed by Suppes and van Fraassen. Although these approaches are seen to fare far better in their ability to accurately represent scientific practice than the received view, they nevertheless also have basic features that make them incompatible with investigating feasibility in science. In fact, I argue that one of these is a feature typically regarded as an advantage of the semantic approach, *viz.*, that it characterizes a theory independently of a particular linguistic formulation. I argue that the actual linguistic formulation of a scientific theory must be represented for a logical reconstruction to explain the reliability of the methods actually used in the practice of applied mathematics to gain knowledge of aspects of the behaviour phenomena. A general reason for this, which we will discuss, is that because many mathematical techniques used in modeling are

applied as rules of symbolic manipulation on particular symbolic constructions, the explanation of the reliability of these techniques requires appeal to 1) those particular symbolic constructions, 2) their mathematical meaning and 3) how the prior two aspects are related. This is a problem for the syntactic and semantic views because the explanation is provided neither in purely semantic terms nor in terms of a canonical linguistic reconstruction.[1] The chapter concludes with a brief consideration of the form a basic formal framework must have to adequately represent and account for feasible scientific inference.

2.0.1 Abstraction and Formal Reconstructions of Science

The formalizations of scientific theories that the so-called received view provided are considered as rational reconstructions of theories, putting them on a canonical rational foundation with the intention of screening out uncertain or metaphysical notions or concepts. In this way, these formalizations can be seen as *models* of scientific theories in essentially the same sense used in applied mathematics, *viz.*, rational reconstructions can be construed as abstract representations of scientific theories that can be studied independently in order to gain insight into their structure and content. Suppe (1974) discusses this explicitly:

> The Received View begins by specifying a canonical formulation for theories in terms of an axiomatic calculus and rules of correspondence. This canonical formulation is claimed to stand in the following relation to a scientific theory: Any given scientific theory could be reformulated in this canonical manner, and such a canonical formulation would capture and preserve the conceptual and structural content of the theory in such a manner as to reveal most clearly and illuminatingly the conceptual or structural nature of that theory (*op. cit.*, 60).

This is a strong claim concerning what sorts of structure the received view is able to represent faithfully. Not only that, it implies that one canonical axiomatic presentation of a theory will best elucidate the conceptual nature of a theory and its relation

[1] Another reason for this that we will not discuss is that different languages have different computational complexity from the point of view of making certain kinds of inference about a given domain, which is very important for feasibility considerations. As a simple example, consider arithmetic calculations performed using positional notation, as we do, versus using Roman numerals. The former drastically reduces the computational complexity of the former, which had significant historical consequences.

to the world. There are general issues with this claim, in that given that very few theories have been presented in this way and, as a result, the state of the evidence for such a general *in principle* claim is quite weak. Our concern here, however, is not so much whether the received view, or any other approach to the formal reconstruction of theories, can accomplish such a task in principle, but rather what such approaches can accomplish in terms of practical investigation of science.

Although reconstructions using mathematical logic do not typically make such strong claims about their epistemological capabilities, they nevertheless claim to be elucidating epistemologically significant features of real theories and their relation to the world. If, then, we are to regard a successful formal reconstruction of a theory and its relation to the world as an abstraction from the details of actual scientific theories and their relation to the world that preserves and elucidates the essential content, then (at least) two things must be the case:

(A) The formulation of a real scientific theory (and its relation to the world) in a reconstruction must faithfully abstract and elucidate the essential structural and conceptual content from a fully detailed description of the theory (and its relation to the world); and

(B) The structural and conceptual content eliminated in the abstraction must be irrelevant or insignificant in relation to the nature of a theory (and its relation to the world).

In order to definitively answer the question of whether any reconstruction process meets these conditions, it is necessary to be able to assess the structural relationship between the reconstructions and fully detailed descriptions of real theories and their relation to the world. Without a characterization of the fully detailed description of a theory, these conditions are impossible to assess rigorously; but we may still assess an approach in relation to these conditions.

Since logical reconstructions are typically formulated *a priori*, in one manner or another, and not through a process of investigation that abstracts content from fully detailed descriptions, a given approach makes a decision *a priori* about what conceptual and structural content is and is not epistemologically relevant. Nevertheless, for a given approach we can assess condition (A) by determining whether the approach

faithfully captures epistemologically important content and we can assess condition (B) in terms of what content the approach evidently determines is *not* epistemologically relevant.

I will argue over the next two sections that neither typical syntactic approaches nor typical semantic approaches meet these two conditions adequately. What these logical approaches do well is provide a means of studying general conceptual or structural properties of theories, particularly where particular real scientific theories can be studied axiomatically. What I will argue that these approaches do not do adequately, however, is elucidate the relationship between scientific theories and the phenomena they are used to describe or represent. A major reason for this is that logical approaches to the epistemology of science typically assume that there is some clear, more or less direct, relationship between theoretical models and phenomena or the world and, consequently, that the details of the methods actually used to gain knowledge about phenomena using scientific theories have little or no relevance to the epistemology of scientific theories. For example, concerning the relationship between theoretical models and the world, Giere (2010*a*) says the following:

> How does one connect abstract models to specific real physical systems? This requires at least two processes which I call "interpretation" and "identification"... . For interpretation, elements of an abstract principled model are provided with general physical interpretations such as "mass," "position," and "velocity." Such interpretations are already present in the statements that characterize the principled models. Scientists do not begin with an "uninterpreted" formalism and then "add" interpretations. For identification, elements of a representational model are identified (or coordinated) with elements of a real system. Do we, as theorists of science, need to give a more detailed account of the processes of interpretation and identification? I think not. We can pass this job off to linguists and cognitive scientists. We know it can be done because it is done. That is enough for our purposes (*op. cit.*, 271).

The central problem with general classical accounts of the theory-world relation such as this is that, though they provide an underlying reason why our methods work when they work, *they tell us nothing about when our methods fail when they do and what underlies the failure.* It is by knowing the conditions under which methods

fail and why they fail that allows us to explain the accuracy and reliability of our methods of applying theories. Since the reasons why methods fail have a great deal to do with the nature of the methods themselves, developing a richer understanding of the reliability of our methods of gaining knowledge about phenomena requires detailed study of the nature of various kinds of scientific methods. In this way, a detailed study of the methods used to gain knowledge in science has a great deal to contribute to epistemology of science. This, I argue, is the main contribution feasible epistemology stands to make to the epistemology of science. And it is a contribution that is generally complementary to the general accounts of the relationship between abstract theoretical descriptions and phenomena provided by classical epistemology. Classical epistemology elucidates the possibility and limitations of theoretical knowledge in general and feasible epistemology elucidates the possibility and limitations of theoretical knowledge in the real world given the methods that we have.

Since a central concern in this study is to assess the relative strength of methods for use in scientific epistemology that are adapted from pure mathematics (or metamathematics) in comparison to those adapted from applied mathematics, we must assess the relative strength for both the purposes of not only feasible epistemology but classical epistemology as well. Consequently, our assessment of how well epistemological methods satisfy conditions (A) and (B) above will involve considerations of methods used for both classical and feasible epistemology, with an emphasis on the latter given the aims of this study.

Given that how well epistemological methods function is in proportion to how well they abstract important structural and conceptual content from real science, we may see that methods in scientific epistemology play a similar role to that of model reduction techniques in applied mathematics. Applied mathematicians often start a modeling project by formulating a complex model on the basis of theories, and then simplify these models to obtain a more analytically and computationally tractable model that still captures the dominant behaviour of the system being examined. These techniques often collapse complex physical behaviour into boundaries and singularities. Given the detailed physical behaviour avoided by this approach, has been dubbed *physics avoidance* by Wilson (1998). The received view, then, is seen to be reducing the full complexity of the syntactic structure of theories and their

relations to data, phenomena and measurement to a canonical symbolic form. If we consider the theories, models, reduction techniques, measurement devices, statistics, computation, and error analysis used in applied mathematics as all being included under the umbrella of *method*, then we can see epistemological methods, including rational reconstructions like the received view, as a form of *methodology avoidance*. The idea of methodology avoidance being to use a reconstruction or representation to reduce or eliminate complex details about the structure and behaviour of real scientific methods. Thus, our central concern in this chapter is the assessment of the efficacy of logical approaches as methodology avoidance strategies.

2.1 Syntactic Approaches to Theories

The central syntactic approach to theories and their relation to the world in the philosophical literatures is still the so-called received view of theories, developed by the logical-empiricist philosophers of the Vienna Circle. This view treated empirical knowledge as certain and statements only as contentful if they could be grounded or directly verified in experience. On this view, the epistemic content of science is to be revealed through the elucidation of the logical structure of scientific language. This required the reconstruction of scientific language in a logical language, which would allow the true content of the theory to be explicated. The metaphysical content of scientific theories was regarded with great scepticism, indeed regarded as meaningless, since metaphysical claims are not verifiable in experience. Thus, an important feature of the reconstruction of scientific language in logical terms was the removal of any metaphysical concepts that obscured the true content of a theory. The general aim of this account, then, was to ground all knowledge in principles that are certain and could be formulated in the language of formal logic.

The received view has been widely discredited since the 1960s, which, for most philosophers of science, has made it a matter of historical interest but of little interest for contemporary approaches to scientific methodology. As we will see, however, there are a number of characteristics of the syntactic approach that are useful, insightful, and worth preserving in future approaches to scientific epistemology, where it is possible to do so. Perhaps the most important feature of this approach that we wish to

preserve is, in contrast to the semantic approaches, the recognition of the epistemological importance of the language of scientific theories. Symbolic constructions are what scientists deal with most concretely and most directly in their reasoning and methods. Consequently, an adequate account of scientific theories in practice must account for the actual symbolic constructions that scientists use and the accuracy and reliability of the methods that employ them. Unlike the received view, the approach we take in this study represents directly and faithfully the actual symbolic constructions used by scientists, correcting what I argue is a genuine limitation of typical syntactic approaches in epistemology of science. Nevertheless, our approach shares with these approaches a recognition of the epistemological importance of scientific language. To clarify the advantages and disadvantages of a syntactic approach, let us consider the received view in more detail.

2.1.1 The Received View of Theories

The picture of theories provided by the received view is tightly bound to experience and the world. This occurs because the logical empiricists held the view that for the theoretical language of a scientific theory to be meaningful, to have *cognitive significance*, it had to be grounded in experience of phenomena. Accordingly, there are not separate views of theories and data; rather, the view of theories and data has the two inseparably entangled together. All cognitively significant discourse about the world was required to be empirically verifiable, at least in principle. Accordingly, all assertions of a theory were required to be reducible, at least partially, to assertions about phenomena in an observation or protocol language. This reduced the problem of how theoretical statements are verified to that of determining how observation or protocol assertions were to be verified.

There was initially debate about the terms in which observational experiences could be described such that observational or protocol sentences were clearly true or false in a particular experiential context. Some members of the Vienna Circle preferred a phenomenalistic language, specified in terms of a sense-datum language. These would be immune to any concerns about truth or falsity, since it was believed that sense-datum descriptions of one's phenomenal experience could be known with absolute certainty. Others preferred a physicalistic language, a language describ-

ing the observable properties of material things. Since the observation or protocol language would refer only to observable properties, assertions about experience of phenomena would be intersubjective, and for verifiable propositions it would be entirely clear whether or not a given thing had the property asserted of it. Physicalistic language won out in the end, causing the received view to endorse materialist metaphysics, even though the content of all assertions was grounded in observation.

Following the formulation in the classic account given by Suppe (1974), the received view takes a *theory* to be formulated in a first order language $\mathscr{L}$, which has an associated logical calculus $\mathscr{K}$, specifying the laws of classical valid inference. The language $\mathscr{L}$ could be augmented by modal operators in order to support the use of counterfactuals. The language $\mathscr{L}$ is then required to have a specific substructure. The nonlogical primitive constants, or terms, of $\mathscr{L}$ are divided into two disjoint classes: the nonempty observation vocabulary V_O and the theoretical vocabulary V_T. The language $\mathscr{L}$ is then divided into three sublanguages, each of which having an associated subcalculus, the restriction of $\mathscr{K}$ to the given sublanguage:

(i) The *observation language*, $\mathscr{L}_{\mathscr{O}}$, which contains no quantifiers or modal operators, contains all the observational vocabulary V_O and none of the theoretical vocabulary V_T. $\mathscr{L}_{\mathscr{O}}$ is the language used to formulate assertions of direct phenomenal experience.

(ii) The *logically extended observation language*, $\mathscr{L}_{\mathscr{O}}'$, an extension of $\mathscr{L}_{\mathscr{O}}$ to include quantifiers and modal operators, *etc.*, of $\mathscr{L}$. This language is used to formulate generalizations about experience and supports the use of counterfactual statements about experience.

(iii) The *theoretical language*, $\mathscr{L}_{\mathscr{T}}$, the sublanguage of $\mathscr{L}$ that contains no V_O terms. $\mathscr{L}_{\mathscr{T}}$ is used to formulate the theory or theoretical framework, independently of its interpretation in experience.

These sublanguages do not exhaust $\mathscr{L}$ because it also contains *mixed sentences*, which contain at least one term from each of V_O and V_T. This sort of sentence is necessary to give the theory an interpretation in experience.

It remains then to give the language $\mathscr{L}$ an interpretation. Given the crucial role of the observation language in giving theoretical content meaning, $\mathscr{L}_{\mathscr{O}}$ must be

given a *semantic interpretation* in a domain of concrete observable events, where the obtaining of properties and relations is directly observable. This can then be extended to an interpretation of $\mathscr{L}_{\mathscr{O}'}$ for empirical generalizations and counterfactual claims. This provides a *partial interpretation* of the observational part of $\mathscr{L}$. The *partial interpretation* of the theoretical terms and $\mathscr{L}$-sentences containing them is given by the following two kinds of postulates:

(i) The *theoretical postulates* T, the axioms of the theory, containing the terms of V_T; and

(ii) The *correspondence rules* C, which are mixed sentences that contain at least one term from each of V_O and V_T. These rules provide a partial interpretation of the theoretical part of $\mathscr{L}$ by grounding it in the semantic interpretation of the observation language.

Letting T and C designate the conjunctions of the theoretical postulates and correspondence rules, respectively, the scientific theory based on $\mathscr{L}$, T and C, is denoted TC. The theory may alternatively be represented as the set of sentences $\mathbf{T}$ formed by the deductive closure $\mathscr{K}(TC)$ of the postulates TC.

Although the form and function of correspondence rules changed during the development of the received view, throughout they function to specify which experimental procedures are admissible for applying a theory to phenomena. The role of the correspondence rules, ultimately, was to specify sufficient conditions for the statements of the theory to be empirically meaningful. This provides clear criteria for determining the observable consequences of the theory that make it testable, and in the process provides a partial interpretation of the terms V_T by specifying their observational content. Terms are partially defined rather than fully defined because the correspondence rules only place conditions on admissible experimental procedures for application to observable phenomena, rather than giving a set of necessary and jointly sufficient conditions in terms of observable phenomena. The latter cannot, in general, be done without making particular measurement devices part of the empirical meaning of the theory, which would mean the theory would change with a change in instrumentation. Moreover, the observable consequences are often understood to be the empirical manifestation the interaction of theoretical entities according to the axioms or laws of

the theory (Suppe, 1974, 25). Consequently, theoretical terms cannot then be defined purely in terms of observational phenomena.

The most basic assumption on any syntactic view of theories is that the logical structure of scientific language is to be elucidated through its reconstruction in a *language* in formal logic, usually first order logic. This leaves the manner in which a semantics is specified for that language open. The received view handled this by using correspondence rules to ground meaning in experience. Since Tarski's work in model theory had a profound effect on both logic and philosophy of science, the semantics of a theory is now generally provided not by interpretation in experience, but rather by sets of **structures** in the universe $\mathfrak{V}$ of sets. A formal *theoretical* semantics of this kind provides no means of discussing the empirical content of a theory or the manner in which it is confirmed or disconfirmed. Consequently a second *empirical* semantics is required to relate the theory to experience and the world. This is now most commonly understood in terms of a relation between the set-theoretic semantics of the theory and some empirical semantics. Generally, since Suppes (1962, 1969), this is understood in terms of a hierarchy of successively less general frameworks of models, down from general scientific theories to models of data and experiment. This then raises issues of the nature of the relation between set-theoretic structures and the world, which can be treated in an empiricist manner, as van Fraassen (2010) advocates, or in a realist manner, as advocated by Da Costa & French (2003), French & Ladyman (2011), and others. This two-stage semantics, an internal *theoretical semantics* in the universe of sets and an external *empirical semantics* in data or phenomena, is how semantic approaches to theories generally specify the theory-world relation. A purely syntactic approach, in the sense of a representation of a theory as a logical language, however, must provide an empirical semantics by relating the language of the theory to experience and perhaps also the world, which is precisely how the received view proceeds.

With the observational consequences of the theory specified, it is then possible to specify adequacy conditions for the theory. The *observational adequacy*[2] of the

[2]I take this term from Muller (2011), which is distinguished from the concept of *empirical adequacy* from van Fraassen's constructive empiricism. According to van Fraassen's definition, a "theory is empirically adequate if it has some model such that all appearances are isomorphic to empirical substructures of that model" (van Fraassen, 1980, 64). In distinction to this definition, the concept of *observational adequacy* depends on the distinction between theoretical and observational concepts,

theory is then defined in terms of the set of empirically verified observational sentences $O_t(\mathbf{T})$, which grows in time t. If $O_t(\mathbf{T})$ is included in the set $\mathbf{T}$ of sentences of the theory, then the theory is confirmed. Confirmation grows as the set $O_t(\mathbf{T})$ grows while still being contained in $\mathbf{T}$, and $\mathbf{T}$ is falsified if $O_t(\mathbf{T})$ contains a sentence not in $\mathbf{T}$. Muller (2011) calls this the *Formal-Linguistic View*.

2.1.2 Limitations of the Syntactic View

One basic limitation of a purely syntactic approach is that the reconstruction of a theory as a logical language does not provide a means for representing reasoning within the theoretical, model-theoretic semantics of that language. The semantics is indeed captured as the set of **models** of a syntactic theory or syntactic model, but the syntactic approach does not provide the ability to *describe* the **models** within the representation provided by the rational reconstruction. This is important from the point of view of representing the feasible reasoning processes that applied mathematicians use in practice, because a great deal of the time the reasoning involved is based on the understanding of the geometry and topology of the spaces the theory relies on, not just the algebraic manipulation of equations.[3] The received view was understood to *allow* the kind of reasoning that applied mathematicians do in terms of the mathematical picture of a theoretical or modeling framework, but it is not able to represent such reasoning, showing that it eliminates epistemologically significant content.

One reason that this is an epistemic limitation of the syntactic approach, and not merely a methodological one, is that the *explanation* of why model reduction methods are effective and reliable is not provided by the syntactic transformation of equations alone; it is necessary to understand the consequences of algebraic manipulation in

depends on the historical development of observations in time, and is defined independently from a semantic view of theories.

[3]It is possible to conceive of a syntactic approach that is augmented with a language to talk about its models, or even a language that constitutes a model following the Henkin method for proofs of completeness. A proponent of the syntactic view would have to demonstrate the viability of this approach, however. There may be good philosophical reasons to favour such an approach, since it would allow one to explain the reliability of semantic language without the need to refer to abstract sets, explaining how mathematical language works without the need to posit the existence of abstract objects. For purposes of practical usability, however, such an approach seems unlikely to be desirable.

the geometric or topological spaces that interpret the algebraic constructions. What will be a recurrent example of this is the technique of linearization, discussed in the last chapter in the context of the ideal simple pendulum model. This technique is applied as an algebraic rule that replaces a given function in an equation with the constant and linear terms of its Taylor series. What explains why this technique is reliable is that the variation of a quantity appearing in a model is often close enough to being constant over a limited range of conditions that treating the variation as constant provides a highly accurate approximation. This is to say that within that limited range of conditions the variation of the quantity is effectively constant. The algebraic technique of linearization works and is reliable, then, because the algebraic manipulation has the geometric effect of replacing a curved quantity (non-constant variation) with a flat one (constant variation) as a result of the correlation between the algebraic form of the function and its interpretation in a geometric space. Since purely syntactic approaches (no description of $\mathfrak{models}$) do not represent the geometric meaning of the functions in equations, they are incapable of explaining why methods like linearization are effective and reliable.

There are, however, much more serious problems with how the syntactic view is typically used to talk about scientific theories. A considerably more severe issue is that theories are rarely reconstructed directly, rather they are usually treated schematically as a set of sentences T of a formal language. This is an issue that Wilson (2006) describes as the *theory T syndrome*. In part the theory T syndrome involves the consideration of theories in general using generic logical terms incompatible with the conceptual and inferential structure of real theories, obscuring the kinds of reasoning that are possible using scientific methods and the real limitations on possible knowledge. The theory T syndrome also involves generic treatment of theories without subsequent checking that general considerations or consequences apply to actual scientific theories. This leads to situations where philosophical considerations of science are wildly out of contact with real science.

The basic reason for this is that representing a real theory as a theory T eliminates all of the inferential structure and conceptual content of an actual theory and its relation to the world, showing that such a representation fails miserably to meet adequacy conditions (A) and (B) above. Not only is the mathematical meaning of

theories eliminated with such an approach, but so also is all the complex *syntactic* structure that is crucial to understanding how scientific inference actually works in practice. Such an approach might be less problematic were it not for the fact that scientific reasoning is rarely presented in axiomatic form and is rife with non-deductive methods of approximation. In reasoning about a theory T, all of the structure and content essential to accounting for and explaining feasible inference is eliminated. Even more than this, however, whatever structure and content of a theory essential for studying the classical epistemology of a theory is also lost. Considering that such an approach is supposed to be gaining epistemological insight by representing a theory in terms of its syntax, it fails miserably, since this general approach actually prevents elucidation of more detailed symbolic structure that can be crucial to understanding and explaining the use and success of methods used in actual scientific theories. This shows that, in the manner it is often used, the syntactic view fails to faithfully represent that part of a theory it is intended to present in a clear and illuminating form. In this respect, syntactic approaches are very strong methodology avoidance strategies, and too strong to be effective tools for elucidating epistemology of science in the manner in which they are often used.

The limitations just identified do not, however, rule out all syntactic approaches as being useful for investigations of feasibility in scientific epistemology. On the contrary, as we will now see, such approaches stand to be quite useful for feasible epistemology when designed to faithfully represent the symbolic manipulation used in real scientific inference.

2.1.3 Modeling Application on the Syntactic Approach

Since the received view is the principal representative of the syntactic view and its development ceased in the 1960s, it is easy to dismiss a syntactic approach to representing theories and their relation to the world. I am arguing, however, that the ability of a formal approach to clarify and explain the efficacy of mathematical methods in science requires the representation of the syntactic structure of a theory *and its methods*. Consequently, we will briefly consider the manner in which a syntactic approach can successfully represent the syntactic structure of theories and their methods of application.

To this end, we will reconsider the case of application of newtonian mechanics to the motion of a simple pendulum. Our consideration of the construction of such a model began with equation (1.1), Newton's second law, reproduced here in a rearranged form:

$$\frac{d^2\mathbf{x}}{dt^2} = \frac{\mathbf{F}}{m}. \tag{2.1}$$

The syntactic structure of this equation can straightforwardly be represented formally, with t being a variable, *i.e.*, symbol ranging over constant terms; $\frac{d}{dt}$ being a unary operator on functors, *i.e.*, function symbols; $\mathbf{x}$ and $\mathbf{F}$ two functor variables, *i.e.*, symbols ranging over functors; $=$ the identity relation; $\vdots$ the division operation; and m a constant term. This captures faithfully the basic algebraic structure of the equation.[4] Its geometric interpetation in terms of state spaces, function spaces and parameter spaces is eliminated entirely in the purely syntactic reconstruction we are considering. Recovering the geometric interpretation would require an additional language that specifies the theoretical semantics of the theory.

Let us suppose that this single equation is the sole axiom of the theory. The set T of postulates would then just contain syntactic representation of equation (2.1) just described. We will suppose that the logical calculi required will be based on first order classical logic. The theoretical postulate is provided with a physical interpretation in terms of the positions of particles in space at a time, the mass of the particles and the total force acting on each particle, which we may suppose is specified empirically in terms of sufficient conditions on the measurement of these quantities.

In the case of the pendulum system, this interpretation could be captured using an observation language $\mathscr{L}_{\mathscr{O}}$ appropriate for the description of observations and manipulations of the physical pendulums and correspondence rules C connecting the theoretical terms to observation terms. The theoretical terms that require interpretation are t, $\mathbf{x}$, m and $\mathbf{F}$. The first three terms can be defined in terms of sufficient conditions for measuring time, position and mass, but since forces are not directly measurable, $\mathbf{F}$ would be defined in terms of measurement procedures for $\mathbf{x}$, m and t. Since the sole axiom of the theory is a differential equation, the language $\mathscr{L}$ and

[4] By faithful translation of the basic algebraic structure of the equation, I mean in the sense that an operator is applied twice to a function, which is then equated to another function divided by a constant.

calculus $\mathscr{K}$ of the theory must include enough classical analysis to handle the treatment of second order ODE. This covers the basic reconstruction of the theory. Now let us consider the modeling process.

We saw that the construction of the model of the ideal simple pendulum required a number of constraints to be imposed, contraints abstracted from the consideration of the physical pendulum being modeled. This could be represented syntactically in terms of observation terms describing aspects of the apparent or effective behaviour of the system, such as the concept of a rigid body, which are then translated into the theoretical framework as theoretical terms. Thus, representing this construction procedure requires rules analogous to correspondence rules, except that they translate certain observational terms into theoretical terms rather than the other way round. Together with T, the equations that characterize the constraints (rigid body, gravitation near Earth's surface, total force, *etc.*) would then be taken to form a set M of assumptions constituting a mathematical model. Any inferences drawn from this model would occur in a subcalculus $\mathscr{K}_M$ of the calculus $\mathscr{K}$ of the theory, which is a derivation system for the ideal simple pendulum model (1.3). In principle, $\mathscr{K}_M$ can be used to derive solutions to the model under any imposed initial and boundary conditions.

Now, the linearization step that converted the original model to the simplified one (1.5) is not deductive. Accordingly, this is not an inferential procedure that can be represented in $\mathscr{K}_M$. The linearized model could be represented, however, by another model M' that is the same as M except that the original nonlinear equation for the total force has been replaced with a linear one. This modified model would then have its own subcalculus $\mathscr{K}_{M'}$ of $\mathscr{K}$. The issue that arises here for the syntactic approach is how to represent the linearization method that transforms M into M'. For a feasible view of theories, it would be natural to represent them as part of the methods of the theory, since they are required in the process of model construction. For a classical view of theories, however, it would be more natural to regard these methods simply as part of the procedure needed to get results out of theories. Either way, these methods require a formal representation on a syntactic view of theories and their relation to the world that accounts for feasible methods.

In order to accommodate these methods, the calculus $\mathscr{K}$ of the theory, or a distinct

calculus specifying the relation between the theory and experience, must include a subcalculus of a very different sort than that considered by the logical empiricists, *viz.*, one that allows *syntactic approximations* of mathematical expressions. This subcalculus of model reduction methods, which we could denote $\mathscr{K}_{reduction}$, could be implemented in a number of ways. Consider, for example, the case of an "inference rule" for linearization. One approach would be to include operations that replace a nonlinear functor with a formal series expansion in powers of $\mathbf{x}$, which would then be truncated at the linear term. Another approach would be to formally calculate the first two (constant and linear) terms of the Taylor series expansion of a functor. Thus, the ability to represent these operations syntactically requires a means of interpreting the mathematical meaning of functors syntactically and a special logical calculus $\mathscr{K}_{reduction}$ that allows syntactic transformations of functors. This is not something that the logical systems considered by the positivists can do, but this is exactly what modern computer algebra systems, like MAPLE, now do to perform symbolic calculations on symbolic expressions using mathematical operations.

Assuming we have obtained the model M' for the linearized simple pendulum model (1.5) and its calculus $\mathscr{K}_{M'}$, the solution to the initial value problem obtained by adding the conditions (1.4) can be derived rigorously in first order logic. All that remains, then, is to model the methods used to test the consequences of the model by experiment. In this case, all of the manipulations of the objects composing the different pendulums and all the measurement devices are directly observable and so can be described in the observation language. Actual measurements would then could be represented as equations specifying particular values of quantities, which taken together constitute a data model D. One more tricky aspect to handle is the statistical calculations needed, which includes not only calculations of mean and standard error but also regression analyses. It would not be difficult in principle, however, to specify yet another calculus $\mathscr{K}_{statistical}$, internal or external to the theory, that is capable of performing any needed statistical manipulation on data models D syntactically. Another tricky aspect to handle is error propagation for calculations involving quantities with error. Again, this could be handled with another calculus $\mathscr{K}_{error}$, which would specify symbolic manipulation rules to handle these computations.

We therefore see that a syntactic approach that moves considerably beyond that

of the logical empiricists could be capable of handling a great deal of the structure involved in feasible theory application. One interesting feature of such an approach is that it requires rules analogous to correspondence rules but that translate from observation to theory, which are required to accurately model the abstraction of constraints from considerations of the phenomenon being modeled. Another interesting feature of such an approach is that handling many of the kinds of methods used by applied mathematicians, including model reduction methods, statistics and error analysis, requires specialized logical calculi that are capable of handling symbol manipulation that goes beyond the deductive inferences syntactic approaches to theory application typically assume. It might be helpful to envisage these special symbolic calculi as *packages* or *libraries*, in the sense used in computer science, that can be added to a theoretical framework as needed in order to perform necessary analytic reasoning.[5]

This serves to show that the limitations or failures of the received view do not vitiate the syntactic approach to theories. As we have already mentioned, however, one of the significant limitations of purely syntactic approaches is that they are not capable of descriptive representation of the internal theoretical semantics of a theory, which in applied mathematics is usually some complex of geometric or topological structures and spaces. This is not an issue for contemporary axiomatic studies of particular theories, or even the more general C^* algebra framework for studying classical and quantum theories, since these approaches are able to make full use of model-theoretic concepts and techniques. The inability to represent the theoretical semantics was one of the major criticisms of the received view, however, and was an important starting point for later methodology avoidance strategies for studing scientific theories, specifically those that focused on representing the **models** of a theory rather than its *language*. It is to these approaches that we now turn.

[5] An analogy that may be useful here to clarify the relation between this sort of syntactic approach and a corresponding semantic one is a similar relation that obtains between classical and abstract analysis in mathematics. The semantic approach is analogous to classical analysis and the syntactic approach to abstract analysis. Once a theory in classical analysis is well-enough developed, an algebraic version of it can be formulated for the purposes of general or axiomatic reasoning. An example of this is the abstract definition of a derivation as an algebraic formulation of a differential operator. Following this sort of analogy, it would seem that the importance of a syntactic approach is that it enables clear general reasoning about stable structures found across a wide range of cases of theory application. In order to elucidate and develop descriptions of such structures, however, the semantic approach would be more useful. Once significant stable structures have been identified, they can then be formulated syntactically.

2.2 Semantic Approaches to Theories

Rather than focusing on the *syntax* of a theory, the semantic view focuses on its **semantics**. The semantic view of theories is another methodology avoidance strategy, but one that reconstructs a theory in set-theoretic or model-theoretic framework rather than the framework of a logical language, treating a theory in terms of sets of **structures**, in the sense of **models** of a logical theory. Thus, for applied mathematical theories, the typical semantic approach does not represent the semantics of a theory natively in the mathematical (typically geometric or topological) terms employed. This buys one generality, since it allows one to cover theories well outside the bounds of applied mathematics, but it does so at the expense of loss of descriptive detail. Such a native representation is captured well, however, by the *state-space approach* that was formalized by Beth (1948, 1949, 1961) and generalized by van Fraassen (1970). Of the logical approaches we consider in this chapter, this approach bears the closest similarity to the background framework we will use in this study. Before assessing the semantic approaches as strategies of methodology avoidance, let us consider the general philosophical outlook and motivation for the view.

The semantic view grounds the meaning of knowledge not in experience but rather in terms of the clear (model-theoretic) meaning of a theoretical structure, in such a way that knowledge is independent of any particular language or linguistic formulation. Unlike the received view, this approach is more open to metaphysics since, *e.g.*, it provides a basis for some forms of structural realism, but it is also compatible with scepticism or empiricism, such as van Fraassen's famous constructive empiricism. Rather than treating observation in terms of observation sentences, the semantic view generally treats observation and the empirical semantics of a theory in terms of the characterization of empirical data, *viz.*, **data structures** or **models** of data.[6] These are **structures** just like those of model theory, and they form the **models** of *theories* of data.

[6]These **data structures**, which are intended to pick out models of empirical data, are not to be confused with data structures from computer science, such as stacks, linked lists, tables, vectors, arrays, records, *etc.* Because **data structures** of the semantic view are **structures** in the sense of sets together with relations on them, computer science data structures are special cases of semantic view ones since instances of stacks, linked lists, *etc.*, are **models** of *theories* of stacks, linked lists, *etc.* The **data structures** of the semantic view, however, are intended to pick out models of empirical data, which can certainly include tables, vectors or arrays of measurements of values of a quantity but also includes, *e.g.*, smooth interpolants of discrete data than can be compared with computed solutions of differential equations.

Since many semantic approaches are realist or "quasi-realist", they often make heavy use of **mappings** that relate the structures of a theory to structures in experience or the world. And a feature shared with the syntactic view is that the aim of the general semantic view is to ground all knowledge in certain principles, except here they are axiomatized model-theoretically, in one manner or another, in the framework of axiomatic set theory.

The symbolic constructions that applied mathematicians work with in the context of mathematical modeling of natural phenomena are usually understood to have a specific mathematical meaning, *e.g.*, most contexts in which ODE and PDE are used in modeling the equations have a well-defined meaning in terms of geometric or topological spaces; and an understanding of this meaning plays an essential role in reasoning processes in mathematical modeling. So in addition to preserving the detailed symbolic structure of the sort discussed in the previous subsection, an adequate representation of feasible theory application must also capture the internal theoretical semantics of theories and mathematical models. It is not sufficient to capture the theoretical semantics, however, since an adequate model of theory application must also include an empirical semantics that specifies the relation between a scientific theory and experience or the world. For the theoretical semantics, the semantic approaches all use some set-theoretic or model-theoretic account. For the empirical semantics, however, there is much more of a range of approaches. We will consider briefly here the approaches of two prominent proponents of the semantic view, which are also relevant in certain respects to the approach we will adopt in this study. The first of these is the approach developed by Suppes, which takes the "sets of models" approach for the theoretical semantics, together with "data-structures" and "hierarchies of models" to specify the relationship of the theory to phenomena. The second is the formal approach developed by van Fraassen, which treats a theory as a "partially interpreted language", and the empirical semantics in terms of "loc-mappings".

2.2.1 Suppes's Sets of Models Approach

Suppes's semantic approach takes its roots from a combination of his work in axiomatic foundations of physics and his direct experience with empirical models and experimentation in the context of mathematical learning theory in psychological re-

search. This combined experience led him to seek an account of epistemology of science that included not only a precise account of theories, but also a precise account of the various types of *models* that are employed in the process of empirical research. Suppes (1961) distinguished between three kinds of models in science: meta-mathematical models, which are set-theoretic **structures** that satisfy the axioms of a theory; so-called "iconic models", which are stand-in representations for a complex phenomenon, such as scale-models or model airplanes; and so-called "**theory models**", which consist of the set of all structures that satisfy the axioms of a theory. Suppes noted that **structures** are simply members of **theory models** and then argued that the "linguistic models" used in mathematical physics, *i.e.*, what we are calling 'mathematical models', and the iconic models used in physics and engineering could be formalized as **theory models** or **structures**, respectively. He concluded, therefore, that the concept of *model* in science, and hence the theoretical semantics of theories, was capable of a clear foundation in terms of Tarski-style semantics.

Suppes (1962) then provided a basic account of how the relationship between theory and phenomena could be represented on such a semantic approach, which accomplished what correspondence rules did for the received view. Since Suppes recognized that the use of models was an essential part of how the analysis of experimental data is handled, and since he could construe all of the three kinds of models in terms of set-theoretic structures, he was able to extend a semantic view of theories to a semantic account of the relationship between theory and data and experiment. He did this by introducing a hierarchy of **theories** and **models** between a general theory and raw data. In his work on mathematical learning theory and statistical analysis of data he recognized that a high amount of abstraction was already needed to structure data into a form such that statistical analysis could be performed. From this example he abstracted a basic schema for how a hierarchy of **theories** and **models** can rigorously account for experimentation using mathematically well-defined scientific theories (see table 2.1). Each **theory** (set of **structures**) provides a framework for the general modalities at each level, with the **models** (**structures**) of each framework being potential realizations of the given **theory** at a given level of abstraction. Suppes supposed that this schematic example provides a strong case for how both models and experimental design could be accommodated rigorously in a semantic framework.

theory	model	typical problems
Scientific Theory	scientific model	estimation of parameters, goodness of fit to models of data
Experimental Structures	model of experiment	number of trials, choice of experimental parameters
Data Structures	model of data	constancy, independence, fit of experimental parameters
Experimental Design	experimental design model	randomization of subjects, choice of measurement devices
Experimental (Lab) Conditions	*ceteris paribus* conditions	noise, interference, physical perturbations

Table 2.1: Suppes's basic schema for a hierarchy of theories, models and problems in experimental scientific methodology, adapted and generalized from Suppes (1962, 259).

The definitive feature of Suppes' approach, then, is the understanding that a scientific theory can be defined semantically in terms of its set $\mathfrak{T}$ of **models**, such that in the context of working with a given scientific theory, the models employed by scientists, or appropriate semantic formalizations of them, are included within the set $\mathfrak{T}$. Since the semantics of a theory is defined in terms of a set of **models**, there is no need to appeal to experience to give the theory meaning, as the received view did. Rather, the connection to experience is formalized separately in terms of hierarchies of theories and models. The correlate in Suppes' view to the set of observation sentences $O_t(\mathbf{T})$ of the received view is a set of **data structures** $\mathfrak{D}_t(\mathfrak{T})$, or models of experimental data, which picks out the "phenomena" that the theory is supposed to describe, explain or predict. Observational adequacy is then defined in terms of every member of the set $\mathfrak{D}_t(\mathfrak{T})$ of data structures being embeddable via morphism in some structure $\mathfrak{S}$ of the theory $\mathfrak{T}$ (Muller, 2011, 92). Confirmation grows as the set of data structures $\mathfrak{D}_t(\mathfrak{T})$ grows while preserving embeddability in $\mathfrak{T}$, and $\mathfrak{T}$ is falsified if $\mathfrak{D}_t(\mathfrak{T})$ comes to have a member that cannot be embedded. This, of course, is viewed by Suppes as an idealization or abstraction of a real theory, but it is supposed that any given theory can be adequately formalized in this way by appropriately specifying the models of data obtained in observation or experiment in an appropriate hierarchy of theories and models.

Suppes' approach to formalization of a theory proceeds by representing all of the objects appearing in a theory, along with all the relations holding among them, set-theoretically. Such a formalization usually produces a canonical n-tuple specifying all this data. This then allows the definition of a *Suppes-predicate*, a set-theoretic predicate, which picks out as its extension all of the 𝔪𝔬𝔡𝔢𝔩𝔰 of the theory. In this way, Suppes relies on the rigorous foundations of set theory in order to specify theory in canonical form without the need for any metamathematical techniques. Suppes' slogan is that "*philosophy of science should use mathematics, and not metamathematics*" (van Fraassen, 1980, 65). This emphasizes the key practical feature of his approach, *viz.*, that it is not necessary to achieve full rigour by beginning axiomatically from first principles since the framework of axiomatic set theory provides full rigour in an informal way, the same informal way that actual mathematicians use all the time.[7]

2.2.2 Van Fraassen's Semi-Interpreted Languages Approach

The semantic approach developed by van Fraassen was also motivated by Tarski's work in model theory but indirectly through the work of Beth (1948, 1949, 1961), himself drawing upon the work of von Neumann, Birkhoff, Destouches and Weyl. This approach is a *state-space approach*, treating a theory fundamentally in terms of the space of possible states of a system and the measurable physical magnitudes that characterize the state of a system. Extending the work of Beth, van Fraassen (1970) developed a general framework intended to characterize the formal structure of nonrelativistic physical theories. Unlike Suppes's approach, which works informally with the notions of set theory in order to specify the set of 𝔪𝔬𝔡𝔢𝔩𝔰 of a theory, van Fraassen explicitly uses a model-theoretic approach to characterize the semantics of theories. We will see that van Fraassen's formal semantic approach bears the closest similarity to the basic framework we will use in this study.

A theory is characterized, on van Fraassen's formal approach, in terms of three fundamental components. One is the *state-space* of a given kind of physical system.

[7] This argument will be important for our purposes, since the constraint approach that I use in this study will make similar use of well-defined mathematical concepts in the interests of modeling theory application.

This includes the system (position, configuration, phase) spaces described in the previous chapter, as well as the Hilbert spaces used in quantum theory. The second component is a certain set of *measurable physical quantities* provided by the theory, which are capable of characterizing a physical system, *e.g.*, position and momentum in classical mechanics. In van Fraassen's picture, this then determines the set of *elementary statements* of the theory. These are statements U that express the proposition that a given physical magnitude m has a certain value r at a certain time t, schematically written, *e.g.*, $U(m, r, t)$. The third component is a function $s(U)$ that specifies a region of the state-space over which U is *satisfied*. This function is called the *satisfaction function*. To illustrate, in the case of the 6-dimensional Euclidean phase space for a classical point particle, the statement $U \Leftrightarrow$ "The x-component p_x of momentum is r" is satisfied over the 5-dimensional hyperplane defined by $p_x = r$. So in this case, $s(U)$ is that 5-dimensional hyperplane. A given statement U about a physical system is then *true* provided that the system's actual state is represented by an element of $s(U)$.

Thus, on van Fraassen's formal approach a theory is understood to specify a set E of elementary statements, a state-space S and a satisfaction function s, which maps the set E of elementary statements to some set of subsets of S.[8] This triple $L = \langle E, S, s \rangle$ is then called a *semi-interpreted language*, since the language E of L is provided with a certain partial interpretation in terms of the satisfaction function $s(U)$. It is partial, because which statements are actually true depends on the actual state of a physical system. In this way, van Fraassen's formal approach specifies a definite theoretical semantics in terms of the state-space S, which provides via $s(U)$ a semi-interpretation for the language E that specifies the possible statements U about a given physical system. Thus, van Fraassen's formal approach provides a limited representation of the symbolic constructions of a physical theory in terms of the elementary sentences E and relates them explicitly to a phase-space semantics via the satisfaction function.

A semi-interpreted language is also given a formal empirical semantics. A *model* for a semi-interpreted language L is a couple $\langle \mathrm{loc}, X \rangle$, where X is a system of the type specified by L and loc is a function that assigns a location to X in S. What X

[8] Typically this set of subsets would be the Borel algebra defined from the topology on S.

is intended to be is not made clear, but it plays a role quite analogous to that of a data model in Suppes's framework, *viz.*, an abstraction from a real physical system that correlates measurement data from experiment with what can be interpreted via the function loc as a state or trajectory in S. This formal empirical semantics is given a truth definition by defining an elementary statement U of L to be *true in* $M = \langle \text{loc}, X \rangle$ iff $\text{loc}(X) \in s(U)$. A set X of statements is true in a model M just in case each member is true in M. This then allows the following:

Definition 2.2.1 (Semantic Entailment) A set $X \subseteq E$ *semantically entails* U in L, written $X \Vdash U$, iff U is true in every model M of L where X is true.

The (semantically) valid sentences then come out as a special degenerate case where U is true in every model of L. The *logic* of a semi-interpreted language is then "essentially a syntactic description of the set of valid sentences and the semantic entailment relation in that language" (van Fraassen, 1970, 335).

The picture of a theory, then, for van Fraassen, at least in the context of physics, is that of Beth, *viz.*, that the postulates of a theory serve to provide "the description of a state-space together with a mapping correlating the state-space with elementary statements about measurable physical magnitudes" (van Fraassen, 1970, 337). The formal description above provides a general formulation of this idea for non-relativistic physical theories. This serves to provide both a theoretical semantics, in terms of the state-space, and an empirical semantics, in terms of models construed as mappings from physical systems (or some abstract representation of them), to locations in the state-space. This faithfully represents the basic theoretical framework of a dynamical system and a specific mathematical model within it that describes the behaviour of a physical system, specified by the model $\langle \text{loc}, X \rangle$ of the semi-interpreted language $L = \langle E, S, s \rangle$. This approach also preserves the natural description of dynamical systems in terms of state spaces and the manner in which trajectories represent the dynamical behaviour of physical systems in the world.

This provides a technical formulation underlying van Fraassen's semantic approach. In his description of his constructive empiricism, van Fraassen (1980) provides an intuitive version of this approach. This is specified as follows (for the corresponding version in terms of semi-interpreted languages see table 2.2). For van Fraassen, just as for Suppes, to present a theory is to specify a set of **structures**. Certain parts

scientific terminology	**constructive empiricism**	**semi-interpreted languages**
theory	set of $\mathfrak{structures}$	set of languages $L = \langle E, S, s \rangle$
model	$\mathfrak{structure}$	semi-interpreted language L
theoretical description of physical system	empirical substructure of a $\mathfrak{structure}$	trajectory or region of state-space S
instance of theory-data relation	isomorphism of appearance to empirical substructure of a $\mathfrak{structure}$	model $M = \langle \text{loc}, X \rangle$
experiential phenomena	appearances	descriptions X

Table 2.2: Roughly equivalent interpretations of scientific concepts as construed in van Fraassen's constructive empiricism (van Fraassen, 1980) and semi-interpreted languages (van Fraassen, 1970) semantic approaches to the theory-world relation.

of these $\mathfrak{structures}$, called the *empirical substructures*, are to be specified as candidates for the direct representation of observable phenomena, which would include trajectories or regions of a system space. Structures that can be described in experimental and measurement reports he calls *appearances*. The appearances, then, are the observed "phenomena" that the theory is supposed to describe, explain or predict. *Observational* adequacy (not to be confused with empirical adequacy) is then defined in terms of a theory $\mathfrak{T}$ having a structure $\mathfrak{S} \in \mathfrak{T}$ such that every member of the set $\mathfrak{A}_t(\mathfrak{T})$ of appearances, the correlate of the observation sentences $O_t(T)$ of the received view, is embeddable via isomorphism to empirical substructures of $\mathfrak{S}$. Confirmation then grows as the set of appearances $\mathfrak{A}_t(\mathfrak{T})$ grows while preserving isomorphic embeddability in some structure $\mathfrak{S}$, and $\mathfrak{T}$ is falsified if $\mathfrak{A}_t(\mathfrak{T})$ comes to have a member that cannot be isomorphically embedded.[9]

Van Fraassen's concept of *empirical* adequacy is stronger than the concept of observational adequacy. A theory $\mathfrak{T}$ is *empirically adequate* if the theory has some model $\mathfrak{M} \in \mathfrak{T}$ such that *all* appearances (within the domain of applicability of the

[9]This view is described by van Fraassen as a 'picture', *viz.*, "something to guide the imagination as we go along" (van Fraassen, 1980, 64), which he recognizes is a very limited description of the relationship between theory and data. The assumption, then, is that, as far as it goes, any given theory can (in principle) be adequately formalized in this way by appropriately specifying the appearances it is meant to capture and their relation to some model of the theory.

theory) are isomorphic to empirical substructures of that model. Empirical adequacy requires not only that the theory captures all appearances that have actually been observed, but also all appearances there have been and any possible appearances in the future. We could denote this set of appearances $\mathfrak{A}_\infty(\mathfrak{T})$, in which case a theory is empirically adequate provided the theory has some model such that each member of $\mathfrak{A}_\infty(\mathfrak{T})$ is isomorphic to an empirical substructure of that model.

Assessing van Fraassen's approach as a methodology avoidance strategy helps to bring into focus some of the main issues in this chapter. Recall the adequacy conditions (A) and (B) on a reconstruction of a theory and its relation to the world considered in the previous section, which pertained, respectively, to preserving essential content and eliminating only inessential content in an abstraction from actual science. The semi-interpreted languages approach is seen to preserve a great deal of the structure of dynamical system models, which are common in theoretical physics, including the state space, dynamics on that space and that part of the symbolic language that specifies states. There is clearly a great deal that is not represented here, however, including the vector field on the state-space, the equations that specify the dynamics, any consideration of solution of equations, and any methods and models used to construct the given model. This is certainly not epistemologically insignificant content but neither is it clearly essential for the purposes of the 1970 paper, which were, *inter alia*, to provide a picture of a physical theory that was in line with work in foundations of physics. The interest is thus a limited task in classical epistemology, *viz.*, the clarification of the foundational structure of physical theories. The structure eliminated is not clearly relevant for this task. Given this limited aim, van Fraassen's approach meets the adequacy conditions.

Another aim of van Fraassen in the 1970 paper, however, is to capture is the relation between such a mathematical model and a physical system, which is done in terms of the "model" $\langle \mathrm{loc}, X \rangle$. This essentially provides an embedding of a description of the dynamics of a physical phenomenon into the mathematical framework of the language L. So, this captures in some way the relation between the measurable quantities of a physical system and the mathematical framework, but it does so without any consideration of how the description X of the system is specified. This serves the classical purposes of picking out a relation between the phenomenon and

the mathematical model that must obtain, in a manner similar to Giere's notions of "interpretation" and "indentification", but it obscures the detailed nature of that relation and the limitations on when such a relation actually obtains. It is here where methodology avoidance turns into *epistemology avoidance*, *i.e.*, eliminating complex details about what, and how, knowledge is gained or provided by a scientific theory. Of course van Fraassen's approach is not intended to characterize precisely the knowledge state-space models provide about physical phenomena, but it is intended to clarify the nature of the knowledge about phenomena provided by physical theories. Whether van Fraassen's approach is adequate for this purpose, even from the point of view of a classical general epistemology of scientific theories is less clear than is its adequacy as a methodology avoidance strategy.

It is useful to note here that included in the structure eliminated in the interests of classical epistemology is structure that is essential to understanding the feasible epistemology of physical theories. This includes the relationship to measurement devices, any consideration of experimental questions and conditions, statistical analysis and error analysis, but also all of the mathematical methods that are essential to account for the *process* of theory application. The semi-interpreted languages approach provides a snapshot of a particular mathematical model and a schematic connection to the world; as a result, it does not represent the process of constructing such a model, the relationships between the models used in the construction process, including those from other theories, or the complex relationship between these models and actual experimental situations. This serves both to highlight the manner in which the interests of classical and feasible epistemology are complementary *and* just how much methodologically and epistemologically significant structure classical accounts leave out.

2.2.3 Limitations of the Semantic View

Setting aside any possible limitations of the two semantic approaches we have considered for the classical purposes for which they were intended, we will now consider their capacity for gaining insight into feasible epistemology. We can see, just from these two cases, that the semantic approach, in general, does a much better job of preserving various structural features of a theory and its relation to the world than the

received view does. In particular, the semantic approach is able to capture directly the intended meaning of equations, which then enables representation of phenomena to be understood non-linguistically. Indeed, one of the commonly cited advantages of this approach is that it captures the theoretical semantics without being tied to any particular formulation in a logical language. This implies a deemphasis of the symbolic presentation of a scientific theory. Although there are clear advantages to this in terms of capturing theoretical semantics, particularly over the received view, it is nevertheless the case that the semantic view tends to describe the linguistic formulation of a theory only as it arises in a Suppes-style axiomatization or something essentially equivalent.

Although this axiomatization strategy has its advantages for classical epistemology, it presents a significant problem from the point of view of descriptive methodology because it loses direct description of what scientists, specifically applied mathematicians, typically deal with most directly, *viz.*, symbolic constructions and the methods of symbol manipulation that they use to elucidate the behaviour of natural phenomena. An example of this is the technique of linearization we have already considered. Such an operation could be represented as a transformation from a $\mathfrak{model}$ in one theory-structure to a $\mathfrak{model}$ in another, or between two $\mathfrak{models}$ in a single overarching theory-structure, which can capture the geometric operation, but the algorithmic nature of the method that effects this shift is lost entirely. Given the importance of such methods in the model construction process in applied mathematics, this is a significant limitation for feasible epistemology. What this demonstrates is that the semantic view as a whole is a methodology avoidance strategy, meaning that it abstracts out only certain content from a full description of the tools and methods of scientific practice; it is not a general framework which faithfully accommodates all aspects of theories and their application, within which particular approaches formulate abstract descriptions of particular aspects. We require such a general framework for the purposes of descriptive methodology, however, since the framework we work within should not decide *a priori* which aspects of theories and their application are methodologically or epistemologically significant.

To illustrate the limitations of the semantic approach as a strategy for focusing on the most epistemologically significant aspects of the theory-world relation, consider

that even to the present the semantic approach lacks a proper treatment of computation. Numerical computation, involving the running of algorithms on a computing device in order to solve mathematical problems, is an increasingly common part of the application process in science. And, for reasons we will discuss in chapter 5, far from being a "merely pragmatic" step of the application process needed to get numbers out of models, which it how computation has historically been considered, it is an *essential* part of the application process in increasingly many cases. From a semantic point of view, the most significant feature of using machines to solve mathematical problems is that the continuous mathematical spaces of mathematical models must be translated into a form that is both discrete and bounded, so that problems can be posed in a form that a computer can solve. In particular, as we shall discuss in chapter 5, this involves representing the continuum of real numbers in terms of a bounded finite number of floating point numbers. This translation process is not something that has been treated by a semantic approach. Indeed, Suppes (2011) acknowledges this limitation, but his solution is to advocate axiomatization of the floating point number system. Though this could help to elucidate general features of floating point arithmetic, such a highly abstract description will not clarify the the extent and reliability of the kinds of knowledge made available by the actual procedures and algorithms used by applied mathematicians, particularly with regard to the nature of the kinds of error involved. We will get a sense in chapters 4 and 5 of just how much structure this approach is eliminating from consideration.

In order to faithfully represent methods used in practice, the semantic approach would have to augment their theory-structures with a particular, not generic, linguistic formulation, along with algorithms that can be performed on such descriptions. And these algorithms must be understood as transformations of, or between, theory-structures, with specific consequences for the image of a **𝔰𝔱𝔯𝔲𝔠𝔱𝔲𝔯𝔢** of the domain theory-structure in the codomain theory-structure. Adding all of this structure at the syntactic level, however, requires a characterization independently of the set of **𝔰𝔱𝔯𝔲𝔠𝔱𝔲𝔯𝔢𝔰** that is supposed to characterize a theory.[10] So, theories that require syntactic algorithms of this sort, push beyond the Suppesian idea of axiomatizing a theory

[10]There may be a way for the semantic approach to continue to treat the language of a theory generically and have such algorithms operate on any instance of the generic language, but such an approach still pushes beyond the ability to characterize a theory simply by specifying its models.

by constructing a predicate. And, even if the semantic approach can be augmented to incorporate reduction methods, it is stretching the meaning of "semantic" to call this a semantic approach. I argue that it would be more appropriate to consider such an approach syntactico-semantic, since *both* the semantics and the syntax are being represented abstractly. Given that these problems show that the semantic approach is too strong a methodology avoidance strategy from the point of view of capturing the structure of feasible theory application, I argue that that logical approaches to modeling feasible theory application need to be syntactico-semantic in the sort of manner indicated here. I will develop this argument further in the next section.

Now, though the semantic approach can *in principle* be extended to provide adequate descriptions of actual methods in mathematical modeling, the manner in which the semantic approach is usually employed has much deeper problems. Theories are rarely axiomatized directly by specifying a set-theoretic predicate for each theory-structure in a hierarchy of models. Rather, they are typically discussed in a generic way in terms of all the structures $\mathfrak{S}$ of a theory $\mathfrak{T}$, or in terms of some generic hierarchy of such theory-structures. This is the semantic version of the theory T syndrome. And just as for the syntactic approach, what is bought in generality is paid for in loss of specificity. Treating theories in this way allows one to talk about theories and data or phenomena in general, and the relationships among theories, data and phenomena, but very easily leads to using a description that outstrips its ability to make adequate contact with actual science and the kinds of knowledge it actually provides. Consequently, this general model of scientific theories and their relation to the world leads to false accounts of science, as I will argue throughout this study.

The tendency to employ logical tools in a generic, fully general way reveals one of the most significant limitations of the logical approaches to representing theories: they have a tendency toward being methodology avoidance strategies in the pejorative sense of avoiding ever having to deal with the details of scientific inference in real science. They provide the epistemologist with a generic template for how to represent a theory *a priori*, no matter what area of science it is from or how much that area relies on the use of mathematics. And in many cases, such logical approaches are considered to reveal the *essence* of a theory, a presentation of the true characteristic features of a theory. In such cases these approaches, in effect, dictate to science

that this is how theories are structured. When we recall that scientific method is the phenomenon that epistemologists, *qua* methodologists, of science are studying, this would be like physicists dictating the structure of nature. The character of this approach is decidedly unscientific. This is not a condemnation of general theories, since there are plenty of general theories in science and there are useful general theories in epistemology, such as the C^* algebra approach in foundations of physics. Such theories, however, are the result of enormous intellectual and experimental effort, ingenuity and time. One does not just come up with general theories of natural phenomena. One must first be able to track the phenomena one wishes to understand and, over time, find ways of zeroing in on what is fundamental. But this is something that is discovered through active investigation and research, not something that is decided *a priori*.

The problem I am identifying here is not a problem with logical approaches *per se* but rather with what they are often used to do. Provided that an approach to reconstructing or representing some aspect of science meets an appropriate form of the adequacy conditions (A) and (B) above, and it is only used for purposes where it captures sufficient structure to be meaningfully applicable to that part of science it is concerned with, logical approaches are fine as far as they go. Conditions of applicability, or the limitations of applicability, of epistemological models are rarely made clear, however, which makes it very easy to use such a model for purposes beyond those for which it provides accurate and meaningful information about science and the knowledge it yields. Moreover, it can be difficult to notice when an approach has become inaccurate or meaningless. The basic problem, then, is not that logical approaches fail to be descriptively accurate or that they are incomplete, but rather than their limitations as tools for elucidating the epistemology of science are in many cases poorly understood.

A central reason for this state of affairs, I argue, is that logical approaches tend to be employed for the purposes of studies in classical epistemology, which uses pure mathematics as a model for scientific epistemology. Given the focus on in principle questions and problems, it is often excusable to have an approach that is highly abstracted from real science, which can make it difficult to compare descriptions with scientific practice or even make comparison to practice seem irrelevant. Logical

approaches to the epistemology of science are often proposed, shown to capture some structure of scientific methods or knowledge, and sometimes shown to solve certain philosophical problems, but they are rarely applied to even simple scientific cases in detail, cases that are supposed to be included under the purview of the approach. We saw in the previous chapter that both syntactic and semantic approaches have difficulty adequately representing the structural features of the detailed treatment of modeling and experiment for a simple pendulum, one of the simplest mathematical modeling tasks imaginable. But, again, this is because these approaches are not typically concerned with studies of feasible epistemology. One of the aims of the present study, however, is to show that not only are there epistemological gains to be made by studying feasibility but that an epistemology grounded in considerations of feasibility can build toward more general epistemological models and theories that are fully consistent with scientific practice.

A successful general scientific theory of epistemology of science ought to be able to be applied to real theories in order to gain insight into important features of their structure and behaviour. From this point of view, then, the scientific value of a general account of science or some area of it that cannot faithfully recover even the simplest examples of its practice is clearly questionable. I argue that the correlate of observation and experiment in scientific epistemology is testing methodological and epistemological models by seeing whether they recover the structure and behaviour of scientific method or knowledge that they are intended to capture. This requires constructing representations of actual methods and knowledge and seeing whether they agree. The aim of the present study is not directly such testing, however, since the aim is to *develop* some initial epistemological models that can then be tested on cases for which they were not constructed to describe. Thus, the present study aims to prepare for scientific testing of this kind and not to initiate it.

The ability to take a general theory and apply it to phenomena in a feasible way that simultaneously provides insight into the phenomenon is precisely what the methods of applied mathematics are tuned to accomplish for general scientific theories formulated in a mathematical language. There is indeed a tension between the search for general models or theories that can be used for philosophical purposes and the search for local models or theories that faithfully capture actual phenomena. But,

again, the models and methods of applied mathematics make it feasible to overcome this tension in real science—they make it feasible to relate general theories to descriptions of real phenomena. The failure of formal approaches in epistemology of science to pursue general models or theories that can faithfully capture actual epistemological phenomena is to pursue an *a priori* strategy that is problematic in many of the same ways for which Francis Bacon criticized Aristotelian metaphysics. Furthermore, if a general epistemology of science cannot feasibly give us insight into real science, then its claims to full generality appear dubious. Consequently, I argue that formal methods in philosophy of science need to adopt a strategy where they develop models and theories that are *shown* to be valid in practice. At least for the moment, such an approach restricts us to limited and more modest models and theories in terms of scope. The picture here is to aim toward a descriptively adequate general theory by first developing a number of stably accurate models covering a reasonably wide range of some area of science. This could then initiate the search for unification, leading to more and more general, but still descriptively accurate, epistemological theories.

Thus, I argue, the development of an adequate formal, general scientific epistemology requires the ability to feasibly translate between descriptions in the general theory and accurate descriptions of scientific practice. The ability to do this requires a formal approach that, *inter alia*, can generate accurate descriptions of the structure and behaviour of scientific methods used to apply theories to phenomena. The purpose of this study is to explore one approach to doing just this in a limited scientific context through the abstraction of patterns from the detailed consideration of the case examples described in the last chapter. Before beginning our consideration of the structure and behaviour of scientific methods and knowledge revealed in the detailed consideration of the case examples, we will consider briefly in the next section some of the general features of our approach. A more detailed description of the foundations of our approach is presented at the beginning of the following chapter.

2.3 Feasible Approaches to Theories

In order understand feasible epistemology and methodology, we need to have a way of picking out the structure or behaviour of scientific methods that pertains to whichever

particular epistemological and methodological questions we might have. To do this we need a way of reducing the full complexity of knowledge and method in practice along particular lines, reducing away irrelevant structure in order to bring out the structural and behavioural features that answer our questions. For example, suppose that we wanted to describe the central features of the structure of methods used to construct and run the LHC near Geneva, where by 'central' I mean those features that are most important for the methods to be reliable and efficient. To do this efficiently we would not want to use set-theoretic structures and axiomatics, we would want ways to simplify the enormous complexity of method in order to pick out dominant structural features that determine and explain the efficiency and reliability of the dominant procedures used in its construction and the dominant methods that allow it to run effectively. Or one may want to know how and why the structures of theory application in physics and chemistry are similar and how and why they are different, or even just the most important structural differences between post-newtonian mechanics and general relativity from the point of view of both feasible and classical applicability. The ability to answer these sorts of questions in a systematic way requires new epistemological methods. In order to develop methods that can answer such questions effectively and efficiently, I argue that we must construct methodology avoidance strategies that are modeled on the methods of applied mathematics.

It is the aim of this study to take some small steps toward a feasible account of the role of mathematics in science by constructing a model of the application of real theories in a very limited area of applied mathematics. Since any scientific theory requires some characterization of the phenomenon that it must recover, a phenomenology of scientific knowledge and method is required to get this process going. This is the reason that the study is based upon the two case studies, double pendulum and near-Earth object modeling, mentioned in the introductory chapter. And since physics avoidance strategies work by operating on a complex description, reducing its structure while focusing in on dominant features, we will attempt, in a similar way, to reduce the complexity of the case studies in ways that reveal dominant structural features of the process of application of real theories in applied mathematics. We seek strategies that can be applied as techniques for the reduction of epistemological complexity, which reduce complexity and variability in knowledge or method that

does not pertain to the given epistemological or methodological questions at issue. The result of the limited investigations of this study will be a *local* model of theory application in applied mathematics. The model is local in quite a complicated sense, since different components of the model have different expected ranges of applicability to scientific practice. That the model is multiply local in this way reflects its *modularity*, *viz.*, it is designed in such a way that the aspects of it that are more specific to the structure of the case studies, and hence likely to be less widely applicable, can be detached without affecting the rest of the local model. The different components of the model can then be applied in different contexts, and where they fail to capture structure or behaviour appropriately, they can be extended, adapted or modified in order to become consistent with method in that context. In this way, models of increasing generality can be built whilst maintaining descriptive accuracy.

This provides, in a nutshell, the character of the approach feasible epistemology that we will develop throughout the study. Other feasible epistemologies could have very different nature and strucure, being tuned for feasible use in quite different scientific contexts. But all such approaches will have a certain number of common features. One feature that is likely to be quite stable, for reasons that I have provided in this chapter, is a modeling approach that is capable of faithfully representing both the syntax and the semantics of theories and their relation to the world.

2.3.1 Logico-Structural Approaches to Theories

Since we are seeking useful formal tools for modeling feasible knowledge and methods in science, it is natural that logic, pure mathematics and applied mathematics would be useful places to look for tools that can be adapted to the task. As we will see throughout the study, some of the tools from the syntactic and semantic approaches in epistemology of science can be usefully easily adapted to the task of a representation or analysis of feasible knowledge and methods. As I have argued throughout this chapter, however, the syntactic and semantic approaches on their own are incomplete. On the most basic level, purely syntactic approaches are incapable of working in an unrestricted way with sets of models of theories and purely semantic approaches are incapable of working in an unrestricted way with accurate representations of mathematical modeling methods. The aims and purposes of feasible epistemology

demand a background theoretical framework that is *in principle* capable of faithfully providing fully detailed descriptions of scientific theories and their application to gain knowledge about phenomena, which includes forms of scientific knowledge and the methods used to obtain them. Accordingly, the background framework needed for feasible epistemology must combine a faithful representation of the semantics of theories and their application with a faithful representation of their syntax. For a number of reasons, however, this is not simply a combination of syntactic and semantic approaches.

As we saw in section 2.1.3, a syntactic approach that is capable of a faithful representation of mathematical modeling methods is not an axiomatic theory in the usual sense. A more accurate description of such an approach would be a set of symbolic manipulation techniques that represent the methods used to transform between mathematical models of, potentially distinct, theories involved in the application process. And this is just a representation of symbolic manipulation in a very simple example of mathematical modeling. More complex scientific examples would require considerably more complex syntactic modeling, which means that it is not clear *a priori* what kind of syntactic approach will be required to faithfully capture the manipulation of symbolic constructions as it occurs in real science.

Adding on a set-theoretic or mathematical semantics to this in the manner required by feasible epistemology is more involved than simply interpreting theories as sets of structures. The reason for this is that mathematical techniques used to modify the algebraic form of a model or theory generally have a quite specific effect on the mathematical interpretation of a theory or model. An example of this we have used repeatedly is linearization, but another example of this are asymptotic methods used to study the relations between different theories, such as those discussed by Batterman (2002*b*). To represent these methods faithfully it is necessary to describe the *correlations* between symbolic constructions used in applied mathematics and their mathematical meaning. In particular, it is necessary to describe the *covariation* of symbolic constructions and their meaning when an algebraic manipulation technique, such as linearization, is applied to a theory or model. Moreover, methods in applied mathematics rely heavily on perturbation methods and approximations, which involves being able to handle variable or varying descriptions. In order to represent

such methods faithfully, it is necessary to have conceptual resources for representing dependencies between syntactic and semantic variations. This allows the faithful representation of both syntactic and semantic reasoning—where the syntax and semantics have clear correspondences, semantic reasoning can be tracked by variations in syntax and the semantic consequences of syntactic reasoning can be clearly traced.

We will use the term *logico-structural* to refer to background frameworks for feasible epistemology that capture the interrelated nature of scientific *language* and the structures that interpret them. This is the general kind of framework that I argue is needed to correct for the limitations of the logical approaches from the point of view of (naturally) admitting descriptions of the application process in applied mathematics that are accurate in full detail. It is not, however, to be assumed that such a background framework will always be sufficient in feasible epistemology, since future investigations may require frameworks with a more complicated or simply different structure. This is a major difference between feasible and classical epistemology, since classical epistemology typically works with some tacit assumption that all cases will fit into a particular framework in principle. For feasible epistemology we do not just require that representations are possible in principle, since we also require that it is possible to use a framework to feasibly gain the desired or required epistemological insight.

For the purposes of this study it is not necessary to specify such a logico-structural framework using anything approaching full first order logic, let alone higher order logics. Because the mathematical models that we are concerned with in the case studies are picked out by equations, we can restrict attention to a framework modeled on the equational theories of universal algebra, rather than the quantificational theories of formal logic. Since the semantics used in applied mathematics is typically given in terms of geometric or topological spaces, we can assume a semantics with a great deal of geometric or topological structure, and in a much broader and less constrained way than used in the state-space approach. This allows us to treat models and theories in terms of equational algebraic constraints on the geometric or topological interpretation of symbols. A system with a correlated syntax and semantics of this sort may be called *algebro-geometric*. And it is a limited system of this simpler sort that will be used or assumed for the purposes of the study, which we will introduce in outline

in the next chapter. Since for the purposes of this study it is rarely necessary to deal with such a framework in detail. Accordingly, we specify the details of such a framework allowing it to be assumed as a background and any required details referred to as required. We will see in the next chapter that this framework supports a natural high level system of logical concepts, which will be useful for conceptual clarification throughout the study.

Now that we have a limited motivation for the algebro-geometric modeling strategy that is used throughout the study, we can turn to the consideration of how this system can clarify various important structural features of the application of theories in applied mathematics that are revealed through the detailed consideration of the case examples.

Part II

Modeling Feasible Theory Application

Chapter 3

Modeling in the Real World: Feasibly Constructing Models

3.1 A Feasible Approach to Theories and Models

Before we begin our consideration of some important structural features that appear in the feasible construction and analysis of mathematical models, we will first consider the basic structure of the formal approach of methodology and epistemology avoidance used in this study. As I have argued in the previous chapter, some of the most important structural features of the process of constructing models are not captured faithfully by either the syntactic view ($\mathscr{L}$-view) or the semantic view ($\mathfrak{S}$-view).[1] I argued that in order to capture these features faithfully, it is necessary to develop a new framework that combines a syntactic representation with an semantic one. This leads to a representation of a theory as a symbolic language with a correlated complex of mathematical structures forming its models. Such an approach is not simply a pairing of logical syntax with set-theoretic semantics in the usual way. This approach requires a language that models the syntax of a theory, its semantics *and* the complex relationships between the two, which includes mathematical methods used to vary the symbolic form and mathematical meaning of a theory or mathematical model. Accordingly, it is not intended to reveal the "essence" of a theory; rather it is intended to elucidate the structure of a theory and its methods, and the relation of a theory to phenomena through its models. As mentioned in the previous chapter, we will call an approach of this sort a *logico-structural* ($\mathcal{L}$-$\mathcal{S}$) such as on theories and

[1] We can think of $\mathscr{L}$ as *language* or a set of *sentences* and $\mathfrak{S}$ as a set of model-theoretic *structures*, capturing the essence of how theories are viewed in the two logical approaches.

their relation to phenomena.

To be more clear about what I mean here, consider a classical semantic account of a theory in terms of its set of $\mathfrak{models}$. Such an account is typically intended to tell us what a theory *is* by specifying its structure; it is not simply intended to specify formally the theoretical semantics of a theory. The sort of $\mathcal{L}$-$\mathcal{S}$ perspective that I am describing here is entirely compatible with such a classical account of the structure of a theory. In this case, an $\mathcal{L}$-$\mathcal{S}$ perspective represents how applied mathematicians reason using a theory to construct models of phenomena and does not provide a competing account of the structure of a theory. However, the $\mathcal{L}$-$\mathcal{S}$ approach does not require a theory to be defined classically in this sort of manner. Indeed, someone could take it to provide an alternative account of the structure of a theory based on the manner in which theories are actually used in applied mathematics. But this is not actually the purpose of the $\mathcal{L}$-$\mathcal{S}$ approach. Its purpose is to provide a means of *representing* the actual reasoning processes that applied mathematicians use symbolic constructions and mathematical structures for, and in particular to represent the actual methods used to construct, manipulate and apply mathematical models of phenomena. Thus, the $\mathcal{L}$-$\mathcal{S}$ approach really is intended as a *modeling approach* for representing the practice of applied mathematicians, not as a competing account of the structure of theories. This is why I describe instances of it in terms of a "perspective" rather than a "view", given the connotations the latter has in the context of the logical approaches.

Now, in order to represent the model reduction methods used throughout applied mathematics, the language will not just be an ordinary logical language with a vocabulary, formula formation rules, axioms and deductive rules of inference, it must also include sets of auxiliary rules of syntactic transformation that transform one linguistic description into another, which for reduction methods will be one that is analytically or computationally more tractable. We discussed this sort of syntactic transformation in the last chapter in the context of an up-to-date syntactic approach. In order to represent the natural, or intended, semantics of a theory, the set of structures forming the models of a theory will not be arbitrary sets; rather they will be sets with a well-defined *mathematical* structure. Thus, rather than describing the semantics of a theory in terms of the concepts associated with set-theoretic structures,

the semantics may be described in terms of the geometric or topological spaces that interpret the symbols of a theory in applied mathematics. But an $\mathcal{L}$-$\mathcal{S}$ perspective of a theory is not just a simple pairing of a language and its semantics in this way; there are also correlations between the syntax and the semantics that must be specified. In case of linearization, the method applied to a differential equation has specific consequences for the dynamics on the system space. Consequently, as we saw in the last chapter, a representation of linearization of a function must be a(n auxiliary) syntactic rule of transformation on a function symbol together with a specific, well-defined transformation of the function interpreting the symbol.

Given that the mathematical models we are considering are constructed by forming equations relating various quantities and functions, we do not need to work with an $\mathcal{L}$-$\mathcal{S}$ perspective that treats the language in terms of first or higher order logic, we can just work with equational languages. This restricts attention from the representation of general logical theories to algebraic theories. And given that the mathematical models we are considering are interpreted in terms of geometric or topological spaces, we do not need to work with an $\mathcal{L}$-$\mathcal{S}$ perspective that treats the semantics in terms of arbitrary set-structures, we can just work with a semantics of geometric or topological spaces. As we have said, this restricts attention from a representation in terms of model-theoretic semantics in set-structures to representation in terms of a semantics of geometric (or topological) spaces. Consequently, the variety of $\mathcal{L}$-$\mathcal{S}$ perspective we will restrict attention to is an *algebro-geometric* perspective ($\mathcal{A}$-$\mathcal{G}$ perspective).[2]

In order to make the methodology avoidance strategy we use effective for feasible modeling of feasible theory application, we will not treat theories and their semantics in terms native to universal algebra. Rather, we will develop a model of theories which is tailored to bring out salient structure in the theory application process in applied mathematics. In this section we will present the basic structure of the $\mathcal{A}$-

[2]One might wonder why we refer to such a perspective as algebro-geometric rather than algebro-topological, given that topological spaces form a more general category than geometric spaces. We do this for a couple of reasons. One is that the representation of phenomena is generally metrical and not only topological. Accordingly, topological spaces in this context are to be viewed as abstractions from geometric spaces. So, from this point of view, an algebro-topological perspective would be derivative, or an abstraction, from an algebro-geometric perspective. The other reason is that geometry is more suggestive of a semantics in experience or the world than is topology, which is again more abstract. Thus, from this heuristic point of view, 'geometry' is the more appropriate term.

$\mathcal{G}$ approach underlying the feasible modeling strategy we adopt for the study. We will also introduce the high level logical presentation of this approach that we develop through this study and use for the purposes of methodology avoidance, adding additional structure as it is needed in later chapters.

3.1.1 Algebro-Geometric Approaches to Modeling Theories

The aim here is to develop an approach, grounded in some kind of formalism, that is capable of capturing all of the relevant methodological and epistemological structure that is screened out by the typical logical approaches yet has the kind of flexibility that will allow simplified presentations of this structure that promote conceptual clarification and that allow it to be adapted to new contexts. This sounds like a tall order, but the flexibility of mathematical methods provides us with some initial tools that are able to accomplish this for present purposes. Let us consider then how this may be done.

Let us recap some of the kinds of structure we need to preserve in a faithful model of mathematical model construction. One of the most obvious features of theories in applied mathematics is that they take on a specific symbolic form and they have a particular geometric or topological meaning. Accordingly, we need the ability to preserve both the particular symbolic form as well as the geometric or topological meaning. Also, model reduction techniques used in applied mathematics are forms of approximation that reduce structure and yield more tractable models. Accordingly, we need the ability to track *inference* in such a way that allows for not only deduction but also approximation and that can track the nature of the structure reduction that is accomplished by these methods. Thus, we require a modeling approach that preserves the symbolic form of theories *and* their mathematical meaning, and that is capable of describing modifications of the symbolic form that are not strictly deductive along with the resulting change in mathematical meaning. This means we have to consider symbolic structures that *covary*, *i.e.*, vary together, with specific geometric structures, along with operations that can be performed on the symbolic structures that simultaneously act on the geometric structures (see figure 3.1).

To do this properly in complete detail with a real theory would be enormously complex, and beyond the aims and needs of this study. Indeed, in the interests of

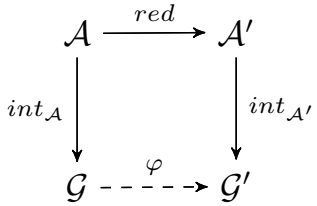

Figure 3.1: Schematic representation of an algebraic model reduction technique applied to an algebro-geometric structure (constraint). $\mathcal{A}$ and $\mathcal{A}'$ represent algebraic constructions, *e.g.*, equations, and $\mathcal{G}$ and $\mathcal{G}'$ represent their geometric interpretations. The arrows labeled *int* indicate geometric interpretation of an algebraic structure and the arrow labeled *red* indicates an algebraic model reduction method. The dashed line indicates the geometric transformation induced by the algebraic operation.

methodology avoidance, we do not want to begin with the complete detail, since all we are after at the moment are broad structural features. These features can be developed or refined in future work. But let us consider some of the basic conceptual materials that are needed in this approach. The basic materials of an $\mathcal{A}$-$\mathcal{G}$ theory are a collection of mathematical objects (numbers, vectors, functions, *etc.*), a collection of mathematical operations on objects (addition, multiplication, differentiation, integration, *etc.*), a collection of mathematical relations (equality, inequalities, *etc.*),[3] a collection of equation formation rules (for theoretical equations, constitutive equations, boundary conditions, initial conditions, *etc.*), and a collection of equation transformation rules (logical rules for equality, model reduction methods, *etc.*). A major difference in this approach from a logical approach is that the objects forming the basic materials from which models are constructed have a great deal more structure.

Consider once again Newton's second law in the form

$$\frac{d^2\mathbf{x}}{dt^2} = \mathbf{f}.$$

We saw before that before the number of degrees of freedom has been determined, $\mathbf{x}$ and $\mathbf{f}$ could potentially be interpreted in infinitely many spaces of the form $C^2(\mathbb{R} \to \mathbb{R}^n)$ and $C(\mathbb{R}^2 \to \mathbb{R}^n)$, respectively. In order to represent this variability in potential

[3]Strictly speaking, algebraic theories are equational, they do not contain inequalities or quantifiers. And in the model construction process we will not use inequalities or quantifiers. For the purposes of putting conditions on models, however, inequalities and quantifiers are sometimes needed. Accordingly, the conditions that may be placed on models are generally treated as distinct from the constraint equations that define the model, and the model itself, once constructed, can be treated as distinct from the theory proper.

interpretation, an *unspecified* object has to be represented as a symbol together with a *range* of possible spaces that could interpret it. The way this sort of thing is handled is that $\mathcal{A}$-$\mathcal{G}$ objects are assigned a space symbol and a value symbol, each of which can have a variable interpretation. The general idea, then, is that the interpretation of the space symbol varies over the infinitely many possible spaces within which the object may lie. And the interpretation of the value symbol varies over the possible values in each of these spaces. Then, through the application of constraints, such objects can be in a wide variety of different levels of specification, from neither the domain nor the value being specified, to both being specified. In order to make this definite, then, the set of possible spaces within which the object can lie must be defined. Thus set of spaces we call the *expanse* of the object.

As an example, let us consider how a real vector $\mathbf{v}$ is described on the $\mathcal{A}$-$\mathcal{G}$ perspective we are working with. The $\mathcal{A}$-$\mathcal{G}$ description of $\mathbf{v}$ includes its value-symbol $\mathbf{v}$, its space-symbol V as well as an expanse $\mathcal{V}$ of spaces within which it could be interpreted. In this case the expanse is the set

$$\mathcal{V} = \{\mathbb{C}^n \,|\, n \in \mathbb{N}, n > 0\}.$$

Thus, the unspecified vector-object $\mathbf{v}$ could be presented as the triple $(\mathbf{v}, V, \mathcal{V})$, or as the pair $(\mathbf{v}, V)$ if the expanse is understood, or simply as $\mathbf{v}$ if both the space and expanse are understood. If a value for n is specified, then V refers to $\mathbb{C}^n$, but $\mathbf{v}$ may remain unspecified. Once $\mathbf{v}$ is specified, however, it will take a particular value in $\mathbb{C}^n$. In a similar way, when the function-objects $\mathbf{x}$ and $\mathbf{f}$ in Newton's second law are unspecified they have as their expanses $\mathcal{X}$ and $\mathcal{F}$ sets of C^2 and C^0 function spaces, respectively. When the number of degrees of freedom has been specified, the expanses of $\mathbf{x}$ and $\mathbf{f}$ reduce to particular C^2 and C^0 function spaces, respectively, and the space-symbols X and F refer to these respective spaces. When the vector field $\mathbf{f}$ has been specified to be a particular function, the value-symbol $\mathbf{f}$ then refers to that function. Determining what the value-symbol $\mathbf{x}$ refers to requires solving the differential equation.

This clarifies the basic manner in which mathematical objects are treated on our $\mathcal{A}$-$\mathcal{G}$ approach. In order to capture the structure of a differential equation, such as

(1.1), in addition to simple objects, functions and operations on functions, we need to be able to specify the $\mathcal{A}$-$\mathcal{G}$ treatment of function application, function composition, arithmetic operations and equation formation. The basic idea of all such implementations of mathematical concepts in this $\mathcal{A}$-$\mathcal{G}$ framework is the same. For any given function or relation, when it makes sense to apply them to $\mathcal{A}$-$\mathcal{G}$ objects, they will be specified in such a way that when all objects are space-specified they will operate in the usual mathematical way. The result is that the $\mathcal{A}$-$\mathcal{G}$ representation of a mathematical construction, such as a differential equation, contains the information about how it was formed from its basic mathematical objects, together with the range of possible values these objects can take that are consistent with the meaning of the construction.

For example, consider the example of the construction of the model of the simple pendulum in the first chapter. There we distinguished two distinct kinds of types of equations: theoretical equations, such as Newton's second law in the form (1.1); and model equations that describe a particular system or class of similar systems, like the equation of motion of the simple pendulum (1.3). The latter has a specific geometric interpretation, such as a particular phase space with a particular vector field on it, *etc.*, and the former has a wide range of possible geometric structures of this sort as its interpretation. Thus, we may see that the later model equation has its system space specified, while the former theoretical equation does not. We will see that for model equations, in the sense of a mathematical model constructed by applied mathematicians, all of the mathematical objects making up the equation have their space specified. This is a necessary but not a sufficient condition for an equation to be a model equation, however, since, as we will see in the next subsection, model equations generally require some degree of value-specification as well.[4]

As will become clear though the study, an important part of the strategy of the $\mathcal{A}$-$\mathcal{G}$ approach we take in this study is that we begin with a consideration of the detailed structure and behaviour of some part of the application process and then abstract away details that are not relevant to the structure or behaviour we are interested in clarifying. For example, if, in the interests of methodology avoidance, we wanted to

[4]For example, specifying a particular function $\mathbf{f}$ for the vector field in the simple pendulum model (1.3) value-specifies the function $\mathbf{f}$ in (1.1).

emphasize the space specification that occurs when a model is constructed, we could reduce the structure of each of the equations to a symbol $\mathcal{E}_X$ and the structure of the expanses to a variable space $\boldsymbol{\mathcal{S}_T}$ for the theory equation and a fixed space $\mathcal{S}_{M^i}$ for the model equations. The complex syntactic transformations between equations can then reduced to *arrows* between equations, and the complex interpretation of the equations can be reduced to an interpretation arrow in the (fixed or variable) space. In this way, the analytic procedure of model construction traces out a *directed graph*, starting with some theoretical equation $\mathcal{E}_T$, where each node of the graph is an equation that has a particular geometric interpretation. If we include the geometric interpretation in the graph structure, as we did in figure 3.1, then it becomes a pair of superimposed graphs, the upper one tracing the analytic modification of the constraints, the lower one tracing the corresponding induced modifications of the geometric interpretations (see figure 3.2).

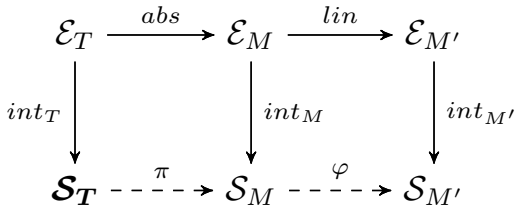

Figure 3.2: Schematic representation of the construction of the equation of motion for the simple pendulum model (1.3) ($\mathcal{E}_M$) from the basic equation of vector mechanics (1.1) ($\mathcal{E}_T$), and its linearization to produce the simplified simple pendulum model equation (1.5) ($\mathcal{E}_{M'}$). The geometric interpretations of the equations in domains are indicated by the $\mathcal{D}_X$, where bold symbols represent variable spaces. For the arrows, *abs* represents the abstraction process involved in construction of a model of a natural phenomenon, with π the induced projection of a large space of potential spaces onto a single space, and *lin* represents the linearization technique, with φ the induced morphism of the space interpreting the model equation.

This, then, provides us with an abstract, schematic algebro-geometric model of model construction—a kind of minimal model. More elaborate model construction procedures, such as those that employ a variety of model reduction techniques, would produce more complicated graphs. And if we wanted to consider more detailed structure in particular cases, or general cases at different levels of abstraction, the simple graph model would have to be "complexified" (structure expansion) in order provide equations, space expanses, transformations or interpretations with more structure. This provides a simple illustration of the advantage of an approach that is intentionally structure-reducing when it promotes clarity of understanding to do so. In

this case, the approach allows us to know exactly what sorts of structure we are eliminating in the abstraction process, providing us with a means of knowing what this minimal model is actually intended to represent in practice. This is because we also know exactly what sorts of structure we are preserving and reducing out in the structure reduction to a graph model of the form of figure 3.2. And as a result, remembering the reduction procedure allows us to *recomplexify*, *i.e.*, to add the right structure back in later on.

The classical syntactic and semantic approaches to theories are sometimes called formal and informal, respectively (*cf.*, Muller, 2011). This is because the received view treated theories in terms of formal languages and formal calculi and the semantic approaches typically work informally by relying on set theoretic concepts to formulate theories. Consequently, the algebro-geometric approach we are considering might be deemed non-formal, since it is even less formal than the semantic approach, relying on methods from applied mathematics like structure reduction. It is still, however, tied to well-characterized mathematical structures and concepts. Similarly to the semantic approach, we are relying on the fact that notions like manifold, structure preserving map or operation, graph, *etc.*, are sufficiently well-defined formally. It is for this reason that I shall refer to this approach as *quasi-formal*. The advantage of this approach is that it achieves clear and distinct meaning, a classical virtue worth preserving, with an economy of explicit formalization, thereby managing the costs of feasible description.

This manner of formalization is all we require for our purposes now; we neither need nor desire a more formal treatment, since at this developmental stage of an algebro-geometric methodology avoidance strategy, flexibility is more valuable than fully rigorous precision. But even more so than this, given that we are interested in effectively *modeling* the application of theory using this sort of strategy, it is not desirable to get bogged down in formal details. As long as we are certain that we are neither working with ill-defined notions nor are we at risk of running into contradiction, then such an approach is well-founded. We may end up working with notions that, for example, are only partially defined over a space expanse, but insofar as this results in apparent contradiction, this is regarded as reaching a boundary of applicability of the model, not as a devastating falsification.

If quasi-formal modeling strategies using an $\mathcal{A}$-$\mathcal{G}$ or $\mathcal{L}$-$\mathcal{S}$ approach can be demonstrated to be useful and effective for elucidating the methodology and epistemology of real science, then it will be necessary to provide a rigorous grounding for background $\mathcal{A}$-$\mathcal{G}$ or $\mathcal{L}$-$\mathcal{S}$ frameworks. A framework of this sort would, *inter alia*, provide a rigorous foundation for the $\mathcal{A}$-$\mathcal{G}$ or $\mathcal{L}$-$\mathcal{S}$ approach and provide a means of clarifying general features of $\mathcal{A}$-$\mathcal{G}$ or $\mathcal{L}$-$\mathcal{S}$ models. Building a proper theoretical framework for $\mathcal{A}$-$\mathcal{G}$ or $\mathcal{L}$-$\mathcal{S}$ models we construct would require tools from universal algebra, model theory and theoretical computer science. Such a theory would be a classical formulation of a feasible modeling framework. The modeling tools we are developing now are then tools that could be refined for the purposes of feasibly applying such a general theory. This is not merely a speculative concern, since I will argue in chapter 6 that the results of our study in part II provide solid evidence that $\mathcal{A}$-$\mathcal{G}$ modeling strategies will be effective for elucidating feasible epistemology. Suppes has a slogan for his semantic approach, *viz.*, "philosophy of science should use mathematics not metamathematics," emphasizing the advantages of an informal semantic approach. Given that the algebro-geometric approach relies on modeling methods adapted from methods used in applied mathematics as a *complement* to the formal methods from foundations of pure mathematics typically appealed to in formal philosophy of science, it could be described using a variation of Suppes' slogan, *viz.*, *philosophy of science should use applied mathematics, not only pure mathematics or metamathematics.*

3.1.2 Algebro-Geometric Constraints

Up to this point we have considered, in a very schematic way, how the $\mathcal{A}$-$\mathcal{G}$ perspective approaches modeling the construction of equations of motion in applied mathematics. It remains to consider how these equations are used to define mathematical *problems* that require solution. The $\mathcal{A}$-$\mathcal{G}$ framework we are considering is set up to treat mathematical problems as *constraint satisfaction problems*. This is to say, one is presented with a set of constraints on a possible solution, and the problem is to determine what solutions in the given framework are compatible with the constraints. The particular reason that it is natural to model mathematical problems in this way is that it is natural to regard equations of motion as constraints on possible solutions to those equations. For equations to be regarded as constraint problems in this way,

however, there must be a well-defined framework of possible interpretation for the equations so that the solution finding problem is *well-constrained*, in a sense we will clarify.

To see what is involved in the shift from general equations to constraint equations consider the equation of motion for the simple pendulum model, *viz.*,

$$\ddot{\theta} = -\frac{g}{\ell}\sin\theta.$$

One way of viewing this equation is that it is a constraint on the functions $\theta(t)$ that can describe the motion of a pendulum of length ℓ. And this, indeed, is how this equation is usually considered. It could just as well, however, be considered to be a constraint on the value of the gravitational constant g given a known function $\theta(t)$ for the motion of a pendulum of length ℓ. Thus, it could be regarded as a constraint that allows one to convert measurements on a pendulum to measurements of g. In order for an equation to qualify as defining a constraint *problem*, then, the quantities that are to be found as solutions to the problem are to be distinguished from the quantities that are not. We will call quantities of the former sort *solution quantities* and quantities of the latter sort *data quantities*. Thus, for the simple pendulum model as we considered it in chapter 1, the data quantities are ℓ and g and the solution quantity is the function $\theta(t)$.

When equations are written as constraints, they are often written in such a way that the left-hand side contains a complex of mathematical objects and the right-hand side takes the value 0. Thus, for the general phase space ODE, the constraint form would be

$$\frac{d\mathbf{x}}{dt} - \mathbf{h} = 0,$$

indicating that the difference between the first order derivative of $\mathbf{x}$ and $\mathbf{h}$ is zero. Strictly speaking '0' denotes the zero function, the function on the shared codomain of $\mathbf{x}$ and $\mathbf{h}$ that takes the value 0 everywhere. In general, whatever the type of mathematical object on the left-hand side, the right-hand side must be the zero object of that type. We will say that a constraint equation written in this form is in *canonical form*. The reasons for writing the constraint are associated with error in introduced in the solution finding process, as will become clear in the last section of

this chapter.

Working with a natural generalization of the case of vector mechanics, then, an algebro-geometric constraint equation establishes a relation between a number of mathematical objects, each of which is constrained to lie in a particular space (first order variable) or within one of a set of such spaces (second order variable). Suppose that the algebraic form of a constraint equation relates n mathematical symbols a_i, then the constraint can be expressed in terms of some algebraic function F of the a_i such that

$$F(a_1, \ldots, a_n) = 0. \tag{3.1}$$

In the example just considered, we would represent this as

$$F(\mathbf{x}, \mathbf{h}) = 0,$$

where $F(\mathbf{x}, \mathbf{h}) = \frac{d\mathbf{x}}{dt} - \mathbf{h}$.

Since this we are considering an algebro-geometric constraint, the constraint also has implications for the geometric interpretation of the symbols a_i. Each symbol a_i must refer to values in some space S_i or in one member of a set of spaces $\mathcal{S}_i$. Accordingly, the level of specification of a constraint depends on the level of specification of the objects related by the constraint. But even if all of the objects appearing in a constraint equation are unconditioned, mathematical consistency will require that certain space or value specifications of the objects will be impossible. Thus, a general unspecified constraint equation still has mathematical content.

These general constraints coming from theory are not sufficient to produce a descriptive mathematical model, since the wide range of possible interpretation of the constraint determined by Newton's second law is compatible with vastly many and contradictory descriptions in a given context. The construction of a mathematical model, then, requires additional constraints on the mathematical objects of a constraint equation in order to determine which space the objects lie in, and for some objects constraints on possible values within their respective spaces. Since allowing an equation to determine a constraint problem requires specifying which objects are input to the problem and which are output, we will assume that the objects appearing in an algebro-geometric constraint can be of two fundamental types:

1. *data-type objects*, which are objects given as input data to a model, *viz.*, that are required to specify a model;

2. *solution-type objects*, which are objects sought as solutions consistent with a system of constraints.

Different constraint problems can be obtained from the same equation, then, by changing which objects are of data-type and which are of solution-type.

Rarely will a single algebro-geometric constraint determine a unique value-specified solution. Usually a variety of different constraints are required. A collection of one or more algebro-geometric constraints is called a *constraint system*. Assuming that a collection of constraints is *consistent*, *i.e.*, that there is at least one simultaneous value-specification of all the objects (data and solution) that satisfies the constraint, then there are three principal classes of constraint systems.[5] The first sort corresponds to the unspecified equations of a general theory, such as Hamilton's equations.

Definition 3.1.1 (Underconstrained System) A system of constraints is *underconstrained* if one or more objects interpreting symbols in a constraint in the system it does not have its space specified.

In order for a constraint problem to be computable, it must both be specifiable, in the sense of fixing values for its data-type objects, and solvable in a finite number of steps. Thus, the definition of a "well-constrained" system is intended to pick out constraint problems that are in principle computable, at least locally and approximately.

Definition 3.1.2 (Well-Constrained System) A system of constraints is *well-constrained* if all objects interpreting symbols in a constraint have their spaces specified and the values of the data-type objects can all be specified by giving values for a finite number of parameters.[6]

[5]There are many more possibilities for how this can be handled, including allowing the ability to talk about well-constrained systems over expanses of spaces, but this level of complexity suits the present purposes.

[6]Note that some data-type objects are functions, which would in general require infinitely many parameters to specify. A system containing data-type functions that are not fully specified can still be well-constrained by this definition, however, if the range of allowable functions is finitely parameterized.

The two model equations for the simple pendulum system we have been considering are both well-constrained in this sense. There is then an important intermediate case, which picks out well-defined spaces of well-constrained systems:

Definition 3.1.3 (Partially Constrained System) A system of constraints is *partially constrained* if all objects have their space specified but the values of the data-type objects cannot all be specified by giving values for a finite number of parameters.

A note about the scope of these definitions is in order here. They have been chosen for the purpose of describing theoretical frameworks in applications of newtonian mechanics and hamiltonian mechanics that rely on numerical methods for ordinary differential equations, and initial value problems in particular, to compute solutions to modeling equations. This is because the modeling aims of the current project are quite limited given that we are just exploring the development of new methods for studying feasible theory application that adapt methods from applied mathematics. Accordingly, the framework specified by these definitions will not be adequate for capturing important structural features in many other areas of applied mathematics. Nor will it be adequate for all applications of newtonian and hamiltonian mechanics. Indeed, Wilson (*e.g.*, in 2006) has shown how many real cases of modeling are not "applications of classical mechanics" in any straightforward sense, since they adopt mixed strategies, combining tools and assumptions from newtonian mechanics, analytical mechanics and continuum mechanics in order to adequately model dominant behaviour under various conditions and at different physical scales. Developing $\mathcal{A}$-$\mathcal{G}$ modeling methods suited to studying applications of this sort will require a very different approach, which, *inter alia*, will require an effective way of treating the mixing of algebraic constraints on very different kinds of mathematical structures. The aim of the current study, however, is to form a simple model of the entire process of applying a theory, including model construction, data handling and computation. To develop an initial model of this kind it is advantageous to begin with simple cases before more complex cases are treated. This is why we are using case examples that are quite simple mathematically, because even for quite simple mathematical models, the entire application process develops its own complexity when all the details of data handling and numerical computation are taken into consideration.

Recognizing the limited scope of the $\mathcal{A}$-$\mathcal{G}$ framework we are developing is important in order to understand the way we will be using language throughout the study when reasoning *using* the $\mathcal{A}$-$\mathcal{G}$ epistemological model. Part of the reason for developing this background framework is to pin down the meaning of language or terms that would otherwise be indefinite. An example of this is the concept of an "underconstrained system of equational constraints". This is given a definite meaning in terms of symbols, spaces, expanses, *etc.*, by the above definitions. The intent of this concept is to pick out the framework defined by theoretical equations from general theories, such as Newton's second law or Hamilton's equations, which yield an underconstrained system of equational constraints in the kinds of applications we are considering. Keeping in mind the considerations of the previous paragraph, this will not necessarily be the case in all cases of modeling. But, it is expected that general equations, whether they be difference equations, differential-algebraic equations, functional differential equations, partial differential equations, stochastic differential equations *etc.*, will have some analogous kind of expanse structure in an adequate $\mathcal{A}$-$\mathcal{G}$ model due to their being too general to describe any particular physical situation. This relates to a general point made by Smith (2002*b*) that general differential equations, such as Newton's law of universal gravitation, do not describe the state or behaviour of any body. In order to do so they require the specification of *constitutive equations*, *viz.*, equations that characterize the state or behaviour of a certain class of body. In our terms here, underconstrained systems of equations formed from general theoretical equations require the specification of additional constraints, including constitutive equations, in order to make any descriptive claims about phenomena. It is for these reasons that we pin down the meaning of the term *theoretical framework* by defining internally to our $\mathcal{A}$-$\mathcal{G}$ model as an underconstrained system, with the understanding that a varied or more general concept of 'underconstrained system' would be required more broadly and that this meaning of the term 'theoretical framework' is more limited than what may be demanded for different epistemological modeling purposes.

With these provisos made clear, let us now consider how we will pin down the meaning of the terms 'theory' and 'model' for the purposes of our $\mathcal{A}$-$\mathcal{G}$ epistemological model. Model equations that describe particular system-types, like the equations for

the simple or double pendulum or for the orbital motion of a near-Earth object, will generally yield a well-constrained system of equational constraints. This is essentially because in order for a model to be capable of making distinct assertions about particular systems, the model equation must be capable of being applied to a specific case by specifying finite data, *i.e.*, a finite number of parameters. The ideal simple pendulum model does just this, since the equation of motion depends on g and ℓ. The quantity g is usually treated as a constant, but away from the surface of the earth it is variable, as it is at different latitudes because the Earth is an oblate spheroid. And ℓ is also usually treated as constant, but it can have a large range of possible values, allowing the model to describe physical pendulums with a variety of different rod lengths. When the values of all the data parameters have been specified, the model equation then describes the dynamics of a particular system, which can be studied under a variety of different conditions. Thus, when we use the term 'mathematical model' we will generally refer to the concept of a well-constrained system. A well-defined framework that can be specified in numerous ways to generate mathematical models of this sort corresponds to the case of a partially constrained system. Consequently, when we use the term 'modeling framework' we will refer to the concept of a partially constrained system. We will not use the term 'theory' in such a definite way, but generally we will take this term to refer to the concept of a theoretical or modeling framework.

One final issue that it is important to be clear on is that a well-constrained system of constraints is generally not sufficiently constrained to pick out definite behaviour. This is because initial conditions and/or boundary conditions must be applied in order to pick out definite behaviour. Such a system must have at least one particular solution, since by definition it is consistent. But there could be, and generally is, a large space of particular solutions consistent with the constraint system. This includes models such as the ideal pendulum models which have no initial or boundary conditions applied, but it also includes such models with an open set of initial data. The set of particular solutions will then contain one solution for each point in the open set. This variety of well-constrained system is called *underdetermined*, the same term used in linear algebra for a problem that has more than one particular solution. In order to render an underdetermined system *well-determined* additional constraints

must be added to ensure a unique solution. A subtlety that arises here is that some models are well-determined for some parameter values but become underdetermined at some critical parameter values where multiple solutions become possible. A simple example of this kind of case is Euler buckling (see, *e.g.*, Zeeman, 1976). Since the point in control space where a model becomes underdetermined can actually carry important physical information, as in the Euler buckling case where it identifies critical loads when buckling occurs, underdetermined models can be valuable. However, it may still be desirable to add constraints to prevent underdetermined systems, such as load constraints to prevent buckling.

Conditions can also be added, however, to turn an underdetermined system into an *overdetermined* one, *i.e.*, a system that has no particular solutions. This is often the case when empirical data is used to constrain a mathematical model. For example, two astronomical cases of this we will see in this study are the orbit determination problem and telescope error models used to convert image plates into celestial coordinates of observed objects. This, strictly speaking, renders the constraint system inconsistent. But rather than meaning that the model is worthless, it means that one must seek solutions of best fit, *i.e.*, solutions that minimize the deviation from the satisfaction of the imposed (non-data) constraints. Least squares solutions are a common variety of this type, which shows that overdetermined systems define a different type of mathematical problem, including inverse problems or parameter estimation problems, which require different techniques of solution.

3.1.3 Local Constraint Logic

The model of mathematical model construction developed in this chapter relies on the algebro-geometric framework we have been considering, in particular on the framework of algebro-geometric constraints. A framework capable of supporting the full model of theory application developed through this study, however, requires a more general sort of logico-structural framework. An example of structure that goes beyond what an algebro-geometric framework of the sort we have considered are the programming languages needed to implement numerical methods on computing devices, which we will discuss the importance of in chapter 5. In the interest of feasible modeling of theory application and effective methodology avoidance, however, there

is no need to consider the specification of such a framework for our present purposes. Instead, for many purposes we can avoid the details of the algebro-geometric framework by presenting it in terms of a "high level" system of logical concepts, meaning the concepts are highly abstracted from the detailed $\mathcal{A}$-$\mathcal{G}$ framework. We will call this system of concepts, which will prove to be very useful throughout the study, *local constraint logic*.

The basic idea of this system of logical concepts is that it provides a means of presenting the framework of algebro-geometric constraints in a perspicuous way in terms of generalized versions of concepts of formal logic. One of the motivating factors for the introduction of an algebro-geometric constraint approach was to be able to represent the equations in theories and models in applied mathematics in such a way that allows for covariation of the algebraic form of equations and their interpretation in terms of geometric (or topological) spaces. In addition to enabling representation of model reduction methods such as linearization, this approach allows for the treatment of error and variation of equations and their geometry. It might not appear that there is a natural logical means of handling this, but the semantic approach of Beth and van Fraaseen considered in the last chapter showed how to relate logical notions to a semantics in terms of state-spaces. The logical concepts we introduce here are both more flexible and more general than van Fraassen's approach, however, as we will see.

We will introduce the basic concepts through examples. Let the theoretical framework be that of newtonian mechanics, so the basic equation of the theory (autonomous case) is $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, where $\mathbf{f}$ is understood to be a force-to-mass ratio. Suppose that we are modeling a spring with one end fixed and the other end left to hang freely with a weight attached. Then suppose that within the general newtonian theoretical framework, we assume that $\mathbf{x}$ is a scalar distance x from the equilibrium position of the spring, m is the mass of the weight attached to the spring, and the spring force is modeled with the scalar Hooke's law, *viz.*, $\mathbf{F} = -kx$, where k is a constant measuring the stiffness of the spring. Under these constraints, $\mathbf{f}(\mathbf{x}) = -\frac{k}{m}x$. Consequently, the point-particle equation of motion is a specialized constraint consistent with the

general $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ constraint. Introducing new notation, we may write this as

$$\mathbf{x} = x, \mathbf{f}(\mathbf{x}) = -kx/m \Vert\!\vdash_{\ddot{\mathbf{x}}=\mathbf{f}(\mathbf{x})} \ddot{x} = -\frac{k}{m}x,$$

which indicates that $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ is the *imposed constraint system*, the constraints on the left are the *assumed constraints* and that $\ddot{x} = -\frac{k}{m}x$ is a *valid specification* of the imposed constraint system given the assumed constraints.

The notation is a logical-style way of representing the specification or solution of an algebro-constraint system. The reason for the choice of a logical-style notation will become more clear later in this subsection. For now, let us consider the rationale for the choice of notation. The ordinary logical symbol for (syntactic) entailment ($\vdash$) has a single vertical stroke, indicating linguistic assertion. The symbol for semantic entailment ($\Vdash$), used in forcing and in van Fraassen's semi-interpreted languages framework, which we saw in the previous chapter, has a double vertical stroke, indicating semantic assertion. Since we are tracing consistency and consequence in *both* algebraic constructions (language) and mathematical (geometric, topological) interpretation, we are tracing both syntactic and semantic consistency and consequence. It is for this reason that we will use a turnstile with a triple vertical stroke to denote algebro-geometric (or logico-structural) consistency. As we will see below, we can also consider consequence in local constraint logic, which is where a stronger connection to ordinary logic is made.

Because the basic notation we are introducing here is denoting consistency, it is important to clarify how it works. In the simple modeling example above we could also write

$$\Vert\!\vdash_{\ddot{\mathbf{x}}=\mathbf{f}(\mathbf{x})} \ddot{x} = -\frac{k}{m}x \tag{3.2}$$

to indicate that $\ddot{x} = -\frac{k}{m}x$ is a *valid*, *i.e.*, consistent, specialization of the constraint system $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$. This is showing how any expression that can appear on the right-hand side of the turnstile with some constraints assumed can also appear there without those constraints assumed. This is the case because if a specification of the imposed constraint system is possible with certain assumed constraints, it continues to be possible with those constraints removed. Another thing to observe is that if the imposed constraint system and the assumed constraints are incompatible or inconsistent, then

there are no valid specifications consistent with the combined system of imposed and assumed constraints. This is the opposite of the situation in classical logic, where a contradiction leads to explosion—any sentence is valid. In local constraint logic a contradiction leads to "implosion", in the sense that no specification (solution) is valid. This shows that unlike ordinary logic, local constraint logic is not *monotonic*, *viz.*, in ordinary logic a consequence of a hypothesis continues to follow if additional assumptions are added, but this is not true in local constraint logic.[7] We will see throughout this study that although the classical logical concept of validity is appropriate in the context of pure mathematics, it is the concept of validity from local constraint logic that is appropriate for *applied mathematics*.

Despite these very significant differences, there is nevertheless a strong, but somewhat subtle, analogy between the ordinary concept of logical validity and the local constraint logic concept of validity. Indeed, the notation and the term 'valid' is, in part, intended to reflect an analogy between the expression (3.2) and the concept of a logically valid proposition p in a logical language $\mathscr{L}$ with a logical calculus $\mathscr{K}$, *i.e.*, where

$$\vdash p.$$

Although this particular "analogy" is entirely superficial, we will see below that overall analogy between classical logic and local constraint logic is anything but superficial.

Now, let us change to the consideration of the model of an ideal simple pendulum that we considered in chapter 1. Once again, if we suppose that we are working in a framework of point-particle mechanics, then we could write[8]

$$\Big\|\Big|_{\ddot{\mathbf{x}}=\mathbf{f}(\mathbf{x})} \ddot{\theta} = -\frac{g}{\ell}\sin\theta$$

to indicate that the ideal simple pendulum model is a valid, *i.e.*, consistent, specification of the point-particle mechanics framework. And we can continue this specification further by shifting to work in the framework where the ideal simple pendulum

[7] I thank Stahis Psillos for pointing this out this technical point.

[8] Strictly speaking the equation of motion for θ uses generalized coordinates, making it an analytical mechanics model. Nevertheless, it can be thought of in terms of a change of variable from cartesian coordinates in the plane to polar coordinates.

constraint is imposed and write

$$\theta(0) = \theta_0, \dot{\theta}(0) = 0 \,\|\!|_{\ddot{\theta}=-\omega^2 \sin\theta} \sin\tfrac{\theta}{2}(t) = A\,\mathrm{sn}(\omega t, A),$$

to indicate a valid, *i.e.*, consistent, solution to the ideal simple pendulum under the conditions of initial angle θ_0 and zero initial angular velocity. The solution is given in terms of the Jacobi elliptic function $\mathrm{sn}(u,k)$, where $A = \sin\frac{\theta_0}{2}$ is the amplitude of $\sin\theta/2$ (Lawden, 1989). In this case it is also the unique solution to the model.

This shows us how the ordinary single turnstile indicates specifications that exactly satisfy, or are exactly consistent with, the imposed constraints. But to model theory application we also require the ability to indicate specifications that only approximately satisfy the imposed constraints. Before we consider how this is handled, we require a new concept. In any situation where error needs to be analyzed for a given mathematical quantity, it is necessary to introduce an appropriate measure of the *size* of the error, which requires a measure of size of the given type of quantity. In mathematics, the size of a quantity is specified by a *norm*, denoted $\|\cdot\|$, which is a function that assigns a real- or complex-valued size to the members of a space of mathematical objects; that it is meaningfully a measure of size is ensured by the satisfaction of certain axioms (see, *e.g.*, Trefethen & Bau III, 1997). In the case of simple quantities, *e.g.*, real or complex numbers, the absolute value/modulus provides a canonical norm. For structured quantities, *e.g.*, vectors, matrices, functions and operators, however, there is no such canonical norm. As a result, which norm is appropriate depends on the context. We need not consider this contextuality for our purposes here, and so we will simply assume that a choice of norm appropriate for the context has been made.

Returning to the consideration of approximate specifications, consider first a simple case of a linear system

$$Ax = b,$$

where A is a matrix of data, b a vector of data and x is the solution vector. Now, suppose that we do not require an exact solution, *i.e.*, a solution a that satisfies

$$Aa - b = 0,$$

but only a solution with a small backward error, *viz.*, a solution $\tilde{a}$ such that for some *error tolerance* $\varepsilon > 0$,

$$\|A\tilde{a} - b\| \lesssim \varepsilon.$$

For an exact solution $x = a$, we would write

$$\Big\|\vdash_{Ax=b} x = a,$$

but for a solution $\tilde{a}$ with an ε-small backward error, we could write

$$\Big\|\vdash^{\|\delta_x\| \lesssim \varepsilon}_{Ax=b} x = \tilde{a},$$

where $\delta_x = A\tilde{a} - b$ is the *defect*, the amount by which $\tilde{a}$ fails to satisfy the constraint.[9] This indicates that allowing a tolerance for error in the solution to a constraint problem means that there will be a *range* of acceptable solutions to the problem in the neighbourhood of the exact solution $x = a$. Any acceptable solution other than $x = a$ is not valid *simpliciter* in this framework because $Ax = b$ will not be satisfied exactly. For such approximate solutions, we may say that $x = \tilde{a}$ is an *effectively valid* solution to the constraint problem $Ax = b$.

We could also consider the same kind of case where we require an ε-small *forward* error, in which case we could write

$$\Big\|\vdash^{\|\epsilon_x\| \lesssim \varepsilon}_{Ax=b} x = \tilde{a},$$

where $\epsilon_x = a - \tilde{a}$ is the *error*, the amount by which the effective solution $\tilde{a}$ differs from the exact solution a. In case we know what ε refers to, we may simply write

$$\Big\|\vdash^{\varepsilon}_{Ax=b} x = \tilde{a}$$

[9] Note that I have been assuming that the backward error and the defect are equal to one another, which is indeed the case in all the cases of interest in the study. It should be noted, however, that *backward error* and *defect* are distinct concepts and they are not in general equal. To see why they are equal for ODE, let the function $P(f(y)) = y$, where $\dot{y} = f(y)$, define the ODE problem. Now let $\tilde{y}$ be a computed solution to the problem. The backward error is the amount by which we need to perturb the data of the problem, which is simply $f(y)$ in this case, so that the computed solution is the exact solution of the problem P with the perturbed data. Since the computed solution satisfies $\dot{\tilde{y}} = f(\tilde{y}) + \delta$ by the definition of the defect δ, it follows from the definition of P that $P(f(\tilde{y}) + \delta) = \tilde{y} + \delta$. Thus δ is also the amount that the data for the problem P must be perturbed for the computed solution to be the exact solution to P for that perturbed data, *i.e.*, δ is also the backward error.

to indicate an effectively valid solution.

In addition to representing approximate specifications or solutions, we also need to be able to represent solutions that are only accurate over a limited range of conditions. To see how this is done, consider once again the framework of the ideal simple pendulum constraint under the condtion that the pendulum starts at rest at an initial angle θ_0. Suppose also that we are content with an approximate solution that is effectively valid at each time t in an interval $[0, T]$.[10] In this case, we do not require the exact solution to the ideal model but only a solution that is *temporarily close* to the exact one, *i.e.*, a solution $\tilde{\theta}(t)$ that does not differ from the exact solution $\theta(t)$ by more than some $\varepsilon > 0$ over some time interval $[0, T]$. Let the difference between the solutions be the *error at* t

$$\epsilon_\theta(t) = \theta(t) - \tilde{\theta}(t).$$

Then we can express the closeness condition by

$$\|\epsilon_\theta(t)\| \lesssim \varepsilon, \quad t \in [0, T].$$

Now, if we only need to study small oscillations of the simple pendulum, *i.e.*, $|\theta| \ll 1$ rad $\approx 57°$, then we can obtain a good approximation by linearizing the nonlinear ideal simple pendulum model to obtain $\ddot{\theta} = -\frac{g}{\ell}\theta$. In this case, the solution is $\tilde{\theta}(t) = \theta_0 \cos \omega t$. This is now a solution that is not valid *simpliciter* in the framework of the nonlinear constraint equation in two distinct senses. First of all, at any given time (other than $t = 0$) such that the condition $\|\epsilon_\theta(t)\| \lesssim \varepsilon$ holds, the approximate solution $\tilde{\theta}(t)$ is only effectively valid. Second of all, the solution is only effectively valid for small angles and short times. So, it is only valid *locally* to that range of conditions. Thus, we can write

$$\theta(0) = \theta_0, \dot{\theta}(0) = 0 \Big\|\!\!\Big|^{\|\epsilon_\theta(t)\| \lesssim \varepsilon, t \in [0,T], |\theta_0| \ll 1}_{\ddot{\theta} = -\omega^2 \sin\theta} \theta(t) = \theta_0 \cos \omega t$$

to indicate that, subject to a tolerance ε for the error in the solution, the linearized

[10] This could be because we are modeling a physical pendulum in terms of the ideal pendulum, which has no friction, so we know that at best the model solution will only be accurate over a short period of time.

solution $\tilde{\theta}(t)$ is *locally effectively valid*, *viz.*, effectively valid in the sense that $\|\epsilon_\theta(t)\| \lesssim \varepsilon$ locally to the restricted range of angles and times where $t \in [0, T]$ and $|\theta_0| \ll 1$.

Note that the linearized model is still a valid point-particle model, so the condition

$$\Big\|\!\!\Big|\overline{\ddot{\mathbf{x}}=\mathbf{f}(\mathbf{x})}\ \ddot{\theta} = -\frac{g}{\ell}\theta,$$

holds, as does the condition

$$\theta(0) = \theta_0, \dot{\theta}(0) = 0 \Big\|\!\!\Big|\overline{\ddot{\theta}=-\omega^2\theta}\ \theta(t) = \theta_0 \cos\omega t,$$

since $\theta(t) = \theta_0 \cos\omega t$ is the exact solution to the linear simple pendulum model for the stated initial conditions. Thus, we can say that the solution $\tilde{\theta}(t) = \theta_0 \cos\omega t$ is only locally (effectively) valid in the framework of the nonlinear simple pendulum but is *globally valid* in the framework of the linear simple pendulum. We could represent this explicitly by writing

$$\theta(0) = \theta_0, \dot{\theta}(0) = 0 \Big\|\!\!\Big|\overline{\ddot{\theta}=-\omega^2\theta}^{\,t\in(-\infty,\infty)}\ \theta(t) = \theta_0 \cos\omega t.$$

Thus, whether a solution is locally or globally valid depends on the constraint framework.

For the dual purpose of conceptual clarification and to demonstrate the generality of this framework of local logic, let us consider how this framework applies in the case of a logical language $\mathscr{L}$ and a logical calculus $\mathscr{K}$, where the combined system is denoted by $\mathscr{L}(\mathscr{K})$. First of all, consider the vocabulary V of $\mathscr{L}$. A valid "state" of the vocabulary is simply a string of symbols from V. So, if p is a string of symbols, then

$$\Big\|\!\!\Big|\overline{\ }_V\ p \quad \Leftrightarrow \quad \text{“}p \text{ is a string of symbols of } V\text{”}.$$

But given that the language $\mathscr{L}$ imposes constraints on what count as acceptable strings, only if p is a well-formed formula does

$$\Big\|\!\!\Big|\overline{\ }_{\mathscr{L}}\ p$$

hold. So the valid "states" of $\mathscr{L}$ are its well-formed formulas.

Now, something different happens when we add the constraint that the calculus $\mathscr{K}$ is applied to the language, which means that the valid "states" are formulae that can be proved from no assumptions using inference rules from $\mathscr{K}$. Or in other terms, the valid "states" are those formulae for which there is a "path" from the empty set of propositions to the formula by applying inference rules from $\mathcal{K}$. Thus, by adding a set of inference rules we shift from considering consistency to considering *consequence*. To distinguish the consistency and consequence forms of local constraint logic, we will refer to the consistency form as the *state mode of presentation* of local constraint logic and to the consequence form as the *path mode of presentation* of local constraint logic.

The upshot of considering the path mode of presentation is that a proposition p is logically valid in $\mathscr{L}(\mathscr{K})$ iff

$$\Vdash_{\mathscr{L}(\mathscr{K})} p, \tag{3.3}$$

which shows that there is a much stronger analogy between the logical concept of validity and validity in local constraint logic when we shift from considering consistency to considering consequence. Now, it is important to recognize that this is *not* identical to the concept of validity from ordinary logic, because the logico-structural notion of validity in local constraint logic is not monotonic, as pointed out above. But, what is interesting to recognize here is that in tracing the language and the semantics, as we do in local constraint logic, the effect of a contradiction does lead to explosion *linguistically*, because anything follows using the logical calculus, but it leads to "implosion" *semantically*, because there are no $\mathfrak{models}$ consistent with the assumptions. This reveals that the "implosion" effect of local constraint logic arising from the non-monotonicity of the logic has a stronger connection to logic than it initially appeared.

The connection to logic can become even stronger if we shift from a treatment in terms of logico-structural constraints to a purely linguistic treatment, but where consequence is still regarded as generically local and effective, then the concept of a logically valid proposition of a logical theory is actually a *special case* of the concept of a *locally effectively valid formula* (or symbolic construction). This concept of a locally effectively valid formula is actually a generalization of logical validity because the equivalent of the concept of validity from logic is the case of local effective validity

where the tolerance ε goes to zero, so that inferences are exact ($\varepsilon \to 0$), and where there are no special conditions under which the formula is valid, so that there is no boundary beyond which an inference ceases to be valid ($\partial \to \infty$). Thus, the concept of validity from mathematical logic is recovered from a "local language logic "in a kind of limit. The consideration of local effective validity only for languages will not be pursued in this study, it is just being used for the purposes of clarifying the connection to mathematical logic. Instead we will use local constraint logic, which tracks both language (symbolic constructions) and semantics (mathematical interpretation). In this chapter and the next the focus will primarily be on the *state* mode of presentation of local constraint logic, but in chapter 5 we will shift to more of a consideration of the path mode of presentation.

We will consider more of the details for local constraint logic in later chapters. For the purposes of this chapter, however, the concept of *local effective validity* is all we will need. And except in rare cases, we will not require the formalism presented here. It is important to make the concepts clear, but it does not add perspicacity for our present purposes to continue to use the formal presentation throughout the study.

3.2 Structural Features of Feasible Model Construction

The modeling process usually begins with a particular *question* for which one seeks an answer or a *problem* that one seeks to solve. Typically in mathematical modeling, we begin with a question, which is posed as a mathematical problem to solve, in order to answer the question. For the case examples we are considering this involves the problem of studying certain aspects of the chaotic behaviour of double pendulum systems and the problem of predicting the probability that a near-Earth object (NEO) will impact Earth. From a feasibility point of view, the strategy is to refine the question or problem so that it admits of a systematic approach to find the answer or solution. This also requires a process of abstraction so that one focuses in on the dominant features of the phenomenon under consideration. Often this involves the choice of a theory or theories that are thought to characterize, or effectively characterize,

the phenomenon in question.[11] Along with some information input concerning the specifics of the phenomenon in question, often in the form of modeling assumptions, the theories, often implicitly, are used to identify a base, or most detailed, model of the phenomenon, which may or may not be constructed explicitly.

There are two basic classes of base models. One is a *classical base model*, which picks out the classically existing "most detailed" valid model in the theoretical framework. We will see that such a model is not necessarily unique. In the case of classical mechanics, there are microscopic, macroscopic and mesoscopic base models. For instance, in the case of hamiltonian mechanics, a microscopic base model would be an atomic-constitutive model of a system that tracks the dynamical evolution of each atom. And a macroscopic base model could treat the components of the systems as macroscopic rigid bodies. Mesoscopic base models would be concerned with dominant behaviour at various intermediate scales between microscopic and macroscopic models. Thus, it is seen that what the base model will be for a given theoretical framework can depend on the *characteristic scale* of the model.

Classical base models exist in the sense of a existence proof in classical mathematics, but we generally have little or no descriptive access to the details of the model. Accordingly, practical modeling requires the explicit construction of a *feasible base model*, which is not unique. It is also less descriptively accurate or less fundamental than the corresponding classical base model, and is picked out by a given practical modeling strategy. An example of a classical base model for the simple pendulum system would be one treated in terms of elastic bodies where all the bodies being modeled are massive and extended in three dimensions, with densities exactly matching the corresponding physical bodies. We see that we do not need to construct this model explicitly in order to feasibly obtain the ideal model (1.3). The ideal model can be obtained directly by a judicious choice of mathematical constraints.

Typically this base model is too complex to get useful information out of, and therefore systematic methods of abstraction are used to reduce the complexity of the model to one better suited to analysis and to focus in on the features of the phenomenon that dominate the behaviour we are interested in for the purposes of

[11] An example of the distinction between effective characterization and characterization would be the treatment of a macroscopic gas as effectively characterized by thermodynamics, even though the gas is actually understood to be characterized by statistical mechanics.

answering the question, or solving the problem, posed. One approach to this, used in both the double pendulum and NEO cases, is to make judicious use of modeling assumptions that lead to a relatively simple analytic form for the model, sometimes obviating a complex base model entirely.[12] Another approach, also used in both cases, involves mathematical techniques to modify the analytic form so as to make the model more amenable to analysis.[13] These model reduction, or variable reduction, methods ultimately result in a particular model, or models, that we analyze in order to understand the behaviour we are interested in understanding. Often this analysis requires finding solutions to the model(s) in particular circumstances.[14] In some cases analytic methods are available in order to find solutions, but increasingly, and in both the cases considered in this study, the solution process requires the construction of numerical methods, supported by machine arithmetic.

The process I have described here is at odds with the common strategy of using tools from mathematical logic in order to describe the analytic reasoning used in theory application or the theory-world relation. In that approach, the application process is typically treated as a *strict logical derivation* from the principles of a theory, or its fundamental equations, which could amount to the same thing. This is typically taken to involve the strict logical derivation of the differential equation of motion, and the further strict logical derivation of the behaviour of the system once the appropriate boundary and/or initial conditions have been supplied. In contrast, the examples just considered show that the construction of the model involves constraints abstracted from the phenomenon and the model reduction methods introduce approximations. Thus model construction does not proceed from theory alone and the methods used to construct models are not strictly deductive. Indeed, they are locally effectively valid inferences, as was pointed out in the last section.

It is also significant that the approximation methods used are not only pragmat-

[12] In the double pendulum case this involves the choice of locally effectively valid, stable macroscopic constraints for the components of the double pendulum, enabling rigid body (hamiltonian) dynamics to be applied to the system of constraints directly. In the NEO case, this involves assuming that the gravitating bodies are effectively spherically symmetric and treating them as point masses.

[13] In the double pendulum case, non-dimensionalization is used to provide a scale invariant model, effectively generalizing the analysis. In the NEO case, asymptotic analysis is used, explicitly or implicitly, in order to justify the particular (post-newtonian) dynamics used in orbit tracking.

[14] This is not always required because certain qualitative and stability properties of solutions can be studied directly from the equation of motion without having explicit solutions.

ically necessary, *i.e.*, they are not used solely for the purposes of obtaining tractable models. An important aim of these methods is indeed to yield descriptions of behaviour that are close enough to the dictates of theory given the precision required for the question or problem posed. But they achieve more than simply useful approximations, since a great amount of *insight* is gained when one is able to isolate those features of the problem that are responsible for the dominant behaviour one seeks to elucidate. Thus, the result of the assumption that model construction is a process of logical derivation is not only a significant distortion of scientific practice, but the assumption is also blind to the important epistemological function of model reduction methods—to give insight into dominant behaviour. And this does not apply only to model reduction methods, since, to quote the famous phrase from Hamming (1962): "the purpose of computing is insight, not numbers." We will see the forms of insight that computing provides in chapter 5.

We will now consider some structural features of the model construction and analysis process in more detail, and show how these features are driven by the need for the handling of complexity.

3.2.1 Handling Complex Structure

In the vast majority of cases where one needs to model a physical system using mathematics the structure of the system as it would be rendered within our most advanced or fundamental theories is vastly more complicated than we can feasibly handle. Any physical system with macroscopically many particles, for example, falls in this class of cases. For this reason alone we require techniques that allow us to simplify the model description to a such a degree that a mathematical analysis becomes feasible. Even if we could handle the full detail feasibly, however, we would still require methods to reduce complexity so that we could focus in on the dominant behaviour pertaining to the question or problem at hand, and to gain insight into why this behaviour is dominant. In this subsection we will consider ways in which the methods of abstraction from the complexity of a phenomenon are epistemologically significant. This will provide evidence that typical logical accounts of the theory-world relation miss epistemologically important structure from the point of view of how we actually relate theories to phenomena and in terms of the epistemological

significance of the methods used.

There are two main types of procedures that are essential for abstraction from complex structure:

1. The selection of modeling assumptions motivated by a consideration of the phenomenon that avoid the need to deal with the complex descriptions of our most advanced or fundamental theories; and
2. The application of analytic model reduction techniques in order to simplify the form of the under- or partially constrained equations of a theory or to reduce the model descriptions of a well-constrained system of equations.[15]

Both kinds of abstraction are epistemologically significant, as we will see, and reveal the manner in which the picture of theory application as a "derivation from theory" is woefully inadequate in terms of giving insight into the epistemology of theory application.

Consider how a model of a physical double pendulum is constructed. Given the simplicity and high symmetry of a double pendulum system (see figure 3.3) there is quite a bit that can be done using modeling assumptions in order to generate a simple yet epistemically useful model. There are a number of macroscopically stable assumptions that can be made that can effect a rather miraculous reduction in structure over how we would be compelled to describe the system in terms of its microscopic constitution. As a first approximation it is valid to treat the pivot point as rigid, the two joints as frictionless, the rods as rigid, the rods as one-dimensional, the rods as massless, the rods of equal length ℓ, and the weights as sizeless. How good these are as approximations depends on the details of the constitution of the system (respectively, how stable the pivot is, how small the joint-friction is, how rigid the rods actually are, how thin/homogeneous the rods are, how large the ratio between the weights and the rods is, and how small the weights are). So, depending on these details, one or more of these assumptions may be better than just a first approximation. From a complexity reduction point of view the use of symmetry assumptions in this way is extremely powerful. Indeed, over an atomic-constitutive model, this method reduces the degrees of freedom from over 10^{25} to 4!

[15] There is a third possibility that reduces to the latter sort, *viz.*, the reduction of the under-constrained equations of on theory to underconstrained equations of another, which is discussed in §3.3.3.

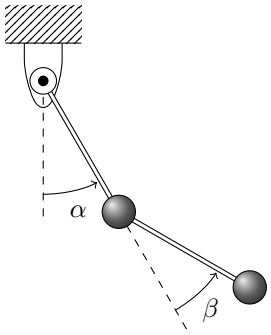

Figure 3.3: An ideal four-dimensional double pendulum system.

An interesting epistemological feature of this approach is that it effects a quite extraordinary reduction in the structure (analytically and geometrically) of the model *before one has decided what theories to use* in the modeling. The decision to treat the pendulum in terms of rigid bodies, frictionless joints, *etc.*, can be made independently of any particular system of mechanics used to model the pendulum. Thus, we may think of this complexity reduction procedure as the translation of phenomenological constraints into constraints within some theoretical framework. The result, then, is that using macroscopic constraints imported from experience it is possible to obtain a useful model without any explicit consideration of a classical base model. Quite far from being a "derivation from theory", then, the model is essentially formulated by direct abstraction from experience into a given theoretical framework. Furthermore, the role of experience in model construction seen here is not accounted for in typical logical approaches, including Suppes's "hierarchy of models" picture of the theory-world relation. Thus, even at the earliest phase of model construction there is already epistemologically significant structure that logical approaches miss, *viz.*, pre-theoretical constraints.

The construction of a model of the orbit of a NEO also uses symmetry and rigidity assumptions. The major bodies in the solar system are very close to being spherically symmetric and their rotation is very close to that of a rigid body. They are so close that the deviations are comparable in size with only the second order of relativistic corrections to the newtonian equations of motion (Kopeikin *et al.*, 2011, §6.3). Accordingly, in most solar system modeling, the bodies are assumed to be spherically symmetric and rigid bodies. A significant feature of this is that rather than having

to deal with the complexity of treating the bodies in terms of the consitutent particles, or with the complexity of partial differential equations (PDE) to treat non-rigid motion, we can use hamiltonian mechanics or post-newtonian vector mechanics, or some combination of the two. This shows that the abstraction of constraints from the phenomenon can allow the use of a theoretical framework that yields simpler model descriptions than a theoretical framework capable of a more detailed and accurate treatment of the behaviour. Provided that the constraints are *effectively* true of the phenomenon, where by 'effectively true' I mean that they provide an accurate description of the behaviour of the phenomenon within a reasonable error tolerance given the scale of the modeling, the use of the simpler theoretical framework can still result in a valid model.

Unlike the double pendulum case, taking the simplifications that follow from symmetry and rigidity assumptions into the context of the constraints of an appropriate theory is not straightforward. Our most advanced theory for the treatment of gravitational phenomena is general relativity. The equations of general relativity, however, are very complex and the mathematics is difficult to handle directly. This, together with the fact that newtonian models are much easier to handle both conceptually[16] and computationally,[17] means that it is more effective to use a *post-newtonian model*, *i.e.*, a newtonian model with relativistic corrections, when relativistic effects need to be taken into account. This is not a straightforward process, because there are a number of different post-newtonian equations of motion that one could use depending on what details need to be taken into account in the modeling. In any case, analytic simplification techniques are required in order to find the post-newtonian corrections (first-order relativistic corrections) to the newtonian equations of motion. Consequently, the symmetry and rigidity assumptions made will determine which post-newtonian limiting reduction of general relativity one will need to con-

[16]Incidentally, an important part of what I mean by conceptual accessibility in the context of mathematics is available reasoning based on the formalism, on the mathematical interpretation, and the ability to move back and forth between the two. If one can only think in one of those two main modes, then one will be limited in the level of understanding one can obtain. The more opaque the mathematics, therefore, the less conceptual accessibility. So, there is some aspect of this that is relative to an individual's capability, but there is also an "objective" side of it that is relative to human capacities for reasoning in general.

[17]There is also the fact that there is considerable technological and human resource investment in vector mechanics methods, both from a conceptual and a computational point of view.

sider. Generally one wishes to make the strongest modeling assumptions possible in the context, because this simplifies the post-newtonian corrections that need to be included in the model. Furthermore, feasibility considerations also require that one use the simplest set of equations possible, both for greatest insight and for fastest computations.

We see here that constraints of feasibility drive the need to work with the strongest simplifying assumptions possible while still allowing for an accurate model. The reasons for this are somewhat subtle, since this not only achieves a combination of ease of analytic manipulation and computation, but it also promotes insight by eliminating all information from the model that does not contribute significantly to the behaviour of interest. The basic idea here being that if it is possible to get an accurate description using a highly abstract model, then it shows that only those features preserved in the abstraction are responsible for the dominant behaviour. By the same token, if very strong simplifying assumptions yield a model that nearly captures the behaviour, this can point to a small effect that needs to be taken into account. For example, one could see that the computed orbit of an object has a small systematic error that is known to not be due to perturbing bodies. This could indicate a significant Yarkovsky effect, *viz.*, where anisotropic thermal emission of a rotating object produces a non-negligible force, meaning that rotational motion must be taken into account.

Let us now turn our attention to the analytic part of the model construction process. In the case of modeling the double pendulum system using hamiltonian mechanics, the symmetry and rigidity assumptions allow a quite simple hamiltonian to be written down for the system, for which solutions may be sought in order to study the behaviour of the system under different conditions. The interest in the double pendulum system comes largely from its providing a very simple example of a chaotic system. As a result, hamiltonian mechanics is a useful theory to use for the purposes of studying the chaotic behaviour of the pendulum. As we will see later in the chapter, modeling chaotic behaviour is particularly problematic for the logical derivation view of the theory-world relation.

The ability to feasibly construct mathematical models in the context of a physical theory is obviously quite crucial for our ability to use theories to gain insight into

the nature of phenomena. And the need to be able to construct models efficiently is reflected in many ways in the syntactic tools that physical theories provide. Indeed, most physical theories, and hamiltonian mechanics is no exception, provide us with an *algorithm* for the construction of a model on the basis of given modeling assumptions. As we discussed in the first chapter, the application algorithm of hamiltonian mechanics involves the following steps:

1. Specify a set of generalized coordinates $\mathbf{q}$ needed to track the motion of the system;

2. Calculate the kinetic $T(\mathbf{q}, \dot{\mathbf{q}})$ energy and potential $V(\mathbf{q}, \dot{\mathbf{q}})$ energy of the system in terms of the generalized velocities $\dot{\mathbf{q}}$, the time derivatives of the generalized coordinates;

3. Calculate the lagrangian $L(\mathbf{q}, \dot{\mathbf{q}})$, the difference between the kinetic and potential energy;

4. Calculate the generalized momenta $\mathbf{p}$ by taking partial derivatives of the lagrangian;

5. Calculate the hamiltonian $H(\mathbf{q}, \mathbf{p})$ from the lagrangian $L(\mathbf{q}, \mathbf{p})$ expressed in terms of the generalized coordinates and momenta;

6. Calculate the equation of motion by substituting the hamiltonian in the hamilton equations:
$$\dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} \equiv \mathbf{f}(\mathbf{q}, \mathbf{p}) \qquad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}} \equiv \mathbf{g}(\mathbf{q}, \mathbf{p}).$$

In terms of the specification of a model from the general theory, the most significant steps are the first, second, and sixth. The first step of the algorithm determines the dimension of the configuration space, which thus determines the dimension of the phase space. Thus, even though the precise dynamics has not been specified, the Euler-Lagrange equations or Hamilton's equations equations become partially specified. Constructing a model from this then requires specifying the form of the dynamics. The specification of the dynamics occurs as a result of specifying the kinetic and potential energy in terms of a finite number of parameters, though the exact

form of the dynamics on phase space is not known until the last step of the algorithm. The result of the process is a hamiltonian model. This shows, clearly, then, that the constraints input to the mathematical framework from the phenomenological considerations, *e.g.*, symmetry and rigidity constraints, come in the first two steps of the algorithm: specifying the generalized coordinates, and specifying the kinetic and potential energy of the system. The rest is an algorithmic process internal to the theory, resulting in the equations of motion.

The model that results from this process is the feasible base model referred to in the introduction to the section. Given our interest in the studying of the chaotic behaviour of double pendulum systems in general, rather than in the study of a particular system, the actual value of the length ℓ of the rods is of little consequence. As a result, it is natural to also *nondimensionalize* the hamiltonian in order to study the scale invariant behaviour of the system. This reduces the number of data-type parameters in the system by one, since ℓ is removed from the hamiltonian. This is not the most powerful application of nondimensionalization. The real power of nondimensionalization is that it can reduce the number of *solution*-type quantities in a problem, leading to a genuine simplification of the problem. The only particularly significant simplification it does offer in this case is that when studying the behaviour of solutions using numerical methods, there is no need to arbitrarily select a value of the length ℓ. This shows why the nondimensionalization step is particularly appropriate for a computer study.

Although the specification of the number of degrees of freedom and the kinetic and potential energies both require input from the consideration of the phenomenon, the specification of the kinetic and potential energies also requires mathematical input from outside of the basic theoretical constraint imposed by Hamilton's equations, since other equations from physics must be employed to compute these quantities. This requires an expansion of the system of equational constraints, introducing a number of other quantities, like g and ℓ in the case of the pendulum models. The specification of the kinetic and potential energies specify the dynamical behaviour of the system, and so are analogous to the specification of constitutive equations or equations of state in other areas of physics.

Equations specifying behaviour require additional assumptions, either modeling

assumptions that are experimentally-motivated, *i.e.*, models drawn from experimental data, or theoretically-motivated modeling assumptions, *i.e.*, models constructed on the basis of some more-or-less principled theoretical considerations of constitution or behavioural phenomenology. In either case, the result is a set of equations added to the constraint system. In the experimental case, behavioural phenomenology must be translated into equations interpretable in the theoretical framework. And in the theoretical case, constitutive equations must be chosen or constructed from the theoretical consideration of special cases of modeling contexts. An example of this is the general treatment of small oscillations in analytical mechanics. In the present case, however, it is just basic kinematics that is used to compute the lagrangian. But the more significant point is that in more complex modeling situations the computation of the lagrangian or hamiltonian may require modeling considerations that go outside of the bounds of theory, which is the case when behavioural equations require actual measurements of behaviour for their specification. This shows yet another way in which a representation of model construction as logical deduction is inadequate.

There is yet another case where input is required from outside the framework of a given theory. This is when another theory entirely is required for an adequate treatment of the behaviour of the phenomenon. If we needed to treat electrodynamic behaviour in addition to mechanical behaviour, we would have to add equations from electrodynamics to the constraint system of whichever classical theory we are using. Where such combination is common new theories are created from the merging of the basic equations of different theories. An example of this in astrophysics is magnetohydrodynamics, which combines fluid mechanics and electrodynamics for the treatment of electrically charged fluids. There are also less sharp combinations of theories, such as climate models and semi-classical theories where part of the system is treated quantum mechanically and classical assumptions are used to treat the effective behaviour of some other part. We are not, however, considering such cases in the present study.

The reason that empirical equations specifying dynamical behaviour are often required in feasible modeling is that our knowledge of how to specify both constitutive behaviour (microscopic theory) and phenomenological behaviour (macroscopic theory) of real systems is quite limited. In the case of our knowledge of constitu-

tive behaviour, in so far as the constraints on behaviour at this level are complete, they are usually of limited use for systems that have large numbers of microcomponents. Particularly when the numbers of microcomponents reaches macroscopic levels, the models of the system specified constitutively become both analytically and computationally intractable. This often forces the use of infinite idealizations in order to leverage the descriptive power of constitutive theory. This is why consideration of intertheoretic relations in scientific explanation, like those studied by Batterman (2002*b*), are feasibly necessary, whether or not they are necessary in principle, as Batterman argues. In the case of our knowledge of phenomenological behaviour, the constraints on behaviour that we can obtain usually only penetrate so far into the nature of the phenomenon. Even the Navier-Stokes equations of fluid dynamics are not complete and interesting generalizations have been developed and studied, such as the equations of Ladyzhenskaya (1967). Part of this is due to the fact that there is an underlying discrete phenomenon. But it is also the case that more fundamental phenomenological equations can be determined locally, and it is not clear that there is any unique phenomenological foundation that could be found, even in quite restricted contexts.[18] So, the limitations on feasible knowledge place strong constraints on what sorts of modeling strategies are possible and often force the need for adoption of constraints on behaviour that are not provided by theory.

Let us now consider how the orbit model is specified for NEO tracking, where we will find a fact that poses a distinct problem for the semantic view of theories. Since post-newtonian effects are not completely negligible for orbit calcuations in the solar system, we must work with some post-newtonan equation of motion for determining the orbit of a NEO. These equations of motion, therefore, will involve some post-newtonian perturbation of the newtonian equations of motion. There are also other perturbations that must be taken into account. These are physical, rather than theoretical, perturbations from other solar system bodies and other non-gravitational forces, such as forces on comets from off-gassing of the cometary nucleus. This brings us to another sort of complexity reduction considerations: the number of bodies that are taken into account in a model.

[18]This ties in to the patch-like facade structure that Wilson (2006) discusses. Also to the blurring of the worlds of thought and sense of the sort characteristic of the Hermetic system and worldview.

There are an enormous number of solar system bodies that we could consider. How many and which ones we consider will be determined by three constraints: which ones are most massive; which are nearest the potential path of the NEO; and how powerful are the computational resources we have at hand. For this problem the last concern is reduced considerably by using ephemerides provided by NASA's Jet Propulsion Laboratory (JPL) at Caltech. This allows very accurate accounting for perturbations from other solar system objects without having to compute their orbits to track the NEO. Thus, the equation of motion for the NEO is a post-newtonian N-body equation of motion of the form

$$\frac{d^2\mathbf{r}}{dt^2} = \mathbf{f}_{so} + \sum_{k=1}^{n_p} \left(\mathbf{f}_{sp_n} + \mathbf{f}_{p_n o}\right) + \mathbf{f}_j, \tag{3.4}$$

where $\mathbf{r}$ is the heliocentric radius vector of the NEO, $\mathbf{f}_{kl}$ indicates the (reduced) force of k on l, s is the sun, o is the object, p_n is the nth perturbing body, n_p is the number of perturbing bodies and $\mathbf{f}_j$ are additional perturbations from relativistic corrections and cometary off-gassing, as required.

We can regard equation (3.4) as an instance of[19]

$$\frac{d^2\mathbf{x}}{dt^2} = \mathbf{f}_M + \mathbf{f}_j, \tag{3.5}$$

where $\mathbf{x}$ is the (heliocentric) radius vector of the NEO, $\mathbf{f}_M$ are the modeled (reduced) forces, and $\mathbf{f}_j$ are other perturbations, including theoretical (relativistic) corrections and the physical perturbations. An interesting feature of this case is that it allows relativistic effects to be taken into account without shifting outside of the theoretical framework of newtonian mechanics. This is accomplished by including the relativistic corrections as another force, effectively shifting the newtonian model that would have been picked out by newtonian mechanics proper given specified physical forces. The result of this shift, however, is nevertheless still a **𝔪𝔬𝔡𝔢𝔩** of newtonian mechanics in the sense of the semantic view; it is just not the **𝔪𝔬𝔡𝔢𝔩** that the theory of newtonian mechanics predicts under the circumstances.

This appears to pose a problem for the semantic view, which defines observational or empirical adequacy in terms of the *existence* of an embedding of the (data model)

[19]To match the form of (3.4), which uses reduced, *i.e.*, mass scaled, forces, the mass is suppressed in this equation.

behaviour of a phenomenon in a 𝔪𝔬𝔡𝔢𝔩 of the theory. What this implies in this case is that long before general relativity, newtonian mechanics was observationally adequate even for the precession of the perihelion of Mercury, since there exists a newtonian 𝔪𝔬𝔡𝔢𝔩 in which that behaviour can be embedded. But this is clearly a mistake, since newtonian mechanics was *not* observationally adequate because any model *constructed* using newtonian mechanics could not account for 43″/century of Mercury's precession, no matter how accurately the solar system bodies were modeled. Correcting the error that the semantic view makes here appears to require appeal to the symbolic form of newtonian mechanics, which is required to construct models of phenomena and which points to particular 𝔪𝔬𝔡𝔢𝔩𝔰 in given modeling circumstances. But this would turn a purely semantic approach into a logico-structural one. If this observation is correct, then it shows that a real theory cannot actually be specified independently of a linguistic formulation, which provides an argument that a logico-structural approach is actually *required* to adequately represent scientific theories.

Before turning to the next variety of complexity handling, let us summarize some of the important structural features that are brought out here. We may consider the modeling process to begin by the identification of a number of basic constraints, which may come from theory or may come from abstraction of macroscopic constraints from empirical knowledge of the structure and behaviour of the phenomenon being modeled. Since the latter are commonly required for feasible model construction, this shows that the construction of equations of motion is not simply a derivation from theory. And an interesting feature in the latter case is that empirical constraints can be imposed *before* any theoretical constraints have been chosen. These constraints are added to the basic underconstrained system of equations of the theory, along with additional constraints on the kinematical and dynamical behaviour of the system, resulting in a well-constrained system, a model. The additional constraints required to construct the model may also require input from experience if the theory does not provide the required constraints. The importance of constraints that are not part of the conceptual resources of the theory itself is a feature that is entirely absent from traditional classical accounts of theory application and shows that the view of model construction as a derivation from theory is inadequate.

There is another important classical-feasible distinction that arises here from the

point of view of models of phenomena. Classically, the basic theoretical constraints determine an enormous space of models, or spaces of spaces of models, which are represented on the $\mathcal{A}$-$\mathcal{G}$ perspective in terms of the unspecified geometric interpretation of algebraic constraints. Classically speaking, when the theory is applicable to a given phenomenon (at a given scale), this space is understood to contain a base model of the phenomenon in question. From a feasibility perspective, however, these models are generally inaccessible, and what matters are those phenomena for which we can actually construct valid models. The extent to which one can select or construct valid behavioural constraints determines a quite different space of models, *viz.*, those that can be feasibly constructed. Thus, in addition to the classical, fixed range of applicability of a theory, there is also a feasible, and variable, *modeling range* of a theory. The extent of this modeling range will vary depending on the tools applied mathematicians have available to feasibly construct constraints that yield valid models in a given theoretical framework. We will consider this issue again in chapter 6.

3.2.2 Handling Complex Dynamics

In the previous subsection, we considered the methods used to reduce the complexity of the *structure* of the real physical system, resulting in significant reduction in the structure of the model we need to deal with in order to study the dominant behaviour. There is another kind of complexity that must be handled properly for effective modeling, *viz.*, complex dynamics. Both of the models constructed above (for the double pendulum and orbit modeling), and the physical systems modeled, exhibit dynamical chaos (extreme sensitivity to initial conditions).[20] Chaos is not an all or nothing phenomenon, however. Some regions of the phase space of a dynamical system can be chaotic while others are not, and certain evolutionary phenomena can be studied in chaotic regions provided that the time scale of the process is shorter than the time scale where exponential divergence becomes significant, and even longer if accurate individual trajectories are not required, as we will see below. Both the double pendulum and the motion of objects in the solar system are examples of the former case, which is actually a property of most chaotic dynamical systems (Wisdom, 1987). The

[20]More technically, the typical standard condition for dynamical chaos is generic, bounded exponential divergence of trajectories.

issue of the time scale over which chaotic effects become significant is a particular concern in the study of the orbits of NEO. Accordingly it is important to be able to determine the time scale over which chaotic behaviour becomes evident or significant. The double pendulum, however, easily exhibits significant chaotic behaviour, which requires a different approach, *viz.*, the asking of different questions.

The problem of estimating the impact probability of a NEO requires that we are able to accurately predict the orbit of a NEO for a significant period into the future. This problem is not solved merely by ensuring that the data collected from the object has certain qualities, since the orbit of the object itself must be stable. This is an issue in the solar system because the orbits of objects are, in general, chaotic. Chaotic behaviour is actually a ubiquitous feature of hamiltonian dynamical systems, so it is a quite general concern for dynamical phenomena (Wisdom, 1987). Consequently, the ability to compute the (first) Lyapunov exponent λ, which characterizes the unpredictability for a dynamical system, for an object is essential in order to be able to know the characteristic time over which chaotic effects will become significant. The Lyapunov exponent measures the maximal log-rate of separation of infinitesimally close trajectories. When $\lambda > 0$, nearby trajectories (on average) diverge exponentially and the dynamics are chaotic. The size of λ, then, determines the rate. The time τ_e taken for the separation of two trajectories to increase by a factor of e is just the reciprocal of the Lyapunov exponent:

$$\tau_e = \frac{1}{\lambda}.$$

This is called the *e-folding time*, and is one measure of the characteristic time scale over which chaotic effects will begin to become significant. A measure of the characteristic time scale of chaos is called a *Lyapunov time*. The e-folding time is one definition of this among others, but the typical one in celestial mechanics. Consequently, we will refer to the e-folding time as *the* Lyapunov time.

One of the significant features of chaotic motion for feasible epistemology is that it implies that precise trajectories can be tracked theoretically only locally to a certain period of time. Since chaos is a general property of hamiltonian dynamical systems, this is not merely a special consideration for some cases. Thus, the generic case

should be that knowledge is *local* to a limited range of conditions. The logical approaches do not provide a conceptual language for such local knowledge. This could be seen as a result of the theory T syndrome, given the fact that even the classical existence and uniqueness theorem for nonlinear dynamical systems only apply locally. Local knowledge is a general issue with nonlinear problems, not even just for chaos. Consequently, an adequate conceptual framework for describing our knowledge of the world should deal generally with local knowledge, which is precisely what the local constraint logic we are using in this study does.

There is also an issue here for the view that model construction is deductive. For real solar system bodies, particularly NEO, analytic computation of the Lyapunov time is not feasible. Consequently, numerical methods must be used to compute this time. This shows that a key part of feasible inference for nonlinear phenomena is to be able to estimate how long the results of a model will be accurate. This does not emerge as a significant concern on the semantic views considered in the last chapter. A common numerical technique used in NEO modeling is a tangent map method developed by Mikkola & Innanen (1999). This method computes the Lyapunov exponents in an analytic way as a byproduct of the main numerical integration.[21] This sort of numerical method, that allows analytic computations to be performed on the basis of numerical ones is called a *semi-analytic method*. The interest of this class of numerical method is that it involves a combination of analytic reasoning and numerical computing, in order to compute quantities not computable analytically to high accuracy. This sort of method, a hybrid of analytic and numerical methods, is quite common in practice and shows yet another way that a picture of the application of theory in terms of deductive inference is inadequate.

An interesting point made by Wiegert *et al.* (2007) is that chaotic behaviour of this sort does not necessarily mean that the model breaks down entirely. Provided that over time scales longer than the Lyapunov time only the uncertainties of an object *within its orbit* become large, then it is possible to study some of the properties of the object over long times, even getting fairly reliable determinations of quantities if they are stable across nearby orbits. This can be studied by numerically integrating an ensemble of objects with nearby initial conditions. Thus, the behaviour of an object

[21] This particular method will be discussed in more detail in chapter 5.

can in some circumstances be studied beyond its Lyapunov time scale at the expense of replacing a deterministic study by a statistical one. Thus, some degree of chaos, or nonlinearity, can be accommodated within the framework of celestial mechanics. Thus, feasible methods have ways of overcoming some of the limitations on knowledge imposed by nonlinearity.

Another interesting feature of this statistical approach is that it shows that by shifting how one thinks about getting information from the model, it can be possible to get useful information from the model even when it can no longer accomplish what it was intended to do. The double pendulum system provides a clearer example of this kind of phenomenon. It is primarily of interest as a theoretical and numerical study of chaotic behaviour in an easily tractable case. Accordingly, the ability to reliably and closely track trajectories is not the only, perhaps not even the central, concern. Even if we were interested in closely tracking the phase trajectory of a real double pendulum system, however, the strongly chaotic behaviour it exhibits under certain conditions would mean that we would have to abandon this aim in certain regions of phase space. In such a case, if we are to obtain useful information from the model we have to ask *different questions* (Corless, 1994*a*). And which questions we can usefully ask are ones where the answers are determined by *stable* quantities.

A chaotic dynamical system is characterized by its dynamical instability, instability determined by extreme sensitivity of the dynamics to variation. If the initial conditions of a chaotic dynamical system are changed slightly, nearby trajectories will diverge at an exponential rate. But, such a dynamical system is also unstable in another sense: if you were to slightly perturb the dynamics (vector field) of the system you would also see exponential divergence of nearby trajectories. This is to say that the dynamics is unstable both in terms of small changes of the initial conditions and in terms of small changes of the dynamics. The reason this matters is that both the analytic model reduction techniques and the numerical solution methods involve variations of the initial conditions and/or the model (vector field). Accordingly, only quantities of the target system that are *stable under such variations* can be feasibly studied using mathematical modeling.

This, then, leads to the question as to what quantities are stable across variations of the model. One likely candidate for a stable quantity over variations of the model

are the Lyapunov exponents. Another is the regions of phase space where the dynamics is chaotic, in the sense that we expect that conditions under which the system becomes chaotic will not change significantly with a small change in the model (vector field). It is for this reason that quite simplified or idealized models, provided that the simplifications or idealizations are not too severe (in an imprecise sense), can still give us useful information about the chaotic behaviour of a real double pendulum system, even when solved numerically. The stability of quantities over variations of the model generally requires proof, but even when proofs are not available *evidence* that a quantity is stable can be obtained by considering an ensemble of nearby models and computing the same quantities to see if they are stable. In this way, a statistical study of the chaotic behaviour of the double pendulum system can be carried out, quite analogously to the study of solar system objects beyond their Lyapunov time scale. Since, classically, models are typically judged in terms of the accuracy of their solutions, chaotic dynamics would in this case be regarded as a breakdown of the model. What this technique of strategically changing the question shows is that looking to quantities that are stable under quite a broad range of variations allows a model that would be viewed as useless from a classical perspective to give useful information about the phenomenon.

Before moving on, it is worth mentioning that another important technique used to study nonlinear dynamical systems is *linearization*, which we have already seen illustrated above on the simple pendulum system. Often the dynamical behaviour of a nonlinear system can be studied over small (in a sense that is relative to the degree of nonlinearity)[22] distances in time and space. This is effectively what is done by tracking a NEO within its Lyapunov time. A useful way of studying this analytically, however, is to expand the vector field in a Taylor series and to retain only the linear term. This yields a linear system which is much more amenable to formal analysis. Thus, a nonlinear system can often be replaced by a linear one at the expense of having a model that is only locally valid in phase space (near the point where the vector field is truncated). Linearization is a very important technique and it finds quite widespread usage as a means of local analysis. It is also a kind of

[22]This can be measured roughly by the Lipschitz constant of the vector field (see *e.g.*, Henrici, 1962, 15*ff.*).

model reduction technique, since linearizing the vector field changes the model from a nonlinear one to a linear one.

There are a wide variety of techniques in applied mathematics that follow the basic algebraic approach of expanding some nonlinear quantity in a power series and retaining only up to the linear term. The linear stability analysis of dynamical systems, for example, uses this technique, as does the standard analytic technique for the computation of the Lyapunov exponents of a dynamical system. So, indeed, do the renormalization group methods used in statistical mechanics for the computation of critical exponents of a system near its critical point, a case considered in detail by Batterman (1998, 2002*b*). Geometrically speaking, this technique of linearization is useful anywhere where over some limited region a *curved* quantity can be accurately approximated by a *flat* one. We will see a variety of kinds of linear methods when we consider the construction of algorithms used in NEO modeling in the following subsection and in the algorithms for computing impact probability in chapter 5.

3.2.3 Handling Complex Computations

So far we have seen methods of handling two kinds of complexity: analytic (algebraic) and structural (geometric) simplicity methods for reducing the complexity of the model needed to capture the dominant behaviour of the phenomenon; and localization methods for handling nonlinear dynamics. A third kind of complexity that needs to be handled arises from difficulties in computations associated with the need to develop efficient algorithms to find solutions to a mathematical problem. In the case of dynamics we are considering, this involves efficiently computing trajectories through the system space of a model. Since analytic methods are quite limited for nonlinear problems, useful analytic algorithms for the solution of mathematical models are rare in practice. Given the steady increase since the 1940s of the power of the computing resources available to scientists, numerical methods implemented on physical machines have been an increasingly important tool for the solution of mathematical problems. The algorithms that numerical methods use are extremely complex in terms of size of the calculations, and most are well beyond the patience of any human being to carry out by hand. Accordingly, one important way we handle the complexity of these computations is to have a machine carry them out for us.

There is another sort of algorithmic complexity reduction that is used in mathematical modeling, however, and it is one that is properly part of the reasoning process used in model construction.

Both the double pendulm and the NEO modeling cases require numerical methods, mainly because of the nonlinearity of the models but also because of limitations of human patience for complex analytic manipulation. In the double pendulum case, at least in terms of how we are considering it, only machine implementation of numerical methods is required in terms of algorithmic complexity handling.[23] In the NEO modeling case, however, the complexity of the analytic algorithm to handle the details of orbit determination and impact probability computation makes certain analytic simplification methods useful for ensuring that the computations can be carried out efficiently on a machine.

Analytic Algorithm Construction in NEO Modeling

To show how complexity reduction is playing a role in the analytic construction of algorithms to solve the orbit determination and impact probability problems for NEO, we first require a digression on spherical astronomy and celestial mechanics. Ordinary (non-ranging) observations of the positions of celestial objects in the sky are interpreted as points on an imaginary *celestial sphere*. A great many stars (sufficiently far away and which exhibit no apparent proper motion) are to a good approximation stationary on this sphere in terms of solar system observations. This allows the celestial sphere to function as a spherical star map. Solar system objects, which have large proper motions, move across this sphere and ground- or space-based observations map their position on this sphere in some spherical coordinate system. The most common ones are the *equatorial celestial coordinate system*, depicted in figure 3.4, and the *ecliptic celestial coordinate system*.[24] Since the celestial sphere is regarded

[23] This is because the full hamiltonian can be constructed explicitly and because numerical methods can be applied directly to Hamilton's equations, as we will see in chapter 5.

[24] Any spherical coordinate system requires the specification of a fundamental plane (for longitudes), a longitudinal point of origin, and a sense of helicity. The fundamental plane for ecliptic coordinates is the plane of the ecliptic, and for equatorial coordinates it is the plane normal to the Earth's rotational axis. The three celestial coordinate systems (there are also *horizontal coordinates*, usually from the perspective of an observer on Earth) take the vernal equinox as the longitudinal point of origin and are usually right-handed (usually defined by the upward normal picked out by the North Pole).

as being at infinity, the coordinate system for it is only two dimensional. In the equatorial system the longitude is called the *right ascension* and the latitude the *declination*. The right ascension is measured toward the east (counterclockwise) from the *vernal equinox*, indicated with the symbol ♈ in figure 3.4. The vernal equinox is one of the two points on the celestial sphere where the great circle of the celestial equator, corresponding to the plane perpendicular to the Earth's rotational axis, intersects the great circle of the ecliptic, corresponding to the plane of the Earth's orbit around the Sun.[25] When the Sun reaches the point ♈ of the vernal equinox on the celestial sphere, spring begins (in the Northern Hemisphere).[26]

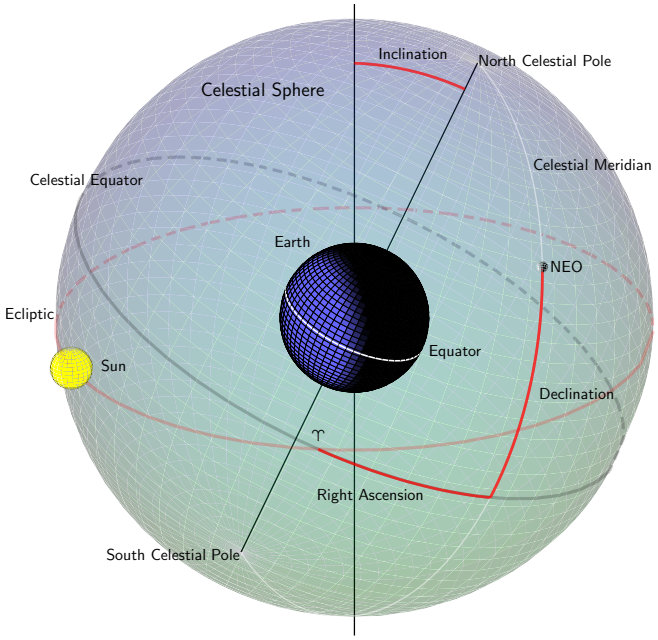

Figure 3.4: The (apparent/geocentric) equatorial celestial coordinate system.

Many issues arise due to the need to make changes between coordinate systems. One issue arises because actual observations are made at some particular point on the surface of the Earth,[27] and simultaneous precise measurements made by different observers can disagree if the object is close enough to the Earth for the parallax between the two points of observation to be significant. This is rarely an issue for stars

[25] Since the plane of the ecliptic is that of the Earth's orbit around the Sun, the Sun appears to travel counterclockwise across the celestial sphere along the ecliptic line once per year.

[26] Hence, we may see that the position of the Sun in figure 3.4 is around January 21.

[27] This requires the use of a fully realizable coordinate system, such as the astronomical geographic coordinate system, which complicates matters further. This issue will be considered in more detail in chapter 4.

and other more distant objects, but can be an issue for objects moving in the solar system, including NEO, where the parallax is non-negligible. A related issue is the need to shift between the topocentric (Earth's surface) coordinates in which observations are made, and geocentric, heliocentric and barycentric coordinates. These shifts involve much more significant parallax, which must be handled with care. And since most coordinate systems are ideal in the sense that they rely on models, handling error propagation in coordinate transformations is quite a complex affair. We will postpone any detailed consideration of these issues until chapter 4, but it is important to understand the complexity of considerations required in real cases.

One of the fundamental problems in dynamical astronomy is how to connect observations with predictions from theory. The way that space and time are handled in observations is very different from the way that they are handled in theory. In terms of time, theoretical time is perfectly uniform but no physical device used to measure time can perfectly realize this. Fortunately, the availability of atomic clocks makes this an issue only for extremely precise measurements, such as those used in the testing of theories. Consequently, we will ignore the matter of time measurement for our purposes. The trickier issues to do with time arise due to the fact that coordinate systems are not inertial and so are specific to a time (fundamental epoch) and so careful modeling is involved in order to correct for non-inertial coordinate systems, a problem compounded with changes of coordinates.

It is the issues of space, however, that are more our present concern. In terms of space, observational locations are two-dimensional and theoretical positions are three-dimensional.[28] To keep this distinction clear, we refer to two-dimensional angular observations as *locations* and the position of an object in space is its *position*. Observations in an equatorial coordinate system (like that illustrated in figure 3.4) can be expressed as a cartesian vector

$$\mathbf{l}(\alpha, \delta) = \begin{pmatrix} \cos\alpha \sin\delta & \sin\alpha \sin\delta & \cos\delta \end{pmatrix},$$

giving a direction in space but no distance from the origin (see figure 3.5). Theory uses three coordinates to fix the position of an object in space, often in terms of a

[28] This issue does not arise in the same way when ranging is possible, like with radar in some cases.

radius vector $\mathbf{r} = (x, y, z)$ in some cartesian coordinate system. If the origin of the observational and theoretical coordinate systems is the same, then the two are related simply by

$$\mathbf{r} = r\mathbf{l}(\alpha, \delta). \tag{3.6}$$

The main trouble here arises from the fact that direct measurements typically do not determine r, together with the fact that in the typical case the observational and theoretical coordinate systems have different origins (*e.g.*, geocentric and heliocentric coordinates, as shown in figure 3.5). This means that measurements cannot be converted into the theoretical coordinate system when parallax is non-negligible, even just as a heliocentric location $\mathbf{l}(\lambda, \beta)$,[29] because such a conversion requires knowing not just the distance between the origins of the coordinate systems, which we know, but also knowing the distance to the object (r in the equivalent of equation (3.6)), which we do not.

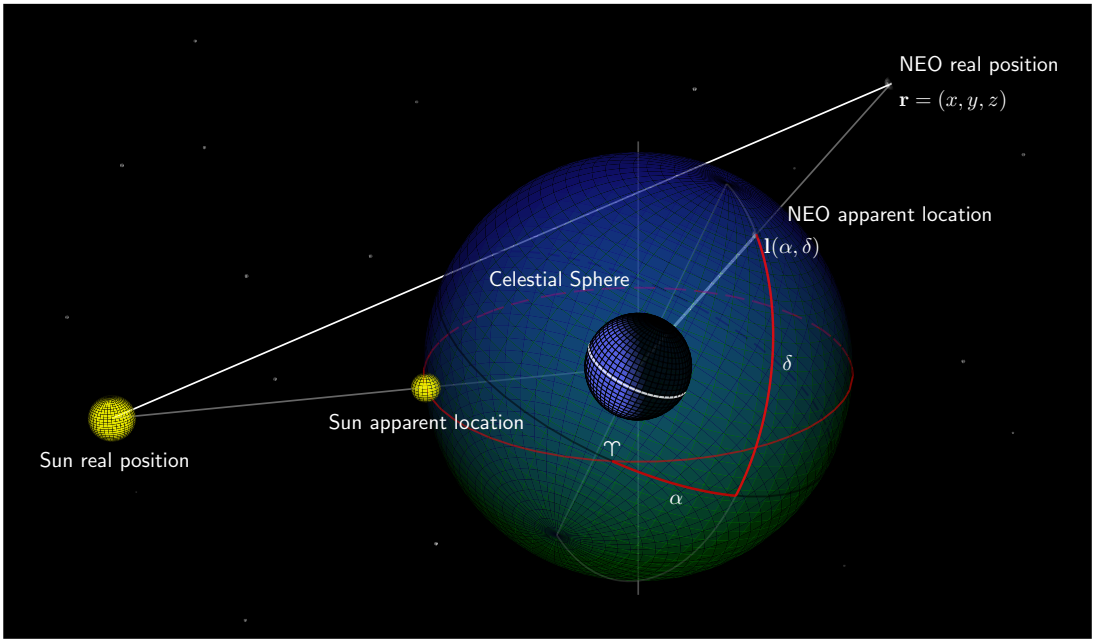

Figure 3.5: The apparent (measured) position $\mathbf{l}(\alpha, \delta)$ in a geocentric equatorial coordinate system and real (physical) position $\mathbf{r} = (x, y, z)$ in a heliocentric cartesian coordinate system of a near-Earth object. This illustrates a typical sort of case for NEO modeling where the observational and dynamical quantities are given in two different coordinate systems with the two origins separated by a significant distance, giving rise to parallax.

[29] Heliocentric locations are naturally measured in ecliptic coordinates, because the natural fundamental plane for a spherical coordinate system is the plane of the ecliptic. The ecliptic longitude is represented by λ and the ecliptic latitude by β.

The basic strategy of an orbit determination algorithm is then the following. The first step is to guess a possible heliocentric position, say $\mathbf{r}_0$, for the object and then calculate what the observed location, say $\mathbf{l}_0(\alpha, \delta)$, would be for an object in that position. We can do this calculation because we can calculate the distances between the origins of the coordinate systems (*e.g.*, Earth-Sun distance and centre of Earth-topocentric place distance) and have formulae for converting quantities in one coordinate system to another. This coordinate transformation is accomplished by the *measurement reduction function* $R(\mathbf{x})$, which will be a major concern in the next chapter. The measurement reduction function takes real phase space positions as input and outputs measured location in a given celestial coordinate system. Given the position $\mathbf{l}(\alpha, \delta)$ from an actual measurement, we can then compute the *measurement residual*[30]

$$\boldsymbol{\xi} = \mathbf{l}_0(\alpha, \delta) - \mathbf{l}(\alpha, \delta) = R(\mathbf{r}_0) - \mathbf{l}(\alpha, \delta).$$

If the guess $\mathbf{r}_0$ turned out to be the exact position of the object then we would find $\boldsymbol{\xi} = \mathbf{0}$. This, of course, never happens in practice, so we look to minimize the value of measurement residuals.

This is the basic strategy. As it stands it could never succeed. First of all, we need to fit at least three observations (giving six quantities) in order to fix the dynamical trajectory (position and momentum). But the probability of obtaining three observations consistent with a unique orbit is zero. Measurement and modeling error ensure that this never happens. Consequently we take a large number of observations and look for an optimal solution of best fit. The method of least squares is typically used to solve this highly overdetermined problem. The most important determinant of the accuracy of a solution (assuming quite precise data) is the time span over which the object was observed, which is called the *data arc*. The length of the data arc is often characterized by the number of *oppositions*, passages by the Earth-Sun line, over which the object was observed.[31] Data arcs several oppositions long usually lead

[30] The measurement residual would otherwise be represented with the symbol '$\boldsymbol{\delta}_{\boldsymbol{\mu_b}}$', but to prevent confusion with the symbol 'δ' for the declination, we use the symbol '$\boldsymbol{\xi}$' instead.

[31] This is because the object is most easily observed when in the vicinity of the Earth-Sun line. For comets the number of *apparitions*, *i.e.*, the number of times the comet passes aphelion, is also important. This is because approaching and receding from perihelion the comet is lit up by the solar wind, causing it to off-gas and acquire a tail, thereby making it much easier to observe. Passing aphelion thus marks the end of one cycle off-gassing and the beginning of the next.

to high accuracy orbits, within the Lyapunov time of the object.

Rather than using position $\mathbf{r}$ and velocity $\dot{\mathbf{r}}$ to pick out the dynamical state of a celestial object, it is common to use a set of six two-body orbit parameters. A common set of these is the classical keplerian set: eccentricity e, perihelion (periapse) distance q, time of perihelion (periapse) passage T_p, right ascension of the ascending node Ω, argument of perihelion (periapse) ω, and inclination to the ecliptic (plane of reference) i (for illustrations of these parameters see figure 3.6). The virtue of these parameters is that most of them are usually slowly varying and so provide a clear geometric interpretation of the orbit. Even if the dynamical coordinates $(\mathbf{r}, \dot{\mathbf{r}})$ are used in the numerical integration of the orbit, six orbital parameters may be used to specify the initial conditions of the orbit, the conditions that render the dynamical model well-determined, and may be sought at future times when the orbit is integrated forward in time to detect close passages to Earth.

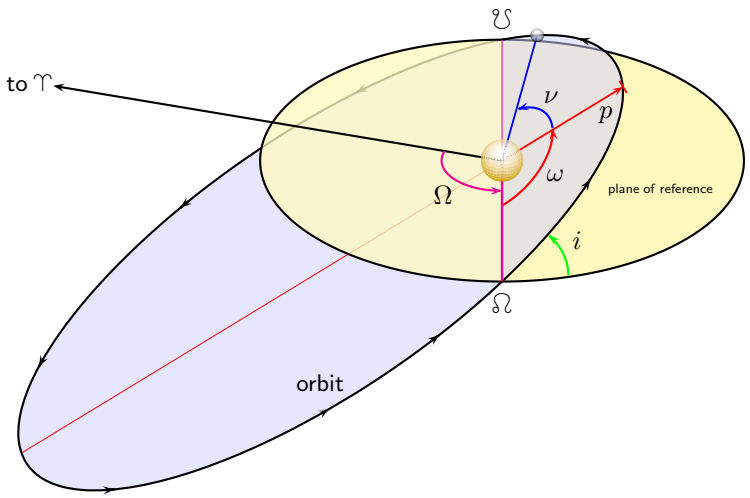

Figure 3.6: Orbital parameters: p (periapse distance), Ω (longitude of the ascending node), ω (parameter of periapse), i (inclination), and ν (true anomaly). The position of the object on the orbit is determined by ν, which is a function of time, or T_p the time since the last periapse passage. The direction of the vernal equinox is indicated by ♈ and the ascending and descending nodes are indicated, respectively, by ☊ and ♉ respectively. The parameters Ω, ω and i are affected by the direction of the vernal equinox because the plane of reference is not inertial; hence these parameters are relative to a particular epoch, the standard now being J2000.

Stated in these terms, the basic orbit determination problem is to estimate the six orbital elements $\mathbf{k}_0$ at some epoch t_0 given a set of n astrometric measurements $m_i(t_i)$, which can be very large.[32] The strategy of least squares solution to this highly

[32]The symbol $\mathbf{k}$ picking out a vector of orbital parameters can be regarded as standing for

overdetermined problem described by Chodas & Yeomans (1999) is to minimize the sum ξ_μ^2 of the measurement residuals $\boldsymbol{\xi}_{\boldsymbol{\mu}_{b,i}}$ for each of the n measurements, weighted by the measurement errors. These are computed by using the evolution model (3.4) to iteratively determine what the time t_{ie} of emission of the light must have been to have been observed on Earth at a particular time t_i, and then computing the residuals for each of the measurements using the measurement reduction function $R_{\mathbf{k}_0}(\mathbf{x}(t_{ie}))$, which depends on the initial conditions $\mathbf{k}_0$ chosen for the evolution model. But since this whole process depends in the choice of initial conditions, and the initial conditions are what we seek as the solutions, this whole process must be repeated as the initial conditions are varied in the search for a minimal solution. The process, thus, must start by seeding the evolution model with a guess $\overline{\mathbf{k}}_0$ of the initial conditions, typically calculated using a two-body approximation for the orbit of the NEO. The solution is then sought by examining small variations around the reference value $\overline{\mathbf{k}}_0$.

This concludes our digression on spherical astronomy and celestial mechanics. We may now consider methods of computational complexity reduction in NEO modeling. One such complexity reduction method is used to simplify the algorithm for seeking the minimal solution to the orbit determination problem. In order to apply a linear least squares method to this highly nonlinear problem (nonlinear dynamics, nonlinear measurement reduction function), it is assumed that the variation of the initial conditions is small enough that the measurement reduction function $R_{\mathbf{k}_0}(\mathbf{x}(t))$ can be treated as a linear transformation, *i.e.*, a matrix. This matrix $\mathbf{R}_{\mathbf{k_0}}$, called the *measurement matrix*, is the jacobian calculated from the reduction function. This can be thought of as a structured, local velocity of the reduction function R. It is the local "rate of change" of the reduction function under changes in the six dimensional control space of orbital parameters $\mathbf{k}$. The effect of this linearization procedure, is that it produces another constraint that the solution $\mathbf{k}_0$ must satisfy, *viz.*, that the difference between the measurement residual $\boldsymbol{\xi}$ and the product of the measurement matrix and the variation of the initial condition must be of the order of the measurement error or smaller. This enables us to see that the strategy of the least squares method is to convert an overdetermined problem (much more data than is consistent with a unique solution) into an underdetermined one (many solutions with small residuals)

k*eplerian*.

and then to add another constraint so as to render the problem well-determined, or *effectively well-determined*, in the sense that the additional constraint is satisfied to the order of the measurement error or better.

That this case of linearization involves the variation of parameters, *i.e.*, components of the vector $\mathbf{k}_0$, of a vector function, *i.e.*, $R_{\mathbf{k}_0}(\mathbf{x})$, in a control space obscures the fundamental similarity to the cases of linearization considered above. The least squares algorithm requires that we know how the nonlinear measurement reduction function *responds to changes* in the initial conditions. The linearity assumption used is essentially that the response of $R_{\mathbf{k}}(\mathbf{x})$ to changes in $\mathbf{k}$ is *constant* sufficiently nearby to the initial value of $\overline{\mathbf{k}}_0$. Another way of thinking of this is that it is an assumption that the *curvature* of R (determined by variations of $\mathbf{k}$) is negligible, *i.e.*, that R is *locally flat*, in a small region of $\mathbf{k}$-space that includes the optimal solution. Achieving convergence then validates the assumption of local flatness. The validity of this assumption means that the "local rate of change" of $R_{\mathbf{k}}(\mathbf{x})$ at $\mathbf{k}_0$ is the same at nearby values of $\mathbf{k}$. This "local rate of change" is just the measurement matrix $\mathbf{R}_{\mathbf{k}_0}(\mathbf{x})$. In this respect, this case is not all that different than a case where the acceleration (a measure of curvature over changes in time) of a moving object is small enough that we can assume that the velocity (linear response to changes in time) of the object is constant for nearby times.

The final component of the NEO tracking problem, the estimation of an impact probability, also often relies on a number of linear approximations. The basic idea of the algorithm to solve this problem is the following. First of all, the solution to the orbit determination problem yields an estimate of the uncertainty in the object's phase space position, using a six dimensional gaussian distribution. The spatial component of this is an ellipsoid, called the *confidence ellipsoid* (Milani *et al.*, 2002). In order to determine whether the object is likely to have an Earth impact in the future, the precision orbit is used along with JPL ephemeris data to detect any close approaches with Earth in the future. If a close approach is detected, then another algorithm is used to compute an impact probability. The basic idea is to evolve the uncertainty ellipsoid forward in time to determine a probability density that the close approach will result in an impact. This is done by computing the probability density in a plane intersecting the Earth and the orbit of the object, called a *target plane*, and

the impact probability is the total probability weight over the disk of the Earth in the target plane.

Several ways of simplifying this calculation that are useful under certain circumstances also use linear approximations. For example, one approach uses the computation of the orbit, along with a method like that of Mikkola & Innanen (1999) mentioned above in connection with Lyapunov exponents, to compute a linear approximation to the map from the epoch of the orbital parameters $\mathbf{k}_0$ to the time of impact. This is simply a linear map of the confidence ellipse from the initial time to the impact time. This then provides the ability to compute the marginal probability density of the object in the target plane, from which the impact probability can be estimated. In a similar way that linearization results in a localization, this linear method is limited in that it is only valid under certain special conditions. One significant limitation is that it is only reliable for high probability collisions, which makes it useful only in certain special circumstances. So, a linear approximation, in general, leads to a simpler calculation, but one that is only valid locally to certain special conditions. The price of the computational speed-up is a restriction in the domain of validity of the calculation.

Thus, we see that not only do methods of model construction and methods of handling complex dynamics require analytic methods of complexity reduction, so sometimes do the analytic methods used to construct algorithms to solve a model problem. We will see that similar kinds of methods of complexity reduction are used in the construction of numerical methods, which are really just analytic methods of model solution of a different kind.[33] We also see a clear pattern in the manner in which linearization is used to handle nonlinearity locally. This is a feature that comes out quite clearly on the $\mathcal{A}$-$\mathcal{G}$ perspective, *viz.*, the fundamental algebraic *and* geometric pattern of assumptions underlying linear methods in quite different areas. Algebraically, linearization involves truncating a Taylor series after the linear term. Geometrically, linearization involves an assumption of local flatness. Given that functions commonly have the property that they are close to constant on sufficiently small scales, it is natural that linearization techniques will yield valid results

[33]Understanding this requires distinguishing numerical methods, which are analytic, from their machine implementations, which are not analytic.

if the nonlinearity is not too severe. Given that the linearization technique can be applied with this sort of understanding, it allows the applied mathematician to look out for conditions where the algebraic assumption will fail. This gives the technique a kind of reliability that is explained in terms of understanding the covariation of the equations and their interpretations under linearization. Thus, the linearization technique is explained in terms of a correlated syntactic and semantic transformation, and one that is only locally (effectively) valid. This shows that the explanation requires representation of both the syntax and the semantics of the method, which the standard logical approaches cannot do.

The need to use linearity assumptions or approximations in the construction of algorithms to solve model problems also shows another way in which a view of theory application as deduction from first principles is lacking in accuracy. The classic aim for the solution of differential equations are analytic methods that can yield a closed form solution in cases where a unique solution is guaranteed by existence and uniqueness theorems. The cases in this subsection show that in cases of non-negligible nonlinearity (non-constant variation) methods of this sort are not available. One reason for this is just that analytic methods of solution for differential equations require either a high degree of symmetry or linear equations. That most realistic problems lack both a high degree of symmetry and linearity means that these methods leave us without analytic solutions to realistic problems. The case considered in this subsection shows another way in which complexity pushes beyond the capacity of purely analytic methods, namely in the consideration of overdetermined problems taking large amounts of data as input. Not only do efficient methods for constructing algorithms to solve such problems require approximate, non-deductive, methods like linearization, the algorithms must be supported by an array of numerical methods that can handle not only the nonlinear equations but also the large calculations resulting from the large amount of data input into the problem. Thus, a central part of the modeling process, *viz.*, model solution, classically viewed as a deduction from laws and initial/boundary conditions is revealed to be heavily entangled with methods of approximation and computation. We are now in a position to see that the structure of these approximation methods is screened out by the pure $\mathscr{L}$-view and $\mathfrak{S}$-view, but is captured readily and naturally with an $\mathcal{A}$-$\mathcal{G}$ perspective. When we turn to the

consideration of numerical methods in chapter 5, we will see that the same is true for methods of computation.

3.3 Sources of Model Error and Their Structural Variation

On an $\mathcal{A}$-$\mathcal{G}$ perspective, any error affects not only the geometric structure of a model, or geometric interpretation of a constraint, but also affects the algebraic presentation of a model or constraint. And since the algebraic presentation and the geometric interpretation are understood to be covarying, any geometrical variation implies a specific algebraic variation and *vice versa*. Accordingly, we require a clear and stable way of assessing the algebraic and geometric effects of the various sorts of error that arise in the process of model construction. It turns out that in applied mathematical practice there can be many ways of analyzing error for a single model, which results from the fact that there is an element of *choice* in how error is analyzed for feasible models. This, I shall argue, reflects the fact that error analysis is actually part of the modeling process. In this section we consider three kinds of error that arise in a mathematical model and its process of construction. In the subsequent two chapters we consider kinds of error that arise in data handling and numerical computing, respectively.

Assessing the nature of these kinds of error from our $\mathcal{A}$-$\mathcal{G}$ perspective that I have adopted for this study, requires considering the concept of *validity* as it is used in applied mathematics. Typically this term is used to describe a model that is shown to be descriptively accurate. Where philosophers would often speak of a model that is shown to be observationally or empirically adequate as being satisfied, applied mathematicians often describe such a model as valid. This reflects the fact that the model construction and reduction methods that applied mathematicians use rely on a variety of approximations and so they rarely consider their models to be descriptively accurate in a strong sense. In this way the concept of *validity* from applied mathematics connotes *justification* and *effectiveness* for certain modeling purposes, rather than being strongly tied to truth. Accordingly, the applied mathematical concept of validity appears to be quite different from the concept of *validity*/𝔳𝔞𝔩𝔦𝔡𝔦𝔱𝔶 used in philosophy, the latter acquiring its meaning from its logical meaning in terms of

truth-preservation. As we will see below, however, the concepts of local constraint logic provides a way of resolving this potential ambiguity.

The term *valid* is usually used in an informal sense in applied mathematics, but in some cases it is used with a more precise meaning. In the context of a book on numerical polynomial algebra, Stetter (2004) presents a technical definition. The purpose of this definition is to track the effect on an algebraic theory or model of parameters or quantities that are determined empirically, and hence are not specified exactly. Since Stetter introduces this definition for the purposes of handling error in the context of an algebraic theory, with a specific geometric interpretation, that is formally exact, it is very well-suited to the purposes of error handling on an $\mathcal{A}$-$\mathcal{G}$ perspective of theories.

The problem with a formally exact theory that is representing an empirical structure or quantity is that the mathematical description can have more precision than the empirical structure or quantity it represents. To illustrate this, consider the case of an eigenvalue equation

$$Ax = \lambda x,$$

where A could be a hermitian operator, presented as an $n \times n$ matrix, and its eigenvalues λ the spectrum of possible values of some physical quantity. Whether a particular real number λ is an eigenvalue of the hermitian operator A, or whether a given n-vector v is an eigenvector, is a precise mathematical question, and has a precise yes or no answer. But if the values of the matrix elements of A are known empirically, then A is an imprecise quantity and it is not possible to answer this question with a clear yes or no. The imprecision in the empirical operator A is tied up with how it is measured, *e.g.*, what devices are used, how the measurement is modeled, *etc.*, and with what is at issue, *e.g.*, how precise the measurement needs to be given the modeling aims and interests. The result of this, as Stetter argues, is that the imprecision cannot be represented as a *precise* range of acceptable values, because we cannot justify an exact boundary for this range. This is a sort of imprecision that will occur whenever mathematics and data or phenomena interact.

The solution to this is to introduce a *validity value*, *i.e.*, a positive real number that provides a continuous transition between clearly affirmative and clearly negative answers to whether or not a particular measured value is admissible for a given

empirical quantity. The idea is that empirical quantities α are represented as pairs $(\overline{\alpha}, \varepsilon)$, not only having a specified value $\overline{\alpha}$ but also a *tolerance* ε that determines the order of magnitude of the imprecision in the specified value. The tolerance provides a rough boundary for what counts as a valid value of the quantity. Suppose that $\tilde{\alpha}$ is a measured value of the quantity α. Then the *validity value* is the quantity

$$\frac{|\overline{\alpha} - \tilde{\alpha}|}{\varepsilon}.$$

If the validity value is less that 1, this means that the measured value is equal to the specified value to within the tolerance ε, in which case $\tilde{\alpha}$ is certainly valid. But, 1 is not a hard boundary, and depending on the error demands of the context, measured values with much larger validity values could be considered valid. In general, however, the larger the validity value, the less likely that the measured value is valid and beyond a certain point, which will depend on contextual modeling considerations, values will be clearly invalid.

As an example, suppose that a model heart rate in a microprocessor of a pacemaker has a specified value of 80.1 bpm and an indetermination in the last digit, the tolerance ε will be of the order of 0.1. As a result, measured values of 80.2 or 80.0 bpm will have validity values of 1 and, hence, will be valid values of this quantity. So also may be the occasional larger deviation, but a value of 81.3 bpm, with a validity value of 12, will be invalid. In general this indicates the (local) failure of the model, but in this particular case it means that the microprocessor needs to adapt the model to the change in measured values.

The boundary around the specified value of an empirical quantity is rough for the same reasons that empirical structures or quantities are imprecise, *viz.*, it is bound up with how things are measured and what is at issue in the modeling. The upshot of this for empirical quantities is that values of empirical data are *tied to a scale*, *i.e.*, they are only *defined* up to a certain precision.[34] The manner in which empirical quantities are tied to a scale is similar in certain respects to how models that make use of empirical constraints become tied to the scale at which those constraints are

[34]Perhaps this could be construed in terms of asymptotic equivalence, *viz.*, differences matter less and less the higher the precision and beyond a certain level of precision, differences are not empirically meaningful and different values are empirically equivalent.

valid. For instance, rigid body constraints can be effectively true at a macroscopic scale, but cease to be so at mesoscopic and microscopic scales. The roughness of the outer boundary reflects the fact that error considerations can be quite complex in general, and so it can be difficult to have such an understanding of sources of error that a very precise tolerance can be set. Accordingly, the boundary of validity of an empirical quantity is imprecise, with the validity value being a reliable guide rather than a strict determination.

Now, this concept of a valid value of a quantity can be seen to be picked out by the concept of effective validity from local constraint logic. Consider a real number α and suppose that our constraint system is simply the equation $\alpha = \overline{\alpha}$ that specifies the value of α to be $\overline{\alpha}$. Now, suppose that we allow an error of around ε in the value of α. Then the values of α that are valid in the sense due to Stetter (2004) are precisely those that satisfy

$$\Big\|\Big|_{\alpha=\overline{\alpha}}^{\varepsilon}\ \alpha = \tilde{\alpha},$$

viz., the effectively valid values of α. This shows that the sense of 'valid' in applied mathematics bears a much stronger relation to the concept of local effective validity than the concept of validity from mathematical logic, since the applied mathematical concept of validity actually *is* a special case of the concept of local effective validity from local constraint logic. Thus, the concept of validity from applied mathematics is naturally incorporated into local constraint logic. And we may now see that rather than having a unique valid value, like mathematical quantities typically do, empirical quantities have a range of effectively valid values.

This simple model of effective validity for a numerical quantity is straightforwardly generalized to structured quantities like vectors, matrices, functions and operators.[35] The complexity that is introduced with this more complex case is that in addition to the specification of a specified value and a tolerance, we need to identify a specific measure of the *size* of errors and a specific measure of the *distance* between values. As we saw in section 3.1.3, the size of a mathematical quantity is specified by a norm, but there is no canonical norm for structured mathematical quantities. Once a norm has been specified, however, there is a canonically induced or natural measure of dis-

[35] I will use the term *quantity* to denote mathematical structures that, at least in principle, can be measured.

tance between quantities.[36] Thus, extending the definition of 'validity' to structured mathematical quantities requires not simply a specified quantity $\overline{\alpha}$ and a tolerance ε, but also a norm $\|\cdot\|$ defined for quantities of the given type.

3.3.1 Measurement Error and Its Structural Variation

The most straightforward of the kinds of error that arise in model construction and analysis is error arising from measured quantities appearing in a model. Whenever a measured quantity appears in a constraint, it is, strictly speaking, automatically an empirical quantity. This means that it will be capable of variation, and we will have to determine the manner in which it varies.[37] Unless there are clear theoretical reasons determining how a quantity varies, there is an element of *choice* for the modeler in terms of how measured quantities are understood to vary.

In order to discuss how error is introduced into a model we need to discuss the manner in which algebro-geometric constraints can vary. Strictly speaking, error will be introduced into any mathematical constraint that has an empirical meaning. In the case of error in algebro-geometric constraints, we must consider the variation of not only specification of values and spaces of the objects of the constraint, but also the algebraic form of the constraint. There are a number of different ways that a mathematical constraint could be varied, which relate to different kinds of error. For the purposes of modeling measurement error we only need a very simple kind of variation. This is a variation of a constraint that leaves the function defining the constraint, denoted F as in equation (3.1), fixed and allows the value-specifications of objects related by F to vary.

To see how measurements are treated, let us consider the simple pendulum model. The constraint equation $\mathcal{C}$ of the ideal simple pendulum model, *viz.*,

$$\ddot{\theta} + \frac{g}{\ell}\sin\theta = 0,$$

[36] If $\|\cdot\|$ defines a norm on a space M, then the function

$$d(x, y) = \|x - y\|$$

defines a metric on M.

[37] Generally, we will assume that mathematical quantities are continuously varying, but it is possible to consider the structure of the variation of mathematical quantities to only allow discrete jumps to other specific values.

relates three objects: $\theta(t)$, g and ℓ. Let us suppose that the value for g is exact and that $\theta(t)$ is a definite function, so that g and $\theta(t)$ are fixed exactly. On the other hand suppose that the value of ℓ is to be measured, in which case ℓ, strictly speaking, must be modeled as empirical and consequently has a tolerance ε.

Now, consider a measurement situation. Suppose that the measured value of ℓ is ℓ_μ and the measurement error is ε_ℓ. Then, we set the identified value $\overline{\ell}$ to be ℓ_μ and the specified tolerance $\overline{\varepsilon}$ to be ε_ℓ. Then, within the mathematical model, any length $\tilde{\ell}$ in the range $|\ell_\mu - \tilde{\ell}| \lesssim \varepsilon_\ell$ is a valid value for the length ℓ. Measuring the value of the length implies that $\ell = \ell_\mu$, which has the effect of adding the constraint

$$\ell - \ell_\mu = 0$$

to the model constraint, forming the constraint system $(\ddot{\theta} + (g/\ell)\sin\theta = 0, \ell - \ell_\mu = 0)$, which we can present in a reduced form as $(\ddot{\theta} + (g/\ell_\mu)\sin\theta = 0)$. But since the actual value of the length need not be equal to the measured value, and if we are dealing with real numbers it (almost) never will be, the actual length ℓ_α will be different from ℓ_μ by some amount δ_ℓ, *viz.*, $\ell_\alpha - \ell_\mu = \delta_\ell$. Since we have now allowed ℓ from the model constraint to vary, we need to consider the effect of the variation on the model constraint.

Since making the length an empirical quantity allows the constraint equation of the model to vary, the constraint equation also becomes empirical. Thus, the error in the empirical quantity is translated to the constraint. Before considering the effect of error in ℓ, let us denote the vector field of the model constraint by $f(\theta, \ell) = -\frac{g}{\ell}\sin\theta$, where we are now regarding ℓ as variable. Then the model constraint $\mathcal{C}$ takes the form

$$\ddot{\theta} - f(\theta, \ell) = 0.$$

We then begin by substituting $\ell_\mu + \delta_\ell$ for ℓ_α in the model constraint equation. After some algebraic manipulation we find that

$$\ddot{\theta} - f(\theta, \ell_\mu) = -\frac{\delta_\ell}{\ell_\mu + \delta_\ell} f(\theta, \ell_\mu).$$

Thus, neglecting terms second order in the relative error $\delta_\ell/\ell\mu$, the defect for the

modified constraint is

$$\delta_{\mathcal{C}} = -\frac{\delta_\ell}{\ell_\mu} f(\theta, \ell_\mu),$$

which will be small compared to $f(\theta, \ell_\mu)$ if the relative error δ_ℓ/ℓ_μ is small, which it will be for precise measurements.

We see from this that the size of $\delta_{\mathcal{C}}$ compared to the size of $f(\theta, \ell_\mu)$ is thus just the relative error δ_ℓ/ℓ_μ. This is what we want to know, *viz.*, that the measurement error produces a small relative error in the dynamics of the model. This means that small relative measurement error will not affect the dynamics of the model in a significant way. In this case we say that the model is *well-conditioned* with respect to measurement error, meaning that as long as the model is *descriptively stable*, *i.e.*, it accurately describes the dynamics of the real pendulum, then inferences about the dynamics drawn on the basis of the model will also apply to the dynamics of the real pendulum. Since we are assuming that the model for the exact length ℓ_α is an exact description, we conclude that the model is well-conditioned with respect to measurement error. This does not come as a surprise, since we would not expect small error in the measurement of the length to have a significant effect on the model. But this is not true in general, and we therefore see that an advantage of this approach to analyzing measurement error by considering the constraint to be variable is that we can *prove* whether or not the model is stable under small errors in measurement.

There is one additional issue for measurement error that arises in the NEO modeling case. Since the model problem requires finding values of the six orbital parameters in order to fix the orbit of the object, the initial conditions and any condition computed from them are empirical. In the context of NEO modeling they do not simply consider tolerances, *i.e.*, measurement error, for these quantities but also a probability distribution over the possible values of the quantities. In terms of the empirical quantities we are using, a gaussian probability distribution is defined over the range of valid values of each quantity with the identified value as the mean. Mathematically, then, the vector $\mathbf{k}$ of orbital parameters is treated as a random vector and the variance is given by the covariance matrix of the joint probability distribution of the six orbital parameters. We may see that this case is *compatible* with the definition of an empirical $\mathcal{A}$-$\mathcal{G}$ vector quantity, since it makes sense to add a probability distribution over valid values and there is no conflict because the definition does not require any

distribution at all. In fact, Stetter (2004) intentionally does not include a probability distribution in his definition of valid value of an empirical quantity. Part of the reason is that leaving out probability makes for a simpler theory. But he also argues that judgements of validity in practice are typically not precise enough to justify the choice of a probability distribution. Consequently Stetter opts for an alternative approach which models variability in the data in terms of variable intervals.

For our purposes here, this debate is significant for its implications for methodology avoidance strategies in epistemological modeling. The significance of this debate about how to model the variability in data is then that it provides an example of a *functionally defined object* that can be implemented in a number of different ways. If we wanted to add a schematic way of talking about the variability of a quantity into our epistemological model, then we could require that this be treated in terms of some formal treatment of the *distribution measure* of valid values around the identified value. This condition on the distribution measure is the functionally defined object, it has to function in such a way as to measure the distribution of valid values around the identified value. Stetter employs a system of variable intervals and a "validity value" to do this in the context of numerical algebra. The difference between this approach and the standard probabilistic modeling of error, then, is that a single functionally defined quantity, the *distribution measure*, can be implemented in different ways in different contexts. Consequently, the general concept of a distribution measure can be implemented in certain mathematical modeling contexts as a probability distribution and associated variance, and in numerical algebra in terms of variable intervals. This is an example of *modularity* in an $\mathcal{A}$-$\mathcal{G}$ methodology avoidance strategy associated with variable implementation of a functionally defined type, which we will see more of in the following chapter.

3.3.2 Reduction Error and Its Structural Variation

This second variety of error arises from the use of mathematical techniques of model reduction, which includes techniques of physics avoidance. In this case we are primarily concerned with techniques that are applied to mathematical constraints in order to obtain an analytically and computationally more tractable constraint and/or to focus in on the dominant behaviour of a natural phenomenon. This can involve tech-

niques such as linearization, which is applied to a function quantity in a constraint in order make it a linear function, thereby facilitates solution or computation. It also includes techniques of dimensional analysis, including nondimensionalization, which either reduces the number of degrees of freedom in the solution that need to be considered or removes unnecessary data-constraints. In terms of the $\mathcal{A}$-$\mathcal{G}$ perspective we are considering, these different techniques can be distinguished in terms of their effect on both the algebraic form of constraints and on the nature of their effect on the geometric interpretation of constraints.

To see how reduction methods affect model constraints, consider once again the model of the ideal simple pendulum, *viz.*,

$$F(\theta, \omega) = \ddot{\theta} + \omega^2 \sin\theta = 0. \tag{3.7}$$

Linearizing this model constraint changes the form of the vector field, yielding the constraint

$$\tilde{F}(\theta, \omega) = \ddot{\theta} + \omega\theta = 0, \tag{3.8}$$

which replaces a nonlinear function object with a linear one. Thus, in this case, rather than changing the value-specification of the objects of a constraint, as occurs for measurement error, we have actually changed the function defining the constraint. We can still evaluate the solution in terms of a defect, however, in the following way.

Suppose that we obtained a solution of the form $\tilde{\theta}(t) = A\cos(\omega t + \phi)$ to the linearized model. This solution will then satisfy the linearized simple pendulum model (3.8) exactly. But we can substitute this solution for θ in the original constraint (3.7), which will not satisfy (3.7) exactly but will leave a *defect* $\delta(t)$, *viz.*,

$$\ddot{\tilde{\theta}} + \omega^2 \sin\tilde{\theta} = \delta(t). \tag{3.9}$$

The solution $\tilde{\theta}(t)$ can thus be regarded as a variant of the (unknown) exact solution $\theta(t)$ to (3.7). And regarding a solution as variable will be reasonable if the original constraint is empirical and has a defined tolerance ε. In this case, the solution $\tilde{\theta}(t)$ to the linearized model (3.8) is an effectively valid solution to the nonlinear model

(3.7) provided that

$$\|\delta(t)\| \lesssim \varepsilon.$$

If this were a real model, ω would not be fixed (since the length ℓ of the pendulum rod is measured) and the value of the tolerance ε would be determined by the scale of other errors, *e.g.*, measurement error, in the model. So, the idea is that if the error in the model constraint due to modeling error is less than or of the order of the other errors in the model then the approximate solution is an effectively valid solution to the model, at least locally to where the modeling approximation is effectively valid.

This type of model reduction leaves fixed the objects to which the constraint function is applied, but allows the function F defining the constraint to vary. An independent way that an an algebraic constraint can be varied, is to keep the function F fixed, but to change the number of objects. In some cases, steps in the model construction process lead to an effective increase in the number of objects, such as when the vector field of the simple pendulum system is specified, which adds the parameters g and ℓ in addition to the function θ on configuration space. This replaces a generic symbol $\mathbf{f}$ for a vector field with the symbol complex $-\frac{g}{\ell}\sin\theta$. Consequently, one symbol is replaced by three, *viz.*, g, ℓ and θ.[38] If a model reduction method increases the number of objects, it will generally also reduce complexity so that more objects need to be specified, but they are much simpler objects to deal with.

Other steps in the model construction process, however, can reduce the number of objects, an example of which is a modeling strategy that reduces a finite-volume rigid body to a point. Another example of this is nondimensionalization, when applied to a constraint that removes the need to specify certain parameters. Thus, we can see that the algebraic effect of model reduction methods on model constraints is to reduce the number of symbols in a constraint, increase the number but reduce the complexity, change the function defining the constraint to a simpler one, remove whole constraints, or some combination of these.

This clarifies the nature of the algebraic variation that model reduction techniques effect. Let us now consider the covarying algebraic and geometric structure of mod-

[38] Arithmetic operations and functions like $\sin x$ can be regarded as *library functions*, and consequently do not need to be treated as symbols. This is the same way that the trajectories of the perturbing bodies in the equation of motion for the NEO orbit model would be treated, given that they are evaluated using ephemerides from JPL.

els that such techniques effect and how the nature of the error thereby introduced can be analyzed in an $\mathcal{A}$-$\mathcal{G}$ framework. We have already seen that the linearization technique, algebraically, replaces a nonlinear function with a linear one, and, geometrically, replaces a curved structure with a flat one. In case that it is applied to a nonlinear ordinary differential equation (vector field), it replaces a nonlinear function in a model constraint with a linear one and a curved vector field on the system space with a flat one. In case that it is applied to a nonlinear measurement reduction function, it replaces a nonlinear operator with a linear one (a matrix) and a function locally curved under variation in control space with one locally flat. In whatever context the linearization method is applied, it has a similar algebraic effect and similar geometric consequences, yielding a more easily handled algebraic constraint and a locally accurate geometric interpretation.[39]

Another technique that we saw in a very simple application was the nondimensionalization of the hamiltonian for the double pendulum model. In this case, algebraically, this removes one data-type parameter symbol and eliminates the need for an algebraic constraint to specify it (specifying the length ℓ of the pendulum rods) and, geometrically, this just removes a characteristic length scale for the vector field. As was mentioned above, however, the main importance of nondimensionalization in applied mathematics, however, is that it is capable of reducing the dimension of a problem, *viz.*, the number of independent quantities that need to be solved for. In this case, the algebraic effect of the technique is to reduce the number of solution-type symbols appearing in model constraints, often simplifying the algebraic constraint in the process, and the geometric effect is to reduce the dimension of the space within which the solution objects are fixed to be in or defined on, usually making it easier to construct or compute a solution. Thus, whereas linearization preserves the dimension of a model control or state space, nondimensionalization generally reduces it. Such a reduction will lead to valid, or effectively valid, solutions provided either that the behaviour of all of the originally specified variables can be recovered, a case of effective invertibility, or the behaviour of the solution in the reduced variables is

[39]The case of the treatment of the orbit determination problem as a linear least squares problem, linearization is effectively being applied at different points along the orbit in order to solve the problem easily, and effectively allowing a linear approximation to not only be valid locally but across the entire data arc.

sufficiently close to the behaviour we would obtain if we knew the exact behaviour of the original variables subject to the original constraints. The geometric idea here is that the projection of the exact solution into the reduced model spaces would yield a small error.

A last case that we will consider here, that we did not consider explicitly, is the calculation of post-newtonian corrections to the newtonian equations of motion. This involves a kind of reduction method since it, *inter alia*, involves reducing the tensor equations of general relativity to effective vector forces in a much simpler newtonian model. This sort of method involves a shift between different theoretical frameworks, *viz.*, entirely different kinds of geometric (or topological) spaces interpreting algebraic constraints. Of course, the algebraic form of the constraints changes drastically as well. This example serves to illustrate that there are reduction methods that are beyond what can be captured by the methodological framework we have developed here. A similar kind of $\mathcal{A}$-$\mathcal{G}$ approach could be applied to include this sort of case, however, by analyzing the algebraic and geometric form of tensor equations and characterizing the algebraic and geometric nature of the constraints and reduction methods used in the calculation of post-newtonian corrections.

This subsection suffices to illustrate how different methods of model reduction can be described on an $\mathcal{A}$-$\mathcal{G}$ perspective, both in terms of their effect on the algebraic form of constraints and on their geometric interpretation. This kind of approach can give us insight into the nature of these kinds of methods, enabling us to understand *why* they are effective for the purposes they are used for, which we have seen illustrated in the simple cases of linearization and nondimensionalization. One of the advantages of this sort of approach is that it could help to render less opaque the qualitative and quantitative nature of modeling error. This kind of approach will be of particular importance for bringing technical precision to explanations of the effectiveness of techniques that are capable of bringing the dominant behaviour out of a model, yet are applicable as algebraic manipulation techniques that do not require clear understanding of their geometric consequences.

3.3.3 Abstraction Error and Its Structural Variation

A third kind of error that arises in model construction is concerned with how well a theoretical framework, modeling framework or model is capable of capturing the (dominant) behaviour of the phenomenon. This kind of error is associated with the nature of the abstraction from the phenomena involved in the selection of theoretical and model constraints. In some cases these are chosen from a particular theoretical framework, one thought to be applicable to the given modeling context. This is done in all of the cases we have considered. In other cases these are chosen by drawing modeling assumptions from our knowledge of the phenomenon being modeled, which includes rigid body and symmetry assumptions. In yet other cases this involves the selection of "library" modeling constraints that have already have been shown to reduce from more complex theoretical frameworks in similar modeling contexts. In each of these three cases of constraint selection, the error involved would be analyzed in a different way.

In the last case, selection of library modeling constraints, which includes the choice of post-newtonian corrections to the newtonian equations of motion, the most immediate kind of error reduces to the reduction error considered in the previous subsection, since such reductions were used by whoever performed the initial construction of the library constraints. Whichever theoretical constraints that these library constraints were originally constructed from are subject to another sort of error, the same kind of error associated with the choice of constraints from a particular theoretical framework. Both this theoretical framework error and the error associated with modeling constraints obtained from assumptions made based on knowledge of the phenomenon, are associated with how well the constraints are able to represent the actual behaviour of phenomena. This raises difficult and sometimes deep issues concerning the relationship between mathematics and the world, since in general the given phenomena do not admit distinct description independently of the theoretical framework we wish to assess in relation to the given phenomena. Some issues concerning the relationship between mathematical frameworks and the phenomena themselves will be considered throughout the remainder of the study; for now, since the three kinds of abstraction error can be seen to reduce to the abstraction error associated with theoretical and modeling constraints, let us consider the different forms of abstraction involved in

these two cases.

For theoretical constraints, provided that they are not capable of a complete, direct representation of every possible aspect of the phenomenon, they necessarily involve some form of abstraction from the actual behaviour of the phenomenon. A simple example of this is classical electrodynamics, which fails to be representative at low photon density and high energy transfer. But this is not specific to "classical" theories, since general relativity and quantum field theory also have such limitations. Provided that the physical behaviour not captured or screened out by the theoretical constraints is not relevant to the dominant behaviour of interest, then the theoretical constraints are likely to admit additional constraint to yield an effectively valid but more easily tractable model when modeling such dominant behaviour. For modeling constraints, which in the feasible case are never capable of complete, direct representation of every aspect of the phenomenon, they also necessarily (from the point of view of feasibility) involve some form of abstraction from actual behaviour. Two basic kinds of abstraction involve what part of all phenomena of the universe one treats as the target system, *viz.*, what part of all phenomena one abstracts out from the rest of the universe, and for that part of the universe one does model, what features do the model assumptions abstract out and what do they screen out. There is also the further question of whether the features abstracted out behave according to the constraints of the theoretical framework chosen. With the nature of the kinds of abstraction error outlined, let us now consider a couple of examples.

Post-newtonian celestial mechanics provides a good example of the kind of error associated with selection of constraints from a theoretical framework. As is well known, newtonian mechanics is only able to adequately describe the behaviour of objects if they are moving slowly relative to the speed of light and moving in weak gravitational fields. But we have seen that even for the motion of bodies in the solar system, which are moving slowly in a weak gravitational field, relativistic corrections are necessary for high precision. So, the use of newtonian mechanics for orbit determination will lead to a distortive effect (error) and predictions of behaviour of objects may not be accurate over long time scales, even within the Lyapunov time. When the consequences of mistakes are the potential destruction of the planet, it is crucial to avoid such distortions. This could be remedied by using general relativity instead,

but this introduces formidable analytic and computational challenges. The strategy used in celestial mechanics is to use general relativity to calculate high precision corrections to newtonian mechanics, thereby extending the useful life of an old, familiar and computationally efficient theoretical framework.

The way that we may think of this in terms of error is that the equations of newtonian mechanics are only (effectively) valid under quite limited conditions and scales of time, energy, space and momentum (also of force or acceleration). In more extreme conditions, the behaviour predicted by newtonian mechanics diverges from the actual behaviour of physical phenomena. This does not mean that the theoretical framework of newtonian mechanics does not contain models that can accurately describe the actual behaviour of physical phenomena, it just means that the models that it naturally generates on its own terms fail to be accurate in more extreme conditions—the theoretical framework becomes descriptively unstable under such conditions. From a feasible point of view, it is prudent to extend the life of old theories if they are analytically or computationally effective, familiar and easier to handle. Even though general relativity is (effectively) valid over a much wider range of physical conditions and scales than newtonian mechanics is, it is much less familiar mathematically, much more complicated to handle both analytically and computationally, and requires the development of new analytic and computational techniques in order to apply readily in practice. This is why applied mathematicians use post-newtonian corrections, and sometimes (second order) post-post-newtonian corrections, in order to continue to apply the theoretical framework of newtonian mechanics to phenomena that newtonian mechanics on its own fails to describe.

The pendulum models we have considered are interesting from a model construction point of view because they employ strong macroscopic constraints that are known to be good approximations to actual macroscopic conditions in order to obtain an enormously simplified model *before choosing a theoretical framework*. The validity of these constraints depends on their capturing the effective behaviour of a phenomenon at a macroscopic scale. Because many materials at macroscopic scales behave like perfectly rigid bodies to a very good approximation, and do so under a variety of conditions, macroscopic rigid body constraints can be very useful for macroscopic modeling. The effective validity of such constraints relies on their being *effectively*

true of the bodies they describe, in the sense that to within reasonable tolerances the behaviour of the actual body is the same as the behaviour characterized by the chosen constraint. For instance, provided that the bending and stretching of a pendulum rod is negligible (relative to the dimensions of the rod) given the forces exerted on the rod under the phase space conditions we are concerned with, then the abstract constraint $\dot{\ell} = 0$ will be effectively true of the rod. This is epistemologically significant because it shows that even though models employ abstractions, there is a distinct sense in which the abstractions are true of the phenomena. Logical approaches view such abstractions as "idealizations" that are false of the phenomena, which makes the use of abstractions in explanation problematic for the DN model which requires the "laws" to be true. Thus, the concepts of local constraint logic are seen to provide a means around this problem by providing distinct conditions under which abstractions are true of phenomena.

This notion of effective truth of an abstraction allows us to see that, similarly to error for theoretical frameworks, error for empirical constraints on behaviour, such as constitutive equations, is treated in terms of whether the behaviour picked out by the constraint matches the behaviour of the phenomenon being modeled, given the scale of physical conditions of the model. In both these cases, the error is negligible provided that the behaviour picked out by the (system of) constraints effectively *covaries* with the behaviour of the actual phenomenon.

Chapter 4

Data Handling in the Real World: Feasibly Describing Phenomena

4.1 A Feasible Approach to Data and Phenomena

4.1.1 Algebro-Geometric Data and Phenomena Constraints

In the last chapter we introduced the general framework of *algebro-geometric constraint systems*. This was used to represent mathematical models and the process of their construction in terms of constraints on possible solutions to a mathematical problem. This chapter concerns the role that data plays in feasible scientific inference. This requires methods of relating the constraints of a mathematical model to constraints imposed on it by available data. As is now well-known, and shown initially by Suppes (1962), in applying a theory, theoretical descriptions are not compared directly with raw observational or experimental data. Thus, the raw observational or experimental data is structured into a form—a data model—that can be compatible with theoretical descriptions. The structured data in such a data model can then be interpreted within the framework of the theory. Data provides an important source of constraints on mathematical model problems by *translating* the data from a data model into the framework of the mathematical model. Or, thought of the other way round, a mathematical model that requires data to find a solution can be "implemented" in a data model, which allows data in that model to be used to constrain the mathematical model problem. Thus, the idea of a *data constraint* is that data provide an important source of constraint on mathematical models, and they do so by translating data from a data model into the framework of a mathematical model.

We will see in this chapter that local constraint logic can provide useful insight into data handling in astronomy by thinking of data in terms of constraints in this way, insight not provided by the usual picture of data models.

We saw an example of using data to constrain a mathematical problem in the previous chapter, *viz.*, the use of measurements of the celestial location of a NEO over a long period of time as constraints that overdetermine a orbit model of the NEO. When we interpret this data as carrying information about the position of a NEO in space, which requires the solution of the orbit determination problem, we obtain information about the phase space position of an object at a particular time. We saw that this phase space position is what is used as the initial condition for the accurate orbit model needed to solve the impact probability estimation problem. Thus, in this case the required constraint on the model, *i.e.*, the initial condition, is not a data constraint *per se*, since it does not come directly from data, but a constraint from knowledge about properties of the phenomenon. Thus, we could consider such constraints to be *phenomena constraints*. In a similar way, such constraints must be "translated" into the framework of the mathematical model. Given that the model of the phenomenon lies within the framework of the mathematical model this might not appear to involve any kind of translation. As we will see in this chapter, however, there is actually quite an elaborate translation process involved.

We may see from this brief consideration of the NEO modeling case that the concern of this chapter is the feasible conversion of data into constraints on a mathematical model. We will see that this process relies on a complex variety of mathematical models in a way that is not accounted for in typical presentations of the "hierarchy of models" picture of the theory-world relation. We will see from the detailed consideration of how data is used to feasibly draw inferences about phenomena in astronomy that the complex variety of models has a quite distinct structure that is not hierarchical. And, interestingly, we will find that the structure of the system of models reflects the methods used to feasibly gain information about NEO from data. In the last section of this chapter I will show how the concepts from local constraint logic clarifies what the stability of such inferences from data relies upon. In particular, I will show that the reliability can be accounted for in terms of a notion of structural stability of transformations between mathematical models that is derived from local

constraint logic.

4.1.2 Local Effective Soundness and Completeness

The concepts of local constraint logic that we considered in the last chapter were restricted to feasible inference within a single constraint framework. For a number of reasons we need to consider translation between constraint frameworks. One case of this that we will consider in this chapter arises in the process of converting a trajectory for a NEO in a theoretical model into the predicted observational data in a data model. Another is involved in the process of conversion of raw image data into observational data in a data model. Another class of cases entirely, which we will consider in the next chapter, arises in the feasible computation of solutions to nonlinear ODE. We will handle such translations in terms of the notion of an *implementation* of a constraint problem in another constraint framework.

We consider a translation between two constraint frameworks where one is identified as the *source framework* $\mathcal{F}_S$ and the other the *target framework* $\mathcal{F}_T$. In the context of modeling in applied mathematics, the source framework will usually be the framework of the model of the target system, and the target system could be a computational model, a data model, or even a representation of the target system itself as distinct from the model of it. For expository purposes here, we will use the notation $x = s$ for a solution to the constraint $\mathcal{C}(x)$ in the source framework $\mathcal{F}_S$. An *implementation* of $\mathcal{C}(x)$ in the target framework $\mathcal{F}_T$ is a constraint system $\tilde{\mathcal{C}}(x)$ in $\mathcal{F}_T$ together with a partial relation R between $\mathcal{F}_T$ and $\mathcal{F}_S$ that allows expressions in $\mathcal{F}_S$ to be interpreted as expressions in $\mathcal{F}_T$ and, in general, *vice versa*. Given an expression E in one framework, the expressions to which it is R-related are the *valid interpretations* of E in the other framework.

As will become more clear throughout the remainder of this study, we require, in general, that translation is possible in both directions because of how implementations are used to feasibly extract information from a model in the source framework that cannot be feasibly extracted directly. The translation from source to target is needed to *specify* a problem where information from the target framework is used as a constraint on a model in the source framework, and the translation from target to source is needed to *solve* the constraint problem set up by the implementation,

thereby allowing information to be extracted from the model in the source framework. Explicit examples of this pattern will be seen in the manner in which data is handled in astronomy (chapter 4) and how numerical computing is handled (chapter 5).

Since the relation R can be many-many, for a given expression E, a *choice of interpretation* is selection of an expression R-related to E. This then allows for a kind of truth definition for a solution in the source framework. A solution $x = s$ to $\mathcal{C}(x)$ in $\mathcal{F}_S$ *obtains* under an implementation $(\tilde{\mathcal{C}}(x), R)$ in $\mathcal{F}_T$ provided that some choice of interpretation of $x = s$ is a solution to $\tilde{\mathcal{C}}(x)$.[1] A solution $x = s$ is then *valid* under an implementation provided that it obtains for all choices of interpretation. In this case, we may write

$$\Vdash_{\tilde{\mathcal{C}}(x)} x = s,$$

where '$x = s$' is here understood to be any of the possible R-interpretations of $x = s$.[2] In analogy to the interpretation of a logical *language* in a $\mathfrak{model}$, the double horizontal stroke indicates a transformation to an *external semantics*, in the sense of interpretation of expressions in the original framework into expressions in a different constraint framework. This translation of a contraint framework to an external semantics is what is effected by an implementation.

An implementation is *sound* if for any set Γ of conditions on $\mathcal{C}(x)$ and valid solution $x = s$,

$$\Gamma \Vdash_{\mathcal{C}(x)} x = s \quad \Rightarrow \quad \Gamma \Vdash_{\tilde{\mathcal{C}}(x)} x = s.$$

And an implementation is *complete* if for any set Γ of conditions on $\mathcal{C}(x)$ and valid solution $x = s$,

$$\Gamma \Vdash_{\tilde{\mathcal{C}}(x)} x = s \quad \Rightarrow \quad \Gamma \Vdash_{\mathcal{C}(x)} x = s.$$

Then an implementation is defined to be *descriptively stable* if it is both sound and complete.

These definitions are, strictly speaking, not error-tolerant. Thus, we require versions of these concepts that are stable when small errors are introduced. The error-

[1] The reason that I choose the term 'obtain' over 'true' is that I wish to reserve the term 'true' for empirically interpreted constraints that are known to be (effectively) true with (effective) certainty, such as data constraints.

[2] Note that there are concerns here about how to deal with the equivalence class of choices of interpretation that obtain. Possible ways of handling this are to choose a canonical representative or to quotient out the equivalence.

tolerant versions of each of these concepts are obtained simply by adding the qualifier 'effective' in front of the terms 'obtains', 'valid', 'sound', 'complete' and 'descriptively stable', with the implied change in meaning. Thus, an implementation is *effectively descriptively stable* if for any set Γ of conditions on $\mathcal{C}(x)$ and effectively valid solution $x = s$,

$$\Gamma \mathrel{|\!|\!\models^{\varepsilon}_{\mathcal{C}(x)}} x = s \quad \Leftrightarrow \quad \Gamma \mathrel{|\!|\!\models^{\tilde{\varepsilon}}_{\tilde{\mathcal{C}}(x)}} x = s,$$

where $\tilde{\varepsilon}$ is the error tolerance in the target framework. In non-mathematical cases, *i.e.*, cases where machine computation or application to data or phenomena is involved, it is this kind of descriptive stability that matters.

These definitions can be considered to be defining *global* notions of obtaining, validity, soundness, completeness and descriptive stability, since for each of these case, the concept applies for all sets of conditions and all valid solutions. It is possible to define *local* notions of all of these concepts by placing restrictions on the solutions $x = s$ of $\mathcal{C}(x)$ or conditions Γ for which the concept applies. We will be interested in two ways in which the concepts 'obtaining', 'validity', 'soundness', 'completeness' and 'descriptive stability' can be local. In one case restrictions are placed on the solutions in the space $\mathcal{S}$ of solutions of $\mathcal{C}(x)$ to which the concept applies; this involves restriction to some subspace S of $\mathcal{S}$. The other case is where the solution $x(p)$ depends on parameters p and restrictions are placed on the values of the parameters to which the concept applies; this involves restriction to some subset P of the parameter space $\mathcal{P}$. In the general case, the concepts are local in both these senses. Thus, we will define an implementation to be *locally descriptively stable* if

$$\Gamma \mathrel{|\!|\!\models^{x\in S, p\in P}_{\mathcal{C}(x)}} x(p) = s \quad \Leftrightarrow \quad \Gamma \mathrel{|\!|\!\models^{x\in S, p\in P}_{\tilde{\mathcal{C}}(x)}} x(p) = s,$$

for any conditions Γ consistent with the solution and parameter space restrictions.[3]

An important issue that is raised when considering approximate solutions to problems has to do with the stability properties of the problem itself. Since the typical case is that of a constraint problem where we cannot compute the exact solution,

[3]Note that the case where this condition only holds for some set of conditions Γ is included as a special case of local descriptive stability so defined. This is so because the conditions are applied to the solution space or parameter space, so where descriptive stability fails because it fails for certain conditions Γ, this just indicates that a boundary to S or P has been reached.

only an approximate one, the backward error, or defect, δ is directly computable or estimable whereas the forward error, or error, ϵ is not. This is because the defect δ can be calculated by substituting the approximate solution in the constraint and calculating the amount by which the constraint fails to be satisfied, but the error ϵ cannot be calculated without knowing the exact solution. Thus, it is useful to know when a small defect implies a small error, *i.e.*, when for a constraint problem $\mathcal{C}(x)$ and effective solution $x = s$,

$$\left|\left|\left|\frac{\|\delta\| \lesssim \varepsilon_b}{\mathcal{C}(x)}\right|\right|\right| x = s \quad \Rightarrow \quad \left|\left|\left|\frac{\|\epsilon\| \lesssim \varepsilon_f}{\mathcal{C}(x)}\right|\right|\right| x = s.$$

Problems that have this property are called *well-conditioned*, which means that a small deviation from the constraint produces only a small change in the (exact) solution. When this condition does not hold, *i.e.*, when

$$\left|\left|\left|\frac{\|\delta\| \lesssim \varepsilon_b}{\mathcal{C}(x)}\right|\right|\right| x = s \quad \Rightarrow \quad \left|\left|\left|\frac{\|\epsilon\| \not\lesssim \varepsilon_f}{\mathcal{C}(x)}\right|\right|\right| x = s,$$

the problem is called *ill-conditioned*. This means that a small change of the constraint leads to a large change in the solution.[4] The ratio of $\|\epsilon\|$ to $\|\delta\|$ is bounded by the *condition number* κ, so that

$$\|\epsilon\| \lesssim \kappa \|\delta\|.$$

Thus, a problem is well-conditioned if the condition number satisfies $\kappa \lesssim 1$.[5]

Recall that in the previous chapter we showed how the concept of local effective validity from local constraint logic provides a strong analogy to the concept of deductive validity from mathematical logic. We may now consider how this extends to the concepts of soundness and completeness from mathematical logic. In order to draw the analogy we will again let $\mathscr{L}$ be a first-order logical language and $\mathscr{K}$ an associated logical calculus (forming $\mathscr{L}(\mathscr{K})$), but we will now assume that the language $\mathscr{L}$ is fully specified, in the sense that it is given an intended interpretation.

[4]Note that we can define forward and backward ill-conditioned problems and that conditioning can be defined in terms of effective soundness and completeness by considering the source framework to be that of the data space and the target framework to be that of the solution space.

[5]Recall that the symbol '$\lesssim$' indicates a bound that is not a strict bound. This symbol can be read as "less than or the order of", indicating that the quantity on the left is either strictly less than or approximately the same size as the quantity on the right.

This allows us to consider translation between interpretations in terms of transformations between constraint frameworks. In this case then, the source framework is that of the logical language with its intended interpretation and the target framework is another framework with the same language but an interpreteation in some other set-theoretic structure of a logical type compatible with that of the language. The implementation of the logical language in the other set-theoretic structures is strongly analogous to the definition of an *interpretation*, in the logical sense, of the language in a set-theoretic semantics.

A choice of interpretation for the implementation is then simply the specification of some new particular $\mathfrak{model}$ of the language. We then see that a proposition p "obtains" under the implementation if it is true in some $\mathfrak{model}$. And, thus, a proposition p is *valid*, in the sense of implementation in local constraint logic, if it is true in every $\mathfrak{model}$, which is formally analogous to the model-theoretic sense of $\mathfrak{validity}$. Thus, in this case we may write

$$\Vdash_{\tilde{\mathscr{L}}(\mathscr{K})} p,$$

where $\tilde{\mathscr{L}}$ is the interpretation of the language. We, therefore, see that an implementation of a language is *sound* in local constraint logic if for any set Γ of sentences of $\mathscr{L}$ and sentence p

$$\Gamma \Vdash_{\mathscr{L}(\mathscr{K}))} p \quad \Rightarrow \quad \Gamma \Vdash_{\tilde{\mathscr{L}}(\mathscr{K})} p,$$

which is seen to be formally analogous to the concept of soundness in mathematical logic. And an implementation is *complete* in local constraint logic if for any set Γ of sentences of $\mathscr{L}$ and sentence p,

$$\Gamma \Vdash_{\tilde{\mathscr{L}}(\mathscr{K})} p \quad \Rightarrow \quad \Gamma \Vdash_{\mathscr{L}(\mathscr{K}))}, .$$

which is seen to be formally analogous to the concept of completeness in mathematical logic. Therefore, the important requirement of a logic that is both sound and complete is seen to be formally analogous to the concept of a globally exact *descriptively stable* implementation of one constraint framework in another. In a manner like we considered in the previous chapter, if we consider the "local language logic" treatment, the the concepts of local effective soundess and completeness are again generalizations of the corresponding logical notions and reduce to the respective logical notions in the

same kind of limit ($\varepsilon \to 0, \partial \to 0$).

Thus, we see that not only is the concept of validity internal to a framework formally analogous to the syntactic concept of validity from logic, as was shown in the last chapter, the concept of validity *external* to a constraint framework is also formally analogous to the *semantic* concept of validity from logic, along with the corresponding concepts of soundness and completeness.

4.2 Structural Features of Feasible Data Handling

Contemporary models of the theory-data or theory-world relation in philosophy of science represent the relationship in terms of a hierarchy of models of different types. Generally speaking, this hierarchy is broken down into a "top half", representing the process of generating hypotheses, generalizations and predictions, and a "bottom half", representing the process of generating data models used to test these hypotheses and predictions (see figure 4.1). The breakdown of the structure of this hierarchy provided by Teller (2010) divides the top half, *i.e.*, the "theory side", into principled models constructed from theory that are used to construct models that represent the target phenomenon, which then yield "relatively specific" hypotheses and predictions. The bottom half, *i.e.*, the "data/phenomena side", is divided into models of experimental setups which are used to construct models of data that can be compared with the hypotheses or predictions from theory, with the world supplying the raw data to be interpreted in the data model. This model is understood to be a general (not universal) pattern of the role that models play in the sciences at a quite high level of abstraction from detail.

Although this is clearly an enormous advance from the ideal classical model of the theory-world relation in terms of logical deductions from theory that are interpreted in and compared with raw observational data, this model still preserves the structure of hypotheses and predictions from theory on one side and data obtained from observation and experiment on the other. In real theory application, however, theory and data are not independent in this way and can be entangled in quite complex ways. Giere (2010*a*), whose models of the theory-data relation Teller based his account on, acknowledges that there are cases where representational models are data driven in

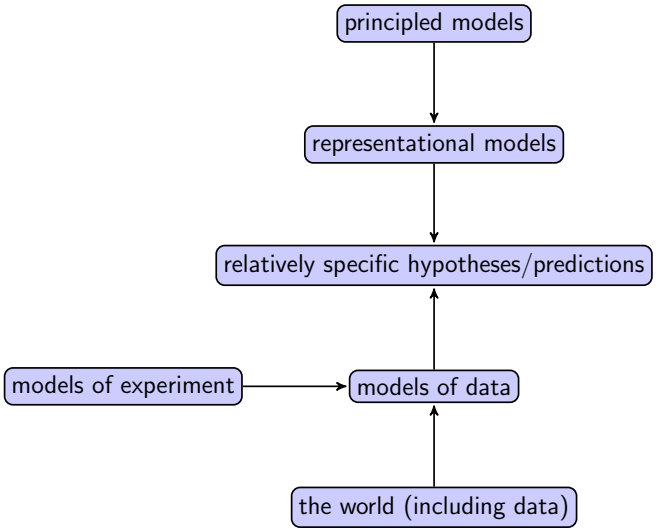

Figure 4.1: A contemporary representation of a hierarchy of models in the theory-world relation (adapted from (Teller, 2010)).

cases where there are no "higher principles" or "theory". In such cases, representational models are constructed from the data along with other empirical models and a variety of mathematical techniques (*op. cit*, 272). Giere takes phenomenological theories and some simulation models to be examples of this kind of case. The discussion from the previous chapter, however, shows that any models constructed from theory that employ macroscopic constraints (which is arguably the vast majority of cases in applied mathematics) also rely on data models for their construction. Thus, even hypotheses and predictions from theory can rely on data models.

The example of impact probability calculations for NEO shows that the entanglements between theory and data run much deeper than this. As we have seen, in order to compute impact probabilities for a NEO, it is necessary to first compute a high-precision orbit on the basis of a long arc of astrometric measurements. The result of this computation is a single point in phase space, the effective position and velocity of the object at a particular point in time. It is only when this point has been computed that a reliable computation of the impact probability can be performed. Moreover, a NEO that is believed to have a significant risk of impact will continue to be monitored so that the error in the computation of the impact probability can be reduced and the length of time that the orbit can be known with high accuracy can be extended. Thus, the relevant prediction in the case of NEO modeling, *viz.*, the time of close approach and the probability of impact, requires a complex *process*

of interaction between theory and data.

In this chapter we will consider this process of interaction in some detail. I will develop a model of the theory-world relation that both reflects this process of interaction and is, I will argue, a more accurate representation of the relationship between the theory and the phenomenon in cases that extend far beyond observational astrophysics. This is a model that represents the theory-phenomenon relation as a recursive process of implementation and reinterpretation involving a complex network of models. This network, naturally, will be seen to resemble the standard sort of hierarchy of models in a number of respects, but I will argue that the process-model that I present reveals epistemologically significant structure that is not captured by the standard view.

The guiding principle underlying the organization of the model of the theory-phenomenon relation I will present is access to information. Since any access to information that can be used in real scientific inference must be feasible, we will be interested in the methods that make access to needed information feasible within the operating constraints of any given case of scientific inference. In the case of NEO modeling, the information we require about a given object is its probability of eventual collision with Earth at some time in the future. To determine this, we must first know the orbital trajectory of the object with high accuracy and known uncertainty. This will allow us to determine impact probabilities. In order to know this however, we must determine a way of using accessible information about the object to compute the orbit. This is accomplished by using a large collection of astrometric observations of the object to compute the position and velocity of the object at a single point in time to high accuracy and with known uncertainty. As we will see, in order to accomplish this we need to accurately model how information is transmitted from the object to the telescope, how the telescope processes the information, and how to reconstruct this information in a data model.

Thus, to make an accurate inference about impact probability, we need to construct a model of the phenomenon, *i.e.*, the orbit of the NEO, and then we need to access a variety of kinds of data and construct/select the models required to tell us with high precision how information transmitted by the object gives rise to the data. Once the required models have been constructed, the inference of the impact proba-

bility for a given object requires an iterative process between the theoretical model, data and measurement. Thus, the example of NEO modeling is one case where inference not only relies on data, but also where a complex interactive process between theory, data and measurement is *required* to make the inference feasible. And since we need to know now whether or not an object will impact the Earth, not simply that it is possible in principle to have this information, it is feasible inference that matters.

4.2.1 Handling Data Structuring

We started with the question of how to determine the probability of a NEO impacting the Earth at some time in the near future. In considering this process, we considered how it begins by constructing a post-Newtonian modeling framework for handling the Sun-NEO orbit subject to perturbations from the Moon, planets and large solar system asteroids. We discussed in the previous chapter how an algorithm is constructed to compute an orbit for a NEO given a large data arc of observed location data and a model, *viz.*, the measurement reduction function, for how the real position of the object in space can be reduced to observed location data. Only once such an orbit can be computed can the orbit be used to construct algorithms to find the probability of a near-Earth pass or an impact of the object with Earth. In this procedure, then, the input from data comes in the form of ephemerides for the Earth, Moon, planets and asteroids that are included in the equation of motion of the NEO, and then the observational data together with the measurement uncertainties.

This might make the NEO impact probability determination case look like the usual pattern of data-in-predictions-out modeling, reminiscent of the hypothetical-deductive (HD) method of confirmation, but this is not so. The nonlinear post-Newtonian model is not actually the starting point for real cases where a NEO is sighted for the first time. When a NEO is sighted for the first time, an observer will take a series of astrometric measurements for as long as possible, but will then seek out the aid of other observatories or amateur astronomers at other places in the world to continue observing the object so that it is not lost. The modeling process begins when a rough initial guess of the orbit is calculated using one of a variety of two-body methods, including the classical methods due to Gauss and Laplace (Chodas &

Yeomans, 1999). It is only when several days of observational data are available that the nonlinear model with perturbations is used to compute a more accurate orbit. Such an orbit is typically still not usable to compute a reliable estimate of impact probability, however, since it may not even be accurate to predict the position of the object on the next opposition.

Data arcs as long as a few months can be accurate enough to predict the position of the object a decade or more into the future, allowing a useful impact probability to be computed. The accuracy of an orbit determination, however, increases significantly when the object can be observed for three or more oppositions. With a data arc of this length, and provided that there have been many accurate individual observations, high-precision orbits can be determined ten years or more into the future and are eligible for receiving an official catalogue number by the International Astronomical Union. There is, however, a sense in which this is possible even with only a week to a month of direct observational data. In such a case, an object can often be linked to a previously observed object that was subsequently lost (Chodas & Yeomans, 1999). And there are even cases where objects observed for a month or more are used to predict previous oppositions in order to locate a previously undetected object in photographic archives.

A dramatic example of this was the case asteroid 1997 XF11, greater than 1 km in size. In March 1998, on the basis of an 88-day data arc, the asteroid was predicted to miss the Earth within a quarter of the lunar distance on October 28, 2028. On March 11, Brian Marsden, director of the Minor Planet Center, announced the prediction, noting that a collision was possible but unlikely, but that the passage within one lunar distance was "virtually certain". The very next day, however, two JPL researchers, Ken Lawrence and Eleanor Helin, found four prediscovery images, immediately increasing the data arc to 8 years, reducing the uncertainty by an order of magnitude and determining that the predicted close approach increased to over 2.5 lunar distances. Chodas & Yeomans (1999) note, however, that a more careful analysis of the 88-day data arc using the methods they describe is able to rule out the possibility of a collision.[6] Thus, the accuracy of a method of orbit determination depends significantly not only on the length of the data arc but also on how much

[6] A collision was only possible at around 55-sigma by their calculations.

thought and care has gone into the details of the algorithm.

Clearly, there is a great deal of back and forth between computation and data handling (modeling, gathering and processing) in order to determine a precise orbit for a NEO, sometimes requiring subtasks of analysis of old data in order to extend the length of the data arc. Thus, even though a given computation of an orbit may bear a certain resemblance to the HD schema, such a representation of prediction is far too simplistic to represent the fact that the data responsible for determining the hypothetical assumptions of the prediction themselves rely on a complex process of modeling and computation, let alone represent the fact that the model used in the "derivation" of the prediction relies on non-deductive approximation methods for its construction and on numerical computation for its solution. Thus, we see that the HD method, along with its twin, the deductive nomological (DN) model of explanation, represent quite significant classical distortions of feasible reasoning. Again, these classical models are not supposed to accurately represent scientific practice, but if the kinds of "derivations" they assume are not feasible over any time period using techniques of logical deduction, it is difficult to see what insight they are supposed to give into scientific methodology or epistemology. The problem I am pointing to is the models' characterization of "derivation" as logical deduction. As a methodology avoidance strategy it makes good sense to represent prediction schematically, but any faithful representation of practice must represent a "derivation" as a much more complex and stable—not certain—process of inference.

To ensure that the data used in the computation of the orbit of the NEO leads to a high-accuracy orbit, three things are required:

1. A high-accuracy model of the observations as a function of the position of the object in phase space;
2. High-precision ephemerides for significant solar system bodies along with high-precision observations of the object; and
3. Assurance that the devices used to collect the observations are accurate.

We consider the latter two cases in the following two subsections. In this subsection we consider how the model of the observations is computed. In order to understand this, we need to consider a more general problem upon which all accurate measurements in astronomy depend: how to accurately model the transmission of information from

an object to the telescope.

Reduction of Observational Data

The structuring of astronomical data collected on the surface of the Earth requires a variety of interrelated models, including models of the shape and density of the Earth; models of heliocentric motion of the Earth; a variety of models of the forces that are exerted by different solar system bodies on the Earth and that lead to the precession of Earth's axis of rotation; models of the effects of motion on emitted and detected light; and models of the effects of the atmosphere on light. This network of interrelated models is fundamental to precision astronomy; the combination of these models that allows high-accuracy modeling of the transmission of light from a celestial object to the objective lens or primary mirror of a telescope. In this way data modeling in astronomy requires a complex network of data models.

Ordinarily, this network of models would be seen as part of the complex system of models used to interpret and test hypotheses and predictions from theory, right in the middle of the bottom half of the hierarchy of models in the theory-world relation, *i.e.*, the half pertaining to structuring data into a data model (see figure 4.1 above). And in a quite clear sense this is precisely the case, since the function of these models is to accurately structure certain observational data on Earth and dynamical phenomena in the solar system so that precise calculations can be performed. These models allow for the computation of significant perturbations of the motion of light through space and of the coordinate systems used to specify that motion. So from this point of view this network of models fits naturally into a standard picture of a hierarchy of models. From the point of view of actual scientific method in astronomy, however, this network of models is entangled with the process of construction of hypotheses and predictions from principles models of theories. To see this we need to consider the structuring of observational data in astronomy in more detail.

The reason that this complex network of models is required in astronomy is that interpreting predictions about the observed location of a solar system body, star, galaxy, *etc.*, requires calculation of the location in the sky at a given place on the surface of the Earth that the object will appear at a particular time given the hypothesized or predicted position of the object in space (in some coordinate system).

The process of converting data in one form to data in another is called *reduction* in astronomy. Thus, the process we are considering here is the reduction of a predicted position in space to a location on the celestial sphere observed from a given location on the surface of the Earth.

The actual model or models used in this calculation depends on the required precision and the purposes for which the calculation is being performed. Many of these models, such as coordinate transformations and analytic models of the shape of the Earth, are specified by rigorous, exact formulae. The rigorous formulae are not always required and are usually replaced by series approximations or numerical calculations if the full precision of the analytical model is not required for the precision requirements of a particular application. In many cases, however, an analytic model is itself the result of a simplification process, and higher precision may require that a given model be computed to a higher level of precision before being incorporated into reduction calculations.

Which models of this complex network are used and how depends on the precision demands of the given case. The precision that can be had in any given case is constrained by the source of error in measurement that has the largest effect on the model. This relates to the conditioning of a mathematical model with respect to a source of error (recall the consideration of error in measurement of ℓ for the simple pendulum in the previous chapter), since a small error when the model is ill-conditioned in relation to that source can lead to a large error in the model. In any case, the modeling context will demand that the error in the data from a given source be smaller than some value to ensure that the solutions of the mathematical model obtained from the data are valid. Let us call this value ν_0. This value will determine what level of approximation of the models required in a reduction will, when combined, result in a reduction transformation that has an error smaller than ν_0. Typically this allows for a variety of approximations of the model formulae that simplify calculation but are still effectively valid to ν_0.

Now let us suppose that we have a specified precision and let us consider the case of the reduction of positions of an object in the solar system.[7] We often need to do this because we are using celestial mechanics to make certain predictions and we want

[7]The procedure for objects outside of the solar system is similar.

to determine how to connect the theoretical predictions with data. In this case, a standard problem is to compute what is called the *topocentric place*, which is the point on the celestial sphere at which an object would be seen from a point on the surface of the Earth using a perfect optical instrument.[8] And in any real case, the connection between the position in space and the topocentric place must be calculated.

For our purposes, we do not need to consider all of the corrections that go into this computation. It suffices to observe that the reduction calculation breaks down into a sequence of subreductions. In the case we are considering, *viz.*, celestial mechanics calculations in the solar system, the "raw predictions" of the theory are the position and velocity of an object at a particular point in (dynamical) time. Moreover, they are referred to an inertial heliocentric coordinate system. Since heliocentric coordinate systems are not actually inertial, this means that heliocentric coordinates are referred to the ecliptic and equinox of a particular epoch, typically the standard fundamental epoch J2000. If the topocentric place is to be computed for a particular observation time t, the heliocentric position $\mathbf{x}(t_e) = (\mathbf{r}, \dot{\mathbf{r}})(t_e)$ at the iteratively computed time of emission t_e is reduced by the following sequence of steps:

1. A trivial reduction to the coordinate system of the solar system's centre of mass, resulting in a barycentric position $r'\mathbf{l}(\lambda, \beta)$ in ecliptic coordinates;

2. A reduction from this position to a barycentric position that corrects for the precession of the Earth's axis from J2000 to the observation time t;

3. A reduction from this position to a geocentric position $r''\mathbf{l}(\alpha, \delta)$ in equatorial coordinates, making additional differential corrections, including effects of the motion of the object and motion of the Earth;

4. A reduction from this position to a topocentric position in equatorial or horizon coordinates, making more differential corrections, including the various effects of the atmosphere on the incident light.

If this is the procedure used to model the observations, then the result is a data model that fits into a standard picture of the hierarchy of models, but with a twist

[8] Specifically, where the fundamental plane of the coordinate system on the celestial sphere is the true celestial equator at the time of observation and the origin of longitude is the true equinox at the time of observation.

that, to my knowledge, has not been emphasized. The literature on data models emphasizes the importance of data models in the conversion of raw data to structured data that can be compared with theoretical predictions. In this case, however, we witness an inverse situation where the raw predictions of the theory are converted to structured predictions that can be compared with data (models) from observation or experiment. In this case, then, we see that it is not accurate to say simply that the theory makes predictions independently of a data model against which it is tested. Rather, the theory makes raw predictions that are converted into structured predictions in a data model, which can then be tested when the relevant raw data is interpreted in that same data model. In this way, in astronomy, data models play an important role in *mediating* between raw data and raw predictions from theory. The reason for this is simply stated, *viz.*, there is a gap in space and time between what the theoretical model predicts and what can be experimentally observed. Thus, making the connection between theoretical predictions and the resulting data requires accurate modeling of how information is transmitted from the object to the measurement device, which requires a network of data models. Let us call this network, then, the *information transmission model*.

At this point, someone could argue that the information transmission model is easily accommodated by the standard picture because it is actually part of the model of experiment in this case. This is so since it plays the role of determining what phenomena the optical observations are actually picking out, *viz.*, the line of sight to an object in space at the time when the light from it was emitted. The information transmission model is then just adding to this information about what the observed line of sight would be given a particular time of observation and location on the surface of the Earth. This would seem to be a perfectly natural way of accommodating this case in a standard hierarchy of models picture. This interpretation is challenged, however, by the role that the information transmission model plays in the problem of computing impact probabilities for NEO.

Recall that the orbit determination problem for a NEO required the specification of a measurement reduction function $R(\mathbf{x})$, which is a model of the kind of measurements that are input into the problem—a data model. It is used to compute a trajectory for the object that minimizes a weighted sum of the squares of

the measurement residuals, which are computed from the observational data input into the problem. Thus, the measurement reduction function is an essential part of the algorithm used to solve the orbit determination problem. This arguably makes the measurement reduction part of the model of the behaviour of the NEO, since its behaviour cannot be predicted without it. And in this respect, we can see how it operates just like data models in experiments do, except that it is acting on the side of the theoretical model to structure the raw predictions of the trajectory model into predictions of observations. But the measurement reduction function is not identical to the information transmission model just discussed, it is only *part* of it.

In order to optimize the computation of orbits for a NEO it is necessary to avoid computing the same information multiple times. This is why, for instance, ephemerides are used for the dynamical trajectories of the solar system objects perturbing the motion of the NEO in the post-newtonian model. For the same reason, it is inefficient to reduce theoretical predictions to each individual location where observations where made, which can be various locations across the globe. Thus, observations from telescopes are reduced to a standard (J2000) system of geocentric equatorial coordinates. This way, a simpler measurement reduction function can be used and only one is needed since reductions to different locations do not need to be made. This is accomplished by running the information transmission model in the opposite direction, which can be done part way even when the distance to the object is not known.[9] This way, the one of the coordinate transformations and many of the differential corrections, including correcting for the precession of the Earth, are accounted for by the model of the measurement data for every observation. This data can then be stored at a central computing facility and used to compute or improve an existing orbit.

As a result of this approach, the measurement reduction function $R(\mathbf{x})$ required for the orbit determination problem only needs to change coordinate systems to account for the change of origin from the centre of the Sun to the centre of the Earth; rotate the system to equatorial coordinates; and make simple differential corrections, *e.g.*, for the

[9] This is done by correcting for the geocentric parallax, *i.e.*, the change in observed location due to the distance between the observer and the centre of the Earth, either by using an estimated distance on the basis of an estimated trajectory or by using observational data from another location or locations (Roy, 2010, 60).

aberration of light transmitted by the object due to its motion relative to the observer. All the other corrections have already been made to the measurement data that is input into the problem. Accordingly, we may now see that in this case one half of the information transmission model is used in the reduction of raw measurement data to a standard data model, and the other half of the information transmission model is used in the reduction of the raw theoretical predictions to the same data model. So, the data meets the theoretical predictions halfway, and *vice versa*. This example shows that the line between where the theoretical model ends and the measurement data model begins is blurred by both the solution of the theoretical model and the measurement data relying on the information transmission model.

If you are not already convinced that this case presents a problem for standard pictures of the hierarchy of models relating theory and data/phenomena, recall that solving the orbit determination problem does not yet result in the inference we are trying to make. What we need to infer is the impact probability for a given object, which requires calculating a high-accuracy orbit for the object. Thus, the actual prediction that results from solving the main problem is whether the object will have any close approaches to Earth and, if so, what its impact probability is. This is the prediction that corresponds to the "hypothesis or prediction" stage of the hierarchy of models. Thus, there is a complex relationship to data that is *part of the inference* of this prediction, a relationship to data that goes far beyond simple measurement of theoretical parameters and requires a complex network of data models.

And even if someone wished to argue that this can be accommodated into a standard model hierarchy, there is still the matter that the relationship to data in the inference of impact probabilities is not static. The impact probability of a NEO is refined over time by monitoring the object and improving its orbit on the basis of an extended data arc. Thus, data is used in the inference of the impact probability in an iterative correction process extended over a potentially long period of time. This then reveals quite explicitly that the feasible inference of impact probabilities for NEO is a fundamentally computational process—there is no way to make a stable inference of the impact probability without this iterative correction process relying on additional observational data. This reveals the manner in which, far from being a prediction made by theory that is independently tested by data, feasible inference can require a

computational process that is deeply entangled with data.

4.2.2 Handling Data Processing

Up to this point we have considered how the data transmitted between an object and the measurement device is modeled, and how this is used to construct the measurement reduction function. Now we turn to consider the matter of the generation of the observational data itself. Any substantive issues concerning the measurement of time are obviated at the level of precision we are concerned with for NEO modeling due to the availability of atomic clocks, which ensures that measured time is effectively dynamical time, *i.e.*, time in post-newtonian mechanics. Any remaining issues to do with treatment of time are effectively changes in coordinate system for time. Thus, we will just focus on the measurement of location on the celestial sphere.

While we continue to think of data in terms of transmission of information, we are now concerned with the transmission of information through the measurement device to generate the raw data, which is then processed into a data model. At this stage we are comfortably within a standard picture of the hierarchy of models at the level of models of data. The feature of this part of the process that is of the greatest epistemological interest in this case is the manner in which measurement error is handled. Rather than being handled purely in terms of some statistical analysis, measurement is handled in terms of an *error model*, which can be solved in a very similar manner to that in which the orbit determination problem is solved, as we shall see.

Before considering this case, let us consider briefly the matter of transmission of information through the measurement device. Since we are considering optical measurements, the measurement devices of interest are optical telescopes. Consequently, the incoming light is processed by a system of refracting lenses and/or reflecting mirrors and is projected onto an image plane where some kind of photographic data is collected. In order to calculate an observed location on the celestial sphere for incident light, it is necessary to know how the telescope processes the incident light before it reaches the image plane. At the level of precision we are concerned with the light dealt with in the measurement process effectively constitutes rays undergoing rectilinear propagation. Thus, geometric optics is used for all of the calculations of

how the telescope processes the light. In order to know how the telescope processes the rays of light, it is therefore necessary to construct a geometric optics model of the lens/mirror system of the telescope.

A fully detailed model of the imaging system would require a quite complex optical model together with a model of the rest of the telescope, including the mounting, drive system, camera and any other objects that are effectively coupled to the telescope. Simply, a fully detailed model of the imaging system would require a complete model of the device. In the interests of physics avoidance, it is desirable to ignore the full detail and focus simply on a simple model of the optical processing. And provided that the telescope has been well designed and is functioning properly, this is all that is required. Thus, a central design aim for a telescope is to create a device that will collect data that can be analyzed using only a simple optical model. We will consider briefly this design phase in the next subsection. Thus, we will work with the assumption that the effect of the device on the light travelling through it is well below the level of precision ν_0 that has been identified for the modeling task.

In this case, we are considering a charge coupled device (CCD), which is used to take an image from the image plane of the telescope. CCD are useful in optical astrometry because they have a high quantum efficiency, *i.e.*, they register a high percentage of incident photons, and they have a fixed geometry that is easily interfaced with computers. After an image of the sky is taken, it is necessary to reduce this raw image data to a location on the celestial sphere. Ideally, we could align the optical axis of the telescope with a known point on the celestial sphere and a known pixel on the CCD, which was made to lie precisely perpendicular to the optical axis in the image plane. These conditions are not feasible, however, so it is necessary to have a reliable and very precise means of correcting for error in positioning and in the imaging process. This is accomplished by the aforementioned ideal model of the geometric optics of the imaging system, together with a model of the relationship between locations on the photographic image and locations on the celestial sphere. In fact, the need to explicitly consider the optical model of the imaging system can be eliminated through the use of numerical methods.

In general, the precise details of how image reduction is done vary depending on the structure of the imaging system of the telescope and the method of image re-

duction, and need not concern us here. What is of interest is the procedure used to establish the effectively correct relationship between locations on the photographic image and locations on the celestial sphere. Effectively the way that this is handled is by modeling the combined error in the imaging system and in the positioning of the camera as an affine transformation (*i.e.*, a linear transformation together with a translation) of a pair of cartesian coordinates of a point on the photographic image. This reduces the task of analyzing the error in the computed celestial location to calculating the required affine transformation needed to correct for the error, followed by the propagation of error from photographic measurements to the computed celestial location.

The affine transformation of the error model is specified by six numbers, so at least six compatible constraints are required to solve the error model to find these six numbers (for details on the following reduction process see (Green, 1985)). In astrometry, this is achieved by ensuring that the image includes at least four known stars that have available ephemerides and that are, ideally, widely separated. With four stars appearing in the image, their location in the image coordinates can be computed (by computer for a CCD image), which gives a system of eight or more constraints, rendering the error model problem overdetermined and a least squares solution to be sought. The more stars that are included in this calculation, the more accurate the solution of the error model. Once the parameters of the error model have been computed, the celestial location of any point on the image can be calculated. In a very similar way to the least squares method in the orbit determination problem, the solution is found by minimizing a sum of squares of measurement residuals. These residuals can then be used to estimate the measurement error for the location of the NEO.

To briefly clarify how this is accomplished, consider that the celestial sphere is curved and the image plane flat. The locations of the known stars obtained from ephemerides can be projected onto a tangent plane to a point on the celestial sphere. The tangent plane and the (ideal) image plane will then be related by a simple linear transformation. The locations of the stars in this tangent plane are used as constraints on the error model, since they can be related to measured locations of the corresponding stars on the photographic image. Using numerical methods, a

vector method can be used to model the required affine transformation for the error model, allowing the location of the NEO in the tangent plane to be computed and the measurement error to be estimated. These two quantities are then converted into the location of the NEO on the celestial sphere and the measurement error there using the inverse of the projection onto the tangent plane. Thus, in this way, what could be an enormously complex process of error analysis and statistics is reduced to the solution of an overdetermined six-parameter model and simple error propagation. And outside of the estimation of the measurement error on the basis of the measurement residuals, no statistical analysis is required.

4.2.3 Handling Data Collection

We will end this section with a brief discussion of the considerations that go into the design of a telescope so that the error analysis procedures described above will produce accurate results. This requires that the telescope be designed in such a way that the device itself does not introduce error in the raw data that goes beyond a specified tolerance for measurement error suited to the purposes for which the telescope is to be used. Thus, a complex error analysis task is involved in the design of a telescope for adequate performance, *viz.*, to have the imaging system and camera effectively independent of the state of the device and any internally or externally produced physical perturbations. And, much like the case of error analysis just considered, a great deal of it also relies on modeling rather than statistics.

In order to ensure that a telescope system is capable of measurements that are accurate to a given precision ν_0, it is necessary to ensure that forces exerted internally or externally to the telescope do not lead to displacements in the image plane of size larger than ν_0. This is a constraint on possible designs of the telescope in much the same way that the hamilton equations are a constraint on possible mathematical models. The basic components of an optical telescope are its optical imaging system, the mounting of the optical system that allows for its orientation with respect to some celestial coordinate system, the electrical drive mechanism that rotates the mounting to match the rate of rotation of the Earth (called a sidereal drive mechanism), and a mounting for an imaging instrument like a CCD camera. Given this simple specification of components and following the discussion in Roth *et al.* (1994), this general

design constraint gives rise to three more specific but simple design constraints:

I. The static constraint: static forces (depending, *e.g.*, on the orientation of the mounting system) shall not cause displacements in the focal plane larger than ν_0;

II. The kinetic constraint: vibrations resulting from any source shall not excite amplitudes of over $\nu_0/2$ in the focal plane and should be attenuated as soon as possible;

III. The kinematic constraint: image drift in the focal plane caused by kinematic inaccuracies (from, *e.g.*, the sidereal drive mechanism) shall not exceed the amount ν_0 for a given observing time t_b.

These three constraints act much like the basic equations of a theory, providing general constraints on the design of an optical telescope that can ensure accuracy to the precision ν_0. This is so because this system of constraints is highly underconstrained. There is a wide range of possibilities of kinds of optical systems (refracting, reflecting, both), kinds of mountings (*e.g.*, horizontal mountings, equatorial mountings), kinds of drives (different drive mechanisms, motors, control circuitry, *etc.*), supporting electrical systems, guiding systems for locating stars, signal processing, *etc.* Any number of combinations of such designs can meet the above three design criteria. Thus, much like the process of model construction detailed in the last chapter, additional constraints must be introduced to construct a design model.

Of course one would select a number of ideal criteria, such as what kind of control of the orientation of the optical system is possible, but these kinds of constraints are flexible. There are other constraints that are less flexible or inflexible, including available finances, available sources of raw materials and components, and accessible facilities for manufacturing and machining along with the technical expertise in these areas. These constraints help to significantly narrow down the possibilities for what design models will be feasible. Aside from these more obvious constraints, the three design constraints above can be used to compute a variety of more detailed constraints on a design model that is able to meet the ν_0-precision design constraint.

To ensure that it is even possible to meet the first two basic design criteria (I and II), it is necessary to model how a given mounting design responds to internally and

externally generated forces. The response of the system to forces must be analyzed in terms of vibrations and displacements of the focal plane. Accordingly, the model of the mounting system is a kind of error model, used to model the measurement error that will result from the measurement device itself. This is another case where an error model does not rely significantly on statistics, but rather on models of elasticity from continuum mechanics. Simple elasticity models allow more detailed design constraints to be computed that guide the solution of the design problem.

It can be shown that the *stiffness* of materials is the dominant concern for meeting design constraints I and II. For a simple linear model of elasticity, the stiffness is simply the constant c of proportionality in the hooke equation

$$\mathbf{F} = c\mathbf{x},$$

where c has units of newtons per metre.[10] There are two reasons for this: stiffness determines the organization of all of the parts so that constraint I is satisfied; and stiffness also plays an important role in determining the response of the telescope to vibration and can be used to control vibration (constraint II) (Roth *et al.*, 1994, 148). In terms of the static constraint, it can be shown using results from the theory of elasticity that n parts connected in series must have n times the required system stiffness and that the most important stiffness constraint is high flexural stiffness, *i.e.*, resistance to bending. In terms of constraint II it can be similarly shown that the design of the mounting should simultaneously meet two conditions: 1) aim for high natural frequencies (which improve vibrational damping properties) and high mechanical impedence (to insulate the system from external vibrations), and 2) design every part with maximal stiffness and minimal mass and avoid unnecessary attachments coupled to the image plane.

In meeting constraint III, statistics becomes a more dominant concern, since meeting this condition has much to do with parts of the system being manufactured, machined and connected to spec, which also requires accurate measurements of the properties of the materials and their arrangement. Consequently, these concerns fall within the domain of concern of the standard theory of measurement and error.

[10] For a more general elasticity model, c is a second rank tensor or matrix.

The interest in considering how error is handled in the cases of image reduction and telescope design is that a great deal of the analysis of error in these cases is handled using error models that are more similar to the mathematical models of the target phenomenon, *i.e.*, the principled models of the "top half" of the model hierarchy, than the models of error used in statistical analysis on the "bottom half". Statistical methods of error analysis are based on mathematical models of randomness and uncertainty, but they are applied as mathematical algorithms for processing raw data into a data model and as algorithms for analyzing the error introduced in direct measurements. Error models for image reduction and for meeting the static and kinetic design constraints for a telescope, on the other hand, are presented as constraint problems for which a solution is to be computed.

In the design case, the combination of principled models of the components of the measurement system and their couplings together with a formulation of the design problem as a constraint problem guides the design process in a way that makes solving the problem of finding a design that meets all of the strict constraints, including the ν_0 level of precision, feasible. Moreover, it provides assurance that this level of precision is *stable*, since it allows the designer, or operator after the fact, to know what assumptions this level of precision relies upon. This then can guide a maintenance schedule to ensure ν_0 precision performance and the troubleshooting process of finding the source of a detected failure to meet the precision requirement. Having a reliable assurance that the precision constraint is met under given operating conditions, the operator can then be assured that the six-parameter error model for image reduction is sufficient for ensuring that the measured value of the celestial position of an observed object is indeed accurate to within around the calculated measurement error. Thus, having a way of knowing whether the measurement system is operating to spec, by knowing how it responds to internal and external physical perturbations, it is possible to ensure that the operator of the device can interface with it using only a very simple error model.

4.3 Data Error and Structural Stability of Inference from Data

The focus of the last section was on how to describe the structure of the roles played by data and data models in the inferential process used to estimate the impact probability of a NEO. This required a clarification of the manner in which both data and models of data are used in the process of solving this problem and how error is handled throughout the process in order to ensure that the eventual estimate of the impact probability is accurate to within a tolerable level of uncertainty. We saw that the part of the process of solving the impact probability problem in which data is playing the most dominant role in driving the inference is in the computation of a precise orbit for the NEO. We also saw that there are an enormous number of factors that must be taken into account in order to ensure that the inference of the orbit is reliable. The focus of this section is to show how the (high level) concepts from local constraint logic can provide clarity and insight into the reliability of the inference from data in this case. The result will be a model of the structure of inference from data in astronomy and astrophysics that is likely to be generalizable to a wider range of cases of inference from data in applied mathematics.

As we saw in the last section, the analysis of the role that data models play in inference from data in astronomy shows that there are important structural features of this inferential process that are not adequately accounted for by typical models of the theory-world relation in terms of hierarchy of models. In the last section I argued that this is the case even for the contemporary informal understanding of such a hierarchy, implying that a more formal conception of it along the lines of how Suppes (1962) envisioned it will also be inadequate. In this section I will be more explicit in my reasons for seeing a formal semantic hierarchy of models as being conceptually inadequate for giving insight into why the data driven methods of inference described in the last section are reliable. I will argue this by showing how the concept of a locally, effectively descriptively stable implementation of a constraint framework can provide conceptual clarity into how inferences from data are reliable. In the final chapter I will argue that this sort of model can also give insight into *why* such methods are reliable. Accordingly, I will argue here and in the final chapter that an explanation of

reliability of inference from data in astronomy can be given in terms of the structural stability of the conceptual means by which data and data models are used to gain knowledge about the behaviour of objects moving through space.

4.3.1 Structural Stability of Data Modeling

As we discussed earlier, pictures of the theory-world relation based upon a hierarchy of models relating theory and data typically represent the relation as a two-sided process: one side moving down from general theory to descriptive models and then to data models, and the other side moving up from the world and raw data to data models, which usually require support from models of experiments in order to convert raw data into a data model (see figure 4.1). I argued in the previous section that because an accurate data model (of how information is transferred from an object to another location in some Earth-related coordinate system) is an integral part of the algorithm for computing the orbit of a NEO, data models are important tools for solving theoretical models, not only for parameter measurement and testing consequences of theoretical models. Moreover, I argued that because computing the orbit of a NEO is a subtask of computing the impact probability of a NEO, the prediction that answers our initial question about the risk posed by a NEO, data and data models are actually entangled in a complex way with determining predictions of a theory, going far beyond simple boundary conditions and initial states. Thus, a picture where, except for the measurement of theoretical parameters, the role of data is only to test the validity of a scientific inference is a distortion of scientific practice (see figure 4.2).

I do not claim that this shows that such hierarchy models are invalid, because they assume a high level of abstraction from scientific practice. It does show, however, that such a picture of the theory-world relation presents a distorted view of the full role that data plays in scientific inference. I see this as another tension between classical epistemology, which focuses on the most general questions of our use of theories in scientific inference, and feasible epistemology, which focuses on an analysis of how scientific inference actually works with an aim to explaining its reliability. Accordingly, the concern with pictures like Teller's and Giere's (shown in figure 4.1) is that because they are so abstracted from scientific practice they present such a

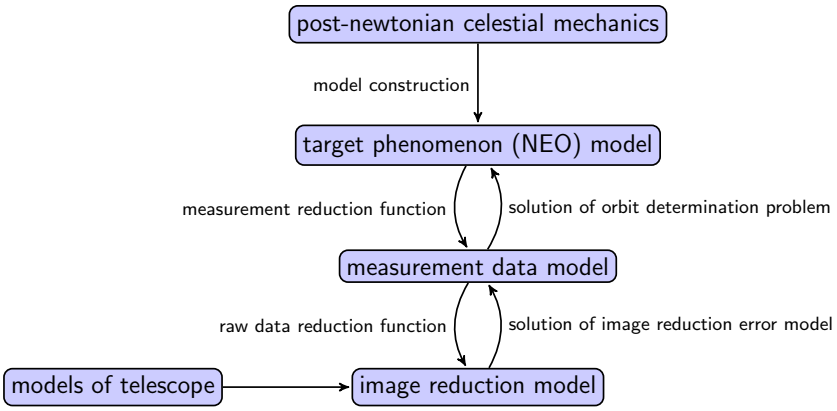

Figure 4.2: A representation of a model of the theory-world relation in the case of near-Earth object modeling using post-newtonian celestial mechanics, showing how data handling is an integrated part of scientific inference that relies on data. In relation to figure 4.1 this figure reveals the distortion of presenting the central inferential task (hypotheses/predictions) as distinct from input from data models.

distorted picture of scientific inference that their conclusions about science become unreliable. I argue that it is important to clarify the structure of scientific inference in real cases to ensure one's representations are valid before proposing or presenting very general abstract models. Although the considerations of data handling in this chapter are too specific to motivate a general pattern in feasible scientific inference, I will develop an argument through the remainder of the study that local constraint logic does reveal a quite robust structure to feasible scientific inference. At the moment I will simply show how local logical concepts can help to clarify the structure of scientific inference that is driven by data in manner that occurs in astronomy.

In the representation of mathematical model construction discussed in the last chapter we actually saw another instance where data models play a role in scientific inference that is not typically emphasized. We saw that structural and behavioural assumptions about a phenomenon are commonly made in the model construction process in the form of symmetry assumptions and constitutive equations or equations of state. These assumptions are now seen to be kinds of data models. Thus, data models play an important role in feasible model construction in applied mathematics. Nevertheless, the result of the model construction process is a system of mathematical constraints on the behaviour of the target system, *i.e.*, those aspects of the target phenomenon being tracked in the modeling. For differential equations models, the common view in philosophy of science is that predictions can be made on the basis

of the model by specifying boundary conditions and an initial state. Since boundary conditions are not required in evolutionary ODE models like those in the NEO modeling case, only the specification of an initial state is required. Once we know an accurate initial state and have an accurate estimate of its error, we can compute the orbit and compute an impact probability for any close approaches. But specifying the initial state is and its error is not possible from a simple measurement; rather, it requires solving another problem, the orbit determination problem, which relies heavily on data models.

So how are we to represent the inferential process used to solve the orbit determination problem? We can do this using a simple representation of the role that data models play in the solution to this problem. As we now know well, the solution to this problem requires the specification of a measurement reduction function $R(\mathbf{x})$, which is an accurate model of the measurement of the celestial location we would obtain (in a specific Earth-related coordinate system) at a given measurement time t. We saw in the last section that this data model is a model of the transmission of information from the object to the location specified in the data model for the measurements. Let us refer to the data model that $R(\mathbf{x})$ relies upon as the *measurement information transmission model* and the data model for the measurements as the *measurement data model*. We described the process of transmission of information in terms of the measurement information transmission model providing a conversion of the raw prediction of the theoretical model into data predictions in the measurement data model. The orbit determination problem is then solved by inputting a long data arc along with the measurement errors into the problem and computing a least squares solution that minimizes some function of the measurement residuals.

Thus, the following inferential process is used to solve the orbit determination problem:

1. construct/select an accurate equation of motion for the object;

2. construct/select an accurate measurement information transmission model for the object;

3. specify the observational data and errors;

4. find a phase space state that determines an orbit that minimizes a weighted sum of squares of the measurement residuals.

Now, since the measurement data model in the NEO modeling case is a geocentric equatorial coordinate system, the measurement data model is a mathematical constraint system. Now, we can regard the measurement information transmission model as an *implementation* of theoretical model framework in the measurement data model framework, *viz.*, it allows states or values in one constraint system to be interpreted in another. In this case, it allows a phase space state to be interpreted as an equatorial location in a geocentric coordinate system.[11] Now, we will consider this implementation to be reliable if it meets the following accuracy condition: given that the NEO is in a state $\mathbf{x}_0$ at time t, the actual measured position in the measurement data model at the time t' that the light from the object reaches the measurement position is $R(\mathbf{x}_0)$ to within around the error of a typical measurement or better. Stated otherwise, the implementation is reliable provided that the measurement predicted by the measurement reduction function would be the effectively true location on the celestial sphere where the object was observed in the coordinate system of the measurement data model, given that the object is in the state $\mathbf{x}_0$ at t. In local logic terms, then, the implementation is reliable provided that it is *effectively sound*.

Just to be entirely clear in this initial case, we will use the logical notation introduced above to express this relationship more precisely. Let $\mathbf{x}_0$ be state of a NEO at some initial time t_0, and suppose that $\mathbf{x}_1$ is the state at some other time t_1. Then, the statement of this in the local logical notation is

$$\mathbf{x}(t_0) = \mathbf{x}_0 \left\|\!\!\left|\frac{d^2\mathbf{r}}{dt^2} = \mathbf{f}_{so} + \sum_{k=1}^{np}(\mathbf{f}_{spn} + \mathbf{f}_{pno}) + \mathbf{f}_j\right.\right. \mathbf{x}(t_1) = \mathbf{x}_1.$$

From now on we will simply use the notation $\mathcal{P}(\mathbf{x})$ to indicate the post-newtonian equation of motion of the model of the target phenomenon. Now, we will consider another constraint framework $\mathcal{M}(\alpha, \delta, t)$, namely that of the measurement data model. Since the measurement reduction function $R(\mathbf{x})$ converts phase space states into measurements, $R(\mathbf{x})$ interprets states of $\mathcal{P}(\mathbf{x})$ as states of $\mathcal{M}(\alpha, \delta, t)$. When we consider

[11] It also allows for the distance to the object to be specified in this coordinate system, but since we are not considering the case where the distance is measurable, which requires radar ranging data, we will not regard distance as being part of the implementation.

an implementation of $\mathcal{P}(\mathbf{x})$ in $\mathcal{M}(\alpha, \delta, t)$, we are considering $\mathcal{M}(\alpha, \delta, t)$ as providing an external semantics for $\mathcal{S}(\mathbf{x})$ for those statements of $\mathcal{P}(\mathbf{x})$ that can be interpreted in $\mathcal{M}(\alpha, \delta, t)$. We will use the notation $\mathcal{M}_{(\mathcal{P}, R(\mathbf{x}))}(\alpha, \delta, t)$ to indicate the implementation of $\mathcal{P}$ in $\mathcal{M}$ via $R(\mathbf{x})$. Then, letting ν denote the error in a typical measurement, the expression that the implementation is effectively sound is written as

$$\mathbf{x}(t_0) = \mathbf{x}_0 \Big\|\!\!\!\overline{}_{\mathcal{P}(\mathbf{x})} \mathbf{x}(t_1) = \mathbf{x}_1 \;\; \Rightarrow \;\; \mathbf{x}(t_0) = \mathbf{x}_0 \Big\|\!\!\!\overline{}^{\nu}_{\mathcal{M}_{(\mathcal{P}, R(\mathbf{x}))}(\alpha, \delta, t)} (\alpha_1, \delta_1, t_1') = R(\mathbf{x}_1),$$

where the primes indicate the times of measurement of the state of the object at the corresponding unprimed time. This statement says that if $\mathbf{x}_1$ is the state at t_1 of the object with initial state $\mathbf{x}_0$, then, under the interpretation $R(\mathbf{x})$, $R(\mathbf{x}_1)$ is an effectively valid observation of the object's state at t_1. Thus, saying the implementation is sound means that the measurement reduction function does what is advertised: it maps phase space points to what would be an effectively valid observation of the object at that phase space point.

Effective completeness in this case means effectively that the measurement reduction function is invertible, *viz.*, that given an observation $(\alpha_1, \delta_1, t_1')$ of an object, then $\mathbf{x}(t_1) = (R^{-1} \circ R)(\mathbf{x}_1(t_1))$. But we know that this cannot be the case, because we would have to already have solved the orbit determination problem to be able to convert an observation into a phase space point at a time. In fact, R is dimension collapsing, since it maps an entire curve of points in phase space to a given measurement. We may see this because observing the right ascension and declination of an object tells us only the direction in which it is to be found, which, accounting for light travel time, will not be a straight line. Thus, the initial implementation of the theoretical model in the measurement data model is effectively sound but not effectively complete, which is to say that the implementation allows translation of phase space locations into celestial locations, but not the other way round. But this is all we require for descriptive stability at the stage of *specifying* the orbit determination problem, because a reliable specification of the problem only requires translation one way (from source to target framework). We do seek an effectively complete implementation, however, since we wish to be able to use observational data to infer the phase space position of an object. Once we solve the orbit determination problem we

are able to do precisely this, since we could then specify an invertible relation between the modeling framework and the data model framework, thereby allowing the translation of celestial locations into phase space locations, or in technical terms effecting an implementation that is not only effectively sound but also effectively complete.

There is another implementation that is involved in this problem. The problem assumes that the measurement data, along with the associated measurement errors, are already available in the measurement data model, which must be the case in order to solve the orbit determination problem. But for the solution of the problem to be reliable, it must be the case that the data that is fed into it is reliable. And ensuring that the measurement data itself is reliable requires that another implementation is effectively sound. This is an implementation of the measurement data model in the image reduction model of the measurement device used to collect the data.

To see how this works, recall that in order to be able to relate the data to the target phenomenon it is necessary to have a model of the transmission of information from the phenomenon to the measurement device, *i.e.*, a telescope. We called this the information transmission model above. And we saw that the measurement reduction function $R(\mathbf{x})$ in the NEO modeling case (measurement information transmission model) is a *proper part* of the information transmission model. In order to connect the measurement data model to the image reduction model of the telescope, we need another transmission model $R'(\alpha, \delta)$ that covers the information transfer from the location of the measurement data model to the input of the measurement device.[12] In general, we then need a model of how the measurement device processes the incoming information, but we saw in the previous section how proper design and functioning of a telescope along with the use of computational methods unburdens the operator of the telescope from having to consider a model of how the measurement device processes information. Thus, the implementation required is a (data model-measurement) *raw data information transmission model* that allows data in the measurement data model to be interpreted in the image reduction model for the telescope. We may call this the *raw data reduction function*, to emphasize the relation to the measurement reduction function (see figure 4.2).

[12]In this case, it is not the information transfer through physical space, it is the information transfer from a point on the celestial sphere in one coordinate system to another, along with the required differential corrections.

If this implementation is effectively sound, then a point on the celestial sphere that would be observed by a perfect device in the coordinate system of the measurement data model would be mapped to an effectively valid measurement of that point on an image plate of the telescope.[13] Provided that the implementation is effectively sound, the image reduction model can be solved to determine what the effective measured value is in the measurement data model.[14] We may now see that in both cases the interpretation of data in the opposite direction of the implementation requires the solution of a least squares problem.

Now, it may seem that this data model-measurement device implementation is not part of the problem we are solving, since it is required in order to generate the data but not to solve the orbit determination problem. Thus, it might seem that all this talk of implementations is just fussing over nothing. Moreover, it would then seem that this whole story fits entirely in a usual hierarchy of models picture. This is so because the data model-measurement device implementation is just part of what is involved in interpreting the raw data in the data model, and the phenomenon-data model implementation is just part of what is involved in connecting the mathematical model to data. Once again, we have the usual picture of theory down to data, and raw data up to the data model. There are reasons why this is still a distortion, as I will argue below. But, the more pertinent reason here is that even if this were correct, as far as it goes, it is not actually a correct representation of the actual orbit determination problem for the purposes of solving the impact probability determination problem.

Recall that solving the impact probability determination problem requires continuous modeling of the NEO for which an orbit has been computed in order to improve the projected orbit and thereby improve the estimate of the impact probability. Thus, the overall solution to the impact probability determination problem is an iterative

[13]Strictly speaking, for this to make sense, a cluster of points must simultaneously have this property, because that is what ensures that the image reduction model can be solved to find the measured value of the point on the celestial sphere.

[14]It is actually a little more complicated than this because, strictly speaking, the solution of the image reduction model gives the point on the celestial sphere at the location of the telescope. Because the geocentric parallax of a NEO is nontrivial, it actually requires, in the most general case, the solution of a problem analogous to the orbit determination problem, where the geocentric parallax is computed using observations of the same object from another geographically distant telescope. This effectively means that two distinct implementations are required in order to reduce data from the image plate to the measurement reduction model. But since this does not contribute much to our discussion, we omit this level of detail.

process between theory and data. As a result, the accurate modeling of the data is actually part of the solution of the problem, *viz.*, the solution of the problem alternates between solving the mathematical problems of orbit and impact probability determination, and the observational problems of orbit tracking, measuring, and reducing image data. Thus, the practical problem of figuring out whether a NEO presents a significant impact risk is a higher order problem involving an iterative correction process that alternates between theory and data.

But there are also other reasons why the theory-phenomenon relation is not adequately represented in terms of a typical hierarchy of models in this case. Recall that the phenomenon-measurement information transmission model, upon which the entire process relies for its stability, depends upon a complex network of interrelated models that can be computed to different levels of accuracy in order to meet a specified degree of precision. In the NEO modeling case, this is the system of models that makes possible accurate calculations of the reduction of a location in phase space in the celestial mechanics model to the topocentric place, along with the system of models of the measurement device used to ensure that it produces accurate data and allows valid reduction of the image data. Thus, the system of models upon which the entire process of solving the problem relies, both on the side of theory and on the side of data, is a complex *network* of data models. It is important to note that there is hierarchical structure if one localizes attention to certain parts of the overall process, as we did when we considered a possible argument that the implementation picture I am providing is entirely consistent with the hierarchy of models picture. There is also hierarchical structure in the consideration of this entire network of models at different levels of precision. Nevertheless, the overall structure is that of a network that must be navigated in order to understand what really underlies the reliability of the methods used to solve problems in astronomy.

Thus, the overall structure of the system of models in the NEO modeling case is the following. Data driven problem solving in astronomy relies, in general, upon a complex network of models of the structure and motion of bodies in the solar system, and models of special and general relativistic motion interpreted as corrections to newtonian models, and more generally the stellar environment of the solar system and the motion of the solar system in the milky way, *etc.* This network of models

provides the ground from which more particular information transmission models are formed and upon which information transmission models rely for their (local effective) validity. It is important to note that this network is not simply some diffuse collection of vaguely specified objects, as it has a very definite structure. The reason I refer to it simply as a network is that its complexity, at least currently, prevents me from describing it in a more precise way.

Now, taking this underlying network as given, we can specify the sort of model of the application of theory that comes out of our considerations of the NEO tracking case. A particular case of theory application used to solve a data driven problem concerning the motion of bodies in the solar system requires a theoretical model of the motion of the body using an effectively valid equation of motion, where effective validity is determined in terms of the precision requirements for the given (data driven) problem. In addition to this, one or more information transmission models accurate to within the precision requirements of the problem is required to relate the motion of the body in the theoretical model to observations and measurements used to constrain the problem. When measurements from a number of geographically separated measurement devices are used, it may be required to reduce the various measurements to a single standard measurement data model, the data from which can then be easily used to constrain methods used to gain detailed information about the target phenomenon and its behaviour. Generally speaking, computational methods will be used, so enough measurements are required to overdetermine any mathematical problem used to gain information about specific features of the phenomenon or its behaviour. Thus, the overall picture is that of a theoretical model of the phenomenon, models of the measurement devices and their data reduction, and information transmission models that connect these models; the measurement data model that contains the data used as a constraint on the theoretical problem stands properly between an overall phenomenon-measurement information transmission model, from the phenomenon to the place where the raw data is collected.

Let us now assume that the measurement data model stands between two distinct parts of an overall phenomenon-measurement information transmission model, as we have in the NEO modeling case. Solving mathematical problems that allow us to gain knowledge of properties of the target phenomenon requires an implementation

of the theoretical model in the measurement data model and reliable data reduced from measurement data. In case that the solution of such a problem requires iterative correction using data, an implementation of the measurement data model in the data reduction models of the measurement devices is required.

Let us now assume that we are considering a problem that requires iterative correction using data, as we have in the NEO modeling case. The overall solution of the problem requires the ability to compute inverses to the implementations that allow raw data to be interpreted in the measurement data model and solutions to the data-constrained problems to be computed, thereby allowing the measurement data to be interpreted in terms of its implications for the behaviour of the target phenomenon. Thus, we may see that, provided the implementations are effectively sound, we are assured that features of the theoretical model accurately determine the structured data that would be observed concerning them, and that such structured data accurately determines valid raw data in measurement devices. We may see, then, that the reliable *specification* of data driven problems, requires effectively sound implementations of these sorts. The reliable *solution* of data driven problems, however, requires that we can compute accurate inverses for these implementations. Only then can we accurately reduce raw data to the measurement model in order to provide accurate constraints on data driven model problems; only then can we accurately find the information about the target phenomenon we seek from the theoretical model. Therefore, the ability to solve a data driven problem requires the ability to (effectively) prove that an implementation known to be effectively sound can also be made to be effectively complete.

Since the ability to make reliable data driven *inferences* requires the ability to not only accurately specify problems but also to accurately solve them, we can see that reliable inference in astronomy requires methods that are *descriptively stable* in the technical sense of allowing feasible construction of effectively sound and effectively complete implementations of the model of the target phenomenon in the (raw or structured) data that constrain the model sufficiently to admit an accurate solution.

It should now be clear that this picture of the application of theory to data in terms of implementations of theory in data via information transmission models, themselves supported by a complex network of data models, does not fit naturally

into a picture of theory application in terms of the sort of hierarchy of models that Suppes envisioned and is typically assumed today. There are other, more specific reasons that we have not yet considered that explain why the kind of hierarchy that Suppes envisioned does not apply here. First of all, the sequence theory→principled model→hypothesis/prediction→data model→models of experiment and raw data, with embedding relations holding between adjacent elements, only obtains in a reasonable sense if we restrict attention to a single data driven mathematical problem, like a single instance of the orbit determination problem, and we ignore the details of how the date drives the hypothesis/prediction and the network of models supporting the problem. In this case we have the usual story: theory allows construction of a principled model, the hypothesis/prediction is deduced based on certain data, and the raw data and model of experiment are used to obtain data in the data model that is used to test the hypothesis/prediction. This story gives us no clear insight into what makes this process reliable, however, which the local logical model of feasible inference does. But if we do not ignore the network of data models, we do not ignore the details of inference from data in real science, and we do not restrict attention to a single data driven mathematical problem, the Suppesian hierarchical model captures less and less of what is really going on in scientific practice and of what explains the reliability of the methods used.

Another problem the Suppesian model has in accounting for this case is that it has no provisions for giving proper account of the role played by the network of data models. Since it is this network that allows the construction of implementations that accurately specify the data driven problems and is therefore crucial for explaining why the methods of solution are reliable, this network ought to occupy an important role in an account of the theory-world relation. The Suppesian model does have a response here, however, because the entire problem could be recast in terms of a detailed model of the solar system developed on the basis of post-newtonian celestial mechanics as well as analytical mechanics, continuum mechanics, *etc.*, for the purposes of more detailed models of particular bodies. And given that the models generated by these theories can be specified to different levels of accuracy, they give rise to hierarchies of models in the usual sort of way, with the less accurate models being embeddable in the more accurate ones. In this case the data models and the particular models of

the measurement devices are much more specific, and we can see them as related to the other models in much the sort of way that Suppes envisioned. Therefore, from a larger scale point of view, the system of models does indeed have a kind of hierarchical structure like Suppes' model and the contemporary versions of it.

This kind of situation is actually to be expected in the case of feasible epistemology in relation to classical epistemology. Classical epistemology has been concerned with very large scale patterns in scientific inference, typically seeking models of scientific methods that were maximally general. Since classical epistemology has been successfully describing structure of scientific method at these scales for at least a century, we would expect to recover similar structural patterns to those seen in classical epistemology when we "zoom out" of a detailed description of feasible method. At the same time, however, it is showing that it will likely be a typical case that classical epistemology will not be able to adequately represent the details of feasible methodology and inference. I argue that the local logical model of inference from data just outlined above also provides evidence that classical epistemological approaches will also not be able to adequately explain the reliability of the methods used in scientific practice.

Explaining the reliability of the methods of inference used in science requires explaining why they continue to work even though the actual inferential process in science is riddled with all sorts of error. Explaining the error tolerance of inferential processes in science requires a detailed analysis of how inference in science is actually rendered feasible, which is not something that can be known *a priori*. It requires very detailed engagement with scientific method in order to understand what is driving and explaining the feasibility and reliability of inferences in real application of theory in science, as this study has already shown. Thus, the concepts and methods of classical epistemology are not sufficiently fine grained to be able to detect this level of detail. The tools of set theory and model theory allow for a different kind of fine detail, but this is a kind of detail that is tuned for precise but very general conceptual specification.

The real world of inference is much more complex and messy than classical mathematical methods would let on. Consequently, real world inference demands different and more flexible conceptual tools in order to detect its structure and to account

for its reliability. The conceptual tools of the local constraint logic that I provide here are just an example of the kinds of tools that are suited to elucidating feasible epistemology. For instance, modeling inference in the cases discussed by Wilson (2006) that mix tools and methods from different systems of mechanics will demand a quite different modeling approach. The conceptual tools required cannot be known *a priori* and must be constructed through the process of a particular investigation, which is precisely what I have done in the process of this study. The expectation is, however, that robust structural patterns in scientific methodology and epistemology will begin to be detected by methods of this sort. Indeed, the next chapter will provide some evidence for this by showing that inference driven by computation can be understood similarly to inference from data, *viz.*, in terms of the descriptive stability of implementations of constraint frameworks.

Before we move on it is necessary to point out an additional level of detail that we have been avoiding for simplicity. I claimed above that we can account for the reliability of feasible inference in astronomy in terms of the effective soundness and effective completeness of implementations of modeling frameworks in data frameworks. Effective soundness and completeness, however, only obtain *locally* to certain ranges of conditions. For example, measurement reduction functions generally rely on assumptions that are only valid over a limited period of time. Accordingly, an implementation based upon a given measurement reduction function will not be effectively sound for all time periods in which it could be used. And on a more localized scale of time, because measurement reduction functions rely on models of the motion of the bodies of the solar system, which are nonlinear, the implementation will only be effectively sound for a limited period of time.

Implementations could fail to be effectively sound for other reasons, however. For instance, if the theoretical model ceases to accurately model the trajectory of a NEO, say over time periods longer than the Lyapunov time, the measurement reduction function will not be effectively sound, even if it is accurate, because the model is telling the measurement reduction function that the object is somewhere where it is not. Or, a systematic error in the measurements could cause the projected orbit to be inaccurate in a time period shorter than the Lyapunov time.

Thus, in all of the cases we have considered, the implementations that feasible

implementations rely upon are only effectively sound and complete locally to a particular time range and range of conditions. This, of course, will be the general case, since the process of feasible model construction inevitably results in approximations that will result in a model of limited applicability. And since a tolerance for small error is also required for feasible inference in the typical case, *effective* soundness and completeness will be typical as well. Generally, it is only *in principle* reasoning in a given theoretical framework that will be global and exact. Thus, by adopting the feasible perspective, as we are in this study, we will also adopt a convention where truth, validity, soundness and completeness *simpliciter* will generally mean their effective and local versions, though it will sometimes be useful to use the modifiers "effective" and "local" for emphasis. In case we are considering ideal versions of these concepts, we will explicitly use the modifiers "exact" and "global".

4.3.2 Structural Stability of Computing the Data Model

As was emphasized in the previous chapter, mathematical descriptions are generally infinitely more precise than is meaningful empirically. In the last chapter we considered the manner in which this can be conceptually misleading and how it can cause problems for handling mathematical descriptions of empirical quantities. This exactness of mathematical descriptions also provides important advantages, however, when considered from the point of view of feasible inference.

If we are in possession of a formula we know is rigorously accurate to all scales then we are guaranteed that values calculated using the formula will be rigorously accurate. Often, however, such rigorously accurate formulae are difficult or time consuming to calculate with. As a result, if we are able to work with a given tolerance for error, then we can use approximation methods in order to find a less complicated formula to compute with that is sufficiently accurate for our needs. The advantage of this is that we obtain a result that is just as valid as the exact value for our given purposes, but we do so with considerably less work. This is precisely the approach used for data reduction models in astronomy. If one knows that extremely high precision is not required in a given application, say because ground-based optical telescopes have fundamental limitations in accuracy due to atmospheric effects, then it is advantageous to work with an approximate data reduction model that is accurate

to within the fundamental limitation in accuracy you have to work with (*e.g.*, ν_0 mentioned above). This way, the calculations are made easier and quicker without any meaningful loss in accuracy.

The shift in thinking here over a classical view of computation, then, is that being in possession of a rigorous formula that is exact but working in a situation with a given error tolerance, there is a wide variety of effectively valid approximations to the formula that will be easier and faster to compute with. And thinking of modeling more generally, working in a situation with a given error tolerance allows for a variety of approximative techniques to be used in the modeling that will provide solutions that are just as valid as the solution to the exact model in the given context. Therefore, we see that allowing a tolerance for error in the computation of values of a function, the computation of solutions to a model, or the making of an inference reduces the *computational complexity* of an evaluation, solution or inference. Thus, we may see that an important part of feasible inference and computation is in optimizing the complexity of an inference or computation given the tolerance for error in a given case.

This is precisely what is involved in the approximation methods used in astronomy. The book on computational methods in spherical astronomy by Taff (1991) has as a central focus the provision of formulae for data reduction that are known to be accurate to within certain tolerances so as to optimize the complexity of the computations required in accurate data reduction. The advantage of the understanding of this point of view is not only that one can simplify computations, it is also that one knows how to adjust the reduction methods if a greater degree of accuracy is required for a certain application. An example of this would be the shift from optical astrometry to radar astrometry and one wished to make full use of the boost in accuracy afforded by the use of radar measurements. To make proper use of the higher accuracy measurements, one would have to use much a much more accurate information transmission model (higher precision measurement reduction function and raw data reduction function).

Now, similar considerations apply to the size of a domain over which an approximation is valid. For instance, one might be in possession of a model that is known to be rigorously accurate for some domain but one might only require the model to

be accurate over some much smaller portion of that domain. An example of this is the case of the nonlinear simple pendulum model. Suppose that we were guaranteed that this model was accurate for all initial angles, but we were only interested in small oscillations. In this case, we would be just as well to use the linearized simple pendulum model, which is easier to solve and compute with, since its results will be valid over the more limited domain that we are concerned with.

Thus, we see that an important advantage of classically, *i.e.*, globally exact, valid descriptions is that they can support locally effectively valid descriptions that are just as valid in the given context but are much easier to compute or draw inferences from. Such a classical description has a precision that is far greater than what is required for effectively stable inference in a given case, so depending on the tolerance for imprecision and the size of the domain over which stable inference is required, we can use the classical description to construct a wide variety of descriptions that are locally effectively valid within the local domain or beyond and to the effective tolerance or better. This, in quite simple terms, explains the importance of classical descriptions for feasible inference and computation.

This suggests that there is a general connection between degrees of feasible inference and both the size of the local domain and the size of the effective tolerance. Specifically, the feasibility of inference will be directly proportional to the size of the effective tolerance, since the greater the tolerance the more leeway there is in the simplifications that can go into the methods used. And the feasibility of inference will be inversely proportional to the size of the local domain over which the results are required to be accurate since, for the larger domain, there are stronger constraints on the methods used. This, then, provides a picture of classical inference as an absolute extreme in opposition to feasibility since it requires global validity over domains and zero tolerance for imprecision. At the same time, however, it suggests that feasible inference has local optimality conditions in any given case because an inference will become less and less useful the smaller the domain over which it is valid and less and less reliably accurate the larger the tolerance for error.

4.3.3 Structural Stability of Measurement Systems

We saw from our considerations of the process of designing a telescope that is capable of measurements that are accurate to within around a given tolerance for error that the most important issue in designing a telescope for a certain inferential task is that the physical perturbations generated internally and externally to the telescope do not give rise to a motion in the image plane greater than the given error tolerance. Although we did not discuss the issue in much detail, we also saw the importance of designing the optical system so that the error in the optical system could be compensated by an affine transformation. This requires that any nonlinear errors in the imaging system are negligible. When a telescope that has been designed to meet these constraints is functioning properly, we saw that by using computational methods the raw image data can be reduced to a measurement data model directly from the affine error model without explicitly considering the geometric optics model of the optical system. The importance of these kinds of considerations for feasible inference are obvious.

What is important about this for the device from a structural stability point of view is that the device is effectively insulated from the effects of likely physical perturbations that are generated by stresses or vibrations generated internally to the telescope or from external sources. This can be understood in terms of the concept of conditioning discussed in the introduction to this chapter. Thinking of the situation in terms of displacements or vibrations due to physical perturbations and their effects in the image plane, the design requirements are ensuring that the device is well-conditioned under internal or external physical perturbations. In fact it is stronger than this, because the device is required to dampen physical perturbations where possible. In this sense, the device is designed to be *superstable* in the sense that errors are not only not amplified, they are reduced. Thus, we may see that a central aspect of the optimal design of a telescope, and measurement devices more broadly, is that they are well-conditioned, and superstable where possible, under physical perturbations of raw data collected from the device.

Thus, the telescope design case provides an example where stability concepts apply outside of the framework of a mathematical constraint system. But, as we saw clearly above, it also provides an example of a non-mathematical constraint

satisfaction problem that is solved through the use of detailed "data models" of the components of the device and how they are related to one another. We saw how this allows the formulation of more specific constraints on design, which together with constraints specific to the purposes for which the telescope is needed and facts about the designer, can be used to "solve" the constraint problem by finding a way of meeting all of the imposed constraints, importantly including the constraints on device mediated error in the image plane. We may see, then, that the example shows that a constraint approach facilitates the solution of design problems in addition to the solution of mathematical problems.

4.3.4 Modularity in Data Handling

One final point to make before moving on to the next chapter is that the search for an understanding of general features of the way that data is handled across applied mathematics will require a modular understanding of the structure of feasible data handling. What I mean here is that because the kinds of data that are dealt with in different cases, as well as the methods for handling data, are so diverse across applied mathematics, a structurally stable but fairly general model of data handling will have to be modular. This is to say that certain structural features must be capable of being treated as independent of others. A useful way of seeing one manner in which this could work can be seen by applying the implementation model of inference from data considered earlier in this section.

Consider the case of the simple pendulum system. In this case, there is no appreciable gap in space and time between the behaviour of the pendulum and the information that is processed by measurement devices that measure angles, length and time because these are measured by direct observations of the pendulum itself. Thus, in this case there is no need for an "information transmission model" and so the "theoretical model" and the "measurement data model" collapse into one another. This shows that the model of the pendulum itself acts as a data model. This makes sense, of course, because the pendulum and its behaviour are "data" because they are directly observable. The only sort of information model we need, then, is the models of how the measurement devices "process" the information they are measuring, and how measurements can be "reduced" to the data model, which is just the pendulum

model.

Thus, the implementation model applies to the simple pendulum at the expense of becoming effectively trivial. But this also makes sense: because the simple pendulum is such a simple case and it can be directly observed, we should expect a more generally applicable model of data handling to apply to it in a trivial way. Aside from showing that the implementation model applies in the simple case of pendulum modeling, this also helps to make a point about modularity in the following way. Considering the move in the opposite direction, *i.e.*, from the pendulum to the NEO modeling case, we see that a data handling model would have to be "expanded" to account for the complexity of data handling in that case. This would require introducing an information transmission model and a measurement data model between the theoretical model of the pendulum and the measurement devices. This structure, *viz.*, the theoretical model→information transmission model→measurement data model, is then a "module" that needs to be included in order to account for data handling in the NEO modeling case.

Now, extrapolating this case a bit, there could easily be cases for which we need to "expand" the simple implementation model discussed above in order to model the structure of more complicated kinds of data handling. But such a case could also require the "collapse" of some of the structure of the implementation model in the same way that the implementation structure must be collapsed to apply it to the pendulum modeling case. Thus, we may expect data handling in applied mathematics more generally to involve a variety of these kinds of expansion and collapse operations, picking out modular structure that is needed only for modeling certain sorts of data handling found in certain sorts of cases. This is how the models of feasible epistemology are expected to work with complex phenomena, *viz.*, we do not have a fully general model that applies to all cases; rather, we have structure modules that clarify certain kinds of structure and the epistemological modeling task is to identify the modules that are needed to account for the epistemology of a given case, or facet, of theory application.

Chapter 5

Computing in the Real World: Feasibly Solving Mathematical Problems

5.1 A Feasible Approach to Scientific Computing

5.1.1 Feasibility Constraints on Computing

Methods of computing solutions to problems is generally a matter that is considered to be of no philosophical concern. This belief derives from an attitude that methods of computing solutions are "merely pragmatic" and so need not be included in philosophical analysis. This view is changing, as is reflected in the work now done by many philosophers, *e.g.*, Winsberg (2001), Humphreys (2009), who study various aspects of mathematical simulations of phenomena. But numerical methods themselves, which are the engine of the application process in real science, and the reliable handling of error that such methods introduce have still received very little philosophical scrutiny. Given the centrality of these methods in the application of theories in modern science, a careful consideration of how these methods fit into the entire application process is overdue, and this is one of the aims of the current chapter. Moreover, despite the increasing interest in mathematical simulations, the view that methods of computing are "merely pragmatic" is quite persistent.

As an aid to clarifying the significance of computing methods, let us consider the classical picture of applying theories by deriving predictions from theory. There are three main stages to this process: the derivation of laws, or equations of motion, from theory; the derivation of predictions (solutions to the equations) given certain imposed conditions (boundary conditions and initial conditions); and the evaluation

of the solutions. The third stage of the process is considered to be so insignificant that it is not even included in this picture of the theory-world relation. But function evaluation is a problem studied in numerical analysis and poses problems in terms of numerical stability. Recall our discussion of the concept of conditioning in the previous chapter. If a function is ill-conditioned, this means that a small change in the value at which we are evaluating the function will lead to a large change in the computed value of the function. Thus, in cases where we need to use numerical computation to evaluate functions, a consideration of the conditioning of the function is essential to ensure accurate results. But function evaluation is an old problem, and one for which computational solution has long been used. A more recent issue is the need to use computation to find solutions to differential equations.

The second state of the process, the solution of equations of motion, is classically viewed as a process of derivation of a closed form solution to the equation. This is what would then be evaluated using the "merely pragmatic" methods of computation. But as we have seen in the consideration of the case examples, when the equations are nonlinear such processes of solution are rarely available. When they are, the systems being modeled have a high degree of symmetry, which allows the dominant behaviour to be captured in much simpler equations that can be solved analytically. Most phenomena do not have a high degree of symmetry and are nonlinear, which forces the use of numerical methods to compute solutions. The use of numerical methods to solve differential equations is not new, with simple methods going back to Newton and to Euler and Lagrange, who developed more sophisticated methods for solving differential equations by solving difference equations (Goldstine, 1977). Despite the long history of numerical methods for differential equations, the ability to feasibly model complex phenomena using numerical solutions is a more recent development that required high-speed computing devices, as was mentioned in chapter 1. Given how powerful computing devices have become, it is increasingly common to use computers to solve equations. And there are phenomena that are understood by no other means.

This shows that the second stage of the process has been taken over by numerical computation in a wide variety of cases. Thus, a task that is classically part of the "derivation" process involved in applying a theory to phenomena is now often per-

formed by computer. I argue that this alone shows that numerical computing has become part of the mathematical modeling process itself, rather than a "merely pragmatic" step tacked on at the end of a modeling task to get numbers out of equations. In this chapter we will see that there are actually much stronger reasons to support this view, since modeling considerations play a significant role in the choice of numerical methods *and* in the choice of how to handle the error introduced in computing. Thus, only the first stage is left within the classical domain of theory.

We may see that this too is changing as a result of the emergence of computer algebra. Computer algebra, or symbolic computation, uses computers to perform analytic manipulations of equations that are otherwise laborious or infeasible to perform by hand. The advantage of computer algebra methods over numerical methods, however, is that the resulting equations are exact or use the same approximation methods that ordinary analytic manipulations use.[1] Computer algebra methods have been used for facilitating computations in celestial mechanics and general relativity since the CAMAL (CAMbridge ALgebra system) system was developed at the University of Cambridge in the late 1960s (Geddes *et al.*, 1992). By 1970, computer algebra systems were also being used for high-energy physics calculations. This shows that even the first stage of the derivation of descriptions or predictions from theory has been infiltrated by computation. Indeed, given the explosive effect the numerical computing has had in terms of what problems become accessible to analysis, computer algebra stands to have a similar impact by allowing the construction of exact models that could not be feasibly constructed by a human being.[2]

Thus, we see that, in general, the entire process of model construction and solution in science uses computational methods. Given that error is introduced in numerical computation it is important to understand the effects of these kinds of error in order to know that a method of computing solutions is numerically stable, *i.e.*, will produce reliable results. We will, thus, consider the effects of numerical error on model problems later in this chapter. We will see how the perspective of algebro-geometric constraints and the concepts of local constraint logic help to clarify the effect of nu-

[1] There are concerns, however, that arise from a different source. Rather than truncation and rounding error, the concern is *expression swell*, *i.e.*, where the form of the algebraic expression produced by the computer becomes unmanageably complex.

[2] Examples already in use by engineers that illustrate the power of these kinds of methods are the advanced design, analysis and simulation systems MapleSim and Simulink.

merical error on the structure of a mathematical model. Following this, we will see that the structural stability of methods of numerical computing involves essentially the same kind of structural stability seen in reliable inference using data.

5.1.2 Computational Inference in Local Constraint Logic

Up to this point we have been using the general concept of local effective validity in terms of the solutions of a constraint system that are consistent with the imposed and assumed constraints. We saw initially in chapter 3, however, that a more specialized concept is available when rules of inference are available. What we will now see that the same is the case when a model specifies dynamical behaviour, which is done in each of the cases considered in this study. In this case, also, there are two basic modes of presentation of local constraint logic. The one we have considered so far treats solutions as fixed objects or states in some space of solutions. This can include the consideration of states in phase space as solutions, or sets of states in phase space, as solutions. Thus, a trajectory is also treated as a fixed object—it is a state in a space of trajectories on the state manifold of the system. This is the view of dynamics on the *state* mode of presentation of local constraint logic. The other mode of presentation is available in cases where we have specified generators of state transformations. In this case, we can consider not only consistent solutions, but also the possible paths through the space of states produced by the selection of generators. The result is the view of dynamics on the *path* mode of presentation of local constraint logic, where we have shifted from considerations of consistency to considerations of (dynamical) *consequence*

For example, consider the case of numerical methods that solve time evolution problems. This requires the computation of a state at a later time given one or more states at earlier times. This is accomplished by a (*discrete*) *state transition map* T_i, which provides the instructions for how to compute the next state given one or more previous states of the system. In this case, we care not only about whether a given state is consistent with imposed constraints, we want to know if a given state is accessible from some other state from the application of the state transition map. Suppose that, in the context of a system of constraints $\mathcal{C}$, that a state y_k is accessible from a state y_0 by k repeated applications of the state transition map. We may the

write

$$y_0 \Vvdash_{\mathcal{C}(T_i)} y_k$$

to express that there is a *valid path* from y_0 to y_k from the repeated application of T_i. Similarly to the state mode of presentation, the concept of a valid path can be made to be error tolerant and locally stable. For example, suppose that the values y_i are the states of some discrete-time dynamical system with initial condition y_0, and suppose that $\tilde{y}_i$ are the values computed by some numerical method. Then we could express the statement that the numerical method accurately computes the values y_i of the dynamical system within a error of around ε for at least N iterations of the method as

$$y_0 \Vvdash^{\|\epsilon\|\lesssim\varepsilon, k\leq N}_{\mathcal{C}(T_i)} y_{k+1},$$

where $\epsilon = y_{k+1} - \tilde{y}_{k+1}$. In dynamical systems terms, this means that the solution generated by the numerical method shadows the exact solution for at least N iterations. In terms of local constraint logic, this expresses the local effective validity of the path generated by the numerical method.

External semantic relations in the path mode are also treated in a similar way to those of the state mode. Suppose that we are using a numerical method with state transition map T_i to compute solutions to a continuous-time dynamical system $\mathcal{D}$ with initial condition y_0. Suppose that the numerical method computes values at a sequence of times t_k and that the solutions generated by the numerical method are $\tilde{y}(t_k)$. Letting $y(t_k)$ be the exact solution to the system $\mathcal{D}$ at the same sequence of times, then we could express the statement that the numerical method accurately computes the values $y(t_k)$ of the dynamical system within an error of around ε for at least N iterations of the method as

$$y_0 \Vvdash\!\!\!=_{\mathcal{D}(T_i)} \tilde{y}(t_k) \qquad \Rightarrow \qquad y_0 \Vvdash^{\|\epsilon\|\lesssim\varepsilon, k\leq N}_{\mathcal{D}(T_i)} \tilde{y}(t_k),$$

where $\epsilon = y(t_k) - \tilde{y}(t_k)$. Once again, in dynamical systems terms, this means that the solution generated by the numerical method shadows the exact solution for at least N iterations. In terms of local constraint logic, however, this expresses the local effective completeness of the path generated by the numerical method as an implementation of the dynamical system.

One final matter that we should address here before we continue on with the chapter concerns the concepts of numerical stability from numerical analysis. Using the concept of a condition number (κ) relating the forward error (ϵ) and backward error (δ), as was discussed in the previous chapter, we can define the stability concepts of numerical analysis in terms of soundness and completeness. Consider the algorithm implementing some numerical method as the source framework and the mathematical model as the target framework. Let T_i denote the state transition map of the numerical method and φ_t the state transition map of the mathematical model. Let $\mathcal{A}[x](T_i)$ denote the algorithm framework, where the constrained quantity is in square brackets and the state transition map is in parentheses. A solution in the algorithm framework is the result of the algorithm for a given input. So, if $x(p) = s(a)$ is the result of using some machine to run the algorithm on input $p = a$, we write

$$p = a \Vdash_{\mathcal{A}[x](T_i)} x(p) = s(a).$$

Now, there are different possibilities that could obtain when a given solution is interpreted back in the framework $\mathcal{M}[x](\varphi_t)$ of the model, which requires the discrete numerical solution to be interpolated to obtain a continuous one. One possibility is that the backward error on input a, *i.e.*, $\delta(a)$ is small. More precisely, $\|\delta(a)\| \lesssim \varepsilon_b$. If this holds for all inputs, then for all values of p

$$\Vdash_{\mathcal{A}[x](T_i)} x(p) = s \quad \Rightarrow \quad \Vdash^{\|\delta(p)\| \lesssim \varepsilon_b}_{\mathcal{M}[x](\varphi_t)} x(p) = s.$$

This condition is called *backward numerical stability*. If in addition to a small backward error, the forward error is also small on all inputs, then the algorithm is *numerically stable*, which we could write as

$$\Vdash_{\mathcal{A}[x](T_i)} x(p) = s \quad \Rightarrow \quad \Vdash^{\|\delta(p)\| \lesssim \varepsilon_b, \|\epsilon(p)\| \lesssim \varepsilon_f}_{\mathcal{M}[x](\varphi_t)} x(p) = s.$$

Forward numerical stability is then defined in a similar way except where the error condition is $\frac{\|\epsilon(p)\|}{\kappa} \lesssim \varepsilon$. Thus, forward, backward and numerical stability can be construed as kinds of effective soundness of an algorithm implementing a model on a machine (or, equivalently, as kinds of effective completeness of a machine implemen-

tation of a model).

5.2 Structural Features of Feasible Computing

Let us recall the basic structure of the model construction process we considered in chapter 3. We begin with a question concerning some phenomenon that we wish to represent as a collection or series of mathematical problems; the solution to these problems can then provide an answer to the question. To construct a mathematical model problem, we select or construct a system of equations that determines a theoretical framework and then select or construct a number of constraints that govern the structure or behaviour of the phenomenon, *e.g.*, symmetry assumptions or constitutive relations/equations of state. This allows the construction of a mathematical model (which may then require model reduction methods to reduce complexity) by eliminating extraneous information, and allows us to focus in on constraints on certain variables that govern the dominant behaviour of the phenomenon being modeled. The result, then, is a model problem that contains information about how the dominant features of the phenomenon behave under various conditions. Accessing the information about behaviour that we require usually requires finding solutions to the model problem under specific conditions. In most contemporary cases, finding solutions requires the use of numerical methods and machine computation. So thinking in terms of feasibility, analytic techniques take us only so far and we require computational aids to extract the information we need from the theory.

In the cases we have been considering, the theoretical framework is capable of both microscopic and macroscopic models of phenomena. The most accurate model of the phenomenon is a microscopic model, which is regarded as being the most "ontologically correct" model of the phenomenon in the framework of the theory. This model may be taken to exist in a classical sense, but we have little or no epistemic access to it. Instead, when considering macroscopic phenomena, we typically start with a macroscopic base model constructed on the basis of macroscopic structural and behavioural constraints, which are chosen on the basis of known or hypothesized structural and behavioural characteristics of the phenomenon. The macroscopic base model is typically a classical base model, *i.e.*, one that we cannot feasibly write down

equations for, but may also be a feasible base model that is too complex to effectively draw inferences from. So, the macroscopic constraints we choose explicitly may not produce the base model directly. And typically, whether we construct the base model explicitly or not, in order to access an equation that gives sufficient insight into the behaviour we seek to model, we use model reduction techniques to give a simple equation that still describes the behaviour we are interested in. In order to then know how the aspect of the phenomenon that the model can describe actually behaves, we need to be able to solve the model problem subject to specific conditions. For most real phenomena, we cannot do this analytically, so the scientific inference incomplete and blocked—we cannot yet access the conclusion of the inference of behaviour.

It is at this point that numerical methods and machine arithmetic come in to overcome the block and allow the completion of the process of scientific inference. This often, and always in the cases we are considering, happens at the expense of discretizing the problem in a quite radical way. To make the inference feasible, the model of the behaviour must be translated into a form in which the behaviour can be deduced using some kind of recursive, or iterative, process that can be run on a microprocessor.

Since the ability to make the inference feasible has such a strong effect on the mathematical framework, rendering continuous space and time both discrete and finite, it is important to have a reliable way of understanding the effect of the computation on the quality of the solution obtained through the computational process. This requires a reliable of analyzing error to ensure that error remains within tolerable limits. Thus, if we want to understand how scientific inference is feasible in contemporary science and what constraints there are on actually possible knowledge, we need to understand how numerical methods allow scientific inferences to be made and what ensures and explains why these inferences are reliable given the methods used to make them.

The logical approaches in the manner in which they are now pursued are not up to the task of gaining insight into the real constraints on possible knowledge imposed by the conceptual and computational tools that are accessible to us. The level of abstraction of the received view and its ignorance of any kinds of feasibility concerns obviously make such an approach inappropriate for understanding real constraints

on possible knowledge given the tools we have available to us. A more up-to-date syntactic approach could be useful in navigating the relationships between different languages used in the process of computation, as we will see below. But, as I will argue, syntactic approaches, given their usual philosophical conception in terms of rigorous formulations in a (typically first order) formal language, lack the flexibility and specificity required to identify and elucidate the strongest determining factors on possible knowledge in real epistemic contexts.

Approaching this problem from a semantic point of view may allow us to gain insight into the general impact of the transformation from continuous domains to finite bounded domains, which could give insight into general constraints on knowledge possible when such transformations must be used. What will not be visible from this point of view, however, are the much tighter restrictions on possible knowledge when such transformations must be used to solve problems when solutions must be found using *limited time and resources*. Thus, on this approach we can gain insight into constraints on knowledge for any possible computational procedure (and some classes of procedure), which does reveal general kinds of epistemic limitations in practice, but we cannot gain insight into the stronger constraints that determine what knowledge is possible for us to actually have in real epistemic contexts.

Accordingly, this section is concerned with clarifying the structure of the process of feasible computation in science. This involves developing an understanding of not only how the the tools used in that process place constraints on possible knowledge but also how they open up avenues of inference that would not otherwise be possible or feasible. A central concern here is how stable inference is ensured despite the introduction of error. The following section is a more detailed consideration of the nature and structure of the error introduced in the process of feasible computation and what is involved in ensuring the reliability of computational methods of scientific inference.

5.2.1 Constructing Algorithms

We saw in chapter 3 that the process of constructing an algorithm to solve a modeling problem begins at an analytic level and can require the use of mathematical simplification techniques just as in the process of model construction. In the case

of constructing algorithms to solve the impact probability determination problem, this requires the use of linearization and linear methods to reduce the complexity of the calculations required. This results in the identification of tasks that require numerical methods capable of performing the required computational task, *i.e.*, that produce results that are (locally) effectively valid when interpreted in the framework of the the mathematical model problem being solved. Methods that are able to do this are called *numerically stable*. Before discussing some of the considerations for selecting or constructing numerically stable methods, we must first clarify the concept of numerical stability.

The general idea of numerical stability is simply that for any given input the numerical method gives you the right answer to within acceptable error tolerances. Part of this requires the construction of an algorithm or method that will output or converge on an effectively correct answer in the absence of error introduced in the process of computation. This is a necessary but not sufficient condition for numerical stability, however, because accuracy or convergence must obtain despite the fact that error is introduced in the process of converting an algorithm into a form that can be run on a computing device. To get a sense of the issues involved, consider the following example.

A very simple nonlinear differential equation is the *logistic equation*

$$\dot{y} = y(1 - y).$$

Its variants are probably best known for their importance in modeling population growth in the presence of a carrying capacity, but they also have applications in many other areas of science. This equation has a closed form solution, the sigmoid curve, but let us suppose that we wish to compute a solution numerically subject to the initial condition $y(t_0) = 1$. The first thing we need to do is *discretize* the time variable to get an expression of the equation as a difference equation. One standard way to do this results in an equation of the form

$$y(t_{n+1}) = y(t_n) + h\Phi(y(t_n), h),$$

where h is the size of the time step, *viz.*, $t_{n+1} = t_n + h$. This function Φ tells us how

to evolve the value of the solution from one time step to the next. The simplest way to do this is to evolve the solution forward in time by assuming that the solution function is linear, so that the difference equation we get is

$$y(t_{n+1}) = y(t_n) + hy(1 - y).$$

This is called the Euler method. Since the solution function $y(t)$ is not linear, it is a curve, this assumption is false, which means that the way we have discretized has introduced an error, called *truncation error*.[3] If we keep the value of h sufficiently small, then it may be possible to keep the error introduced in this approximation small, however, in which case the numerical method will be numerically stable. How small h needs to be, however, depends on the problem and the period of time over which the solution needs to be calculated accurately. In this case, choosing h to be, say, 0.01 results in a quite accurate solution. For other problems, however, h might have to be reduced to such a small value that it could take more than a human lifetime to compute. Thus, the first thing we need to ensure is that the method or algorithm we construct can give us accurate solutions for the problem at hand. It does not end here, however, since the computation required to a get a solution from a method can rarely be computed by hand; instead, it must be performed by a computer processor.

In order to actually perform the computation, we need to implement the numerical method on a machine. Assuming that we have correctly implemented the algorithm of a numerical method, we have still not ensured that the implementation will be numerically stable. The simple reason for this is that machines cannot represent the real line, nor can they represent the rationals, since they must be represented in a way that is discrete and bounded. The result of this is that machines can neither represent arbitrarily large quantities nor arbitrarily small ones. Thus, if the result of an arithmetical computation is a number that is between the two nearest numbers in the discrete representation of the reals, it must be rounded to the nearest value. This kind of error is called *rounding error*. Thus, numerical stability not only requires stability under truncation error, it also requires stability under rounding error, and the effects of both can be quite subtle.

[3]The reason for this term will be made clear in the next section.

A simple way of showing problems that can arise is for simple arithmetical computations. Suppose that we wish to compute the result of

$$x + y - z$$

for $x = 10^{17}$, $y = 1.618$, $z = 10^{17}$. Clearly, the result is 1.618 in this case, but because of the way that the standard implementations of floating point arithmetic work, if we performed this computation as

$$(x \oplus y) \ominus z,$$

where the circled operations indicate their floating point versions, then the result would be 0. This is because standard floating point only has a precision of 16 digits after the decimal place, so when 1.618 is added to 10^{17}, the result in floating point arithmetic is exactly 10^{17}; the 1.618 is lost because there is not enough precision to capture the large integer part and the relatively small fractional part.[4] If we performed this computation as

$$(x \ominus z) \oplus y,$$

however, then we would obtain the correct result because the two large values cancel. The arithmetical versions of the last two formulae are mathematically equivalent, but the floating point versions are not. Thus, the ordinary laws of associativity and commutativity do not hold in floating point arithmetic.

Thus, numerical stability requires that the algorithm will give sufficiently accurate results not only in exact arithmetic, but also given how it is implemented on a computing device. This requires understanding the how the algorithm works in principle, to know when it will be reliable, and how to implement it on a computing device so that the results from the use of machine arithmetic will still be accurate.

It is important to distinguish the notion of numerical stability from that of conditioning, introduced above. Recall that conditioning pertains to the sensitivity of a problem under changes of the input. A problem for which small changes in the input produce small changes in the output is well-conditioned, and one that produces large

[4]To be clear here, it is important to recognize that it is the relative size of the quantities that is the problem in this case. If the values were $x = 1$, $y = 1.618 \cdot 10^{-17}$, $z = 1$, exactly the same problem would arise.

changes in the output is ill-conditioned. Recall also that a measure of conditioning, the condition number κ relates the forward and backward error by an equation of the form

$$\|\epsilon\| \lesssim \kappa \|\delta\|,$$

where δ is the backward error and ϵ is the forward error. Although both numerical stability and conditioning involve concepts of stability, they are types of stability under different kinds of variations. Conditioning pertains to the sensitivity of the problem and numerical stability pertains to the stability of the method used to solve it. The relationship between conditioning and numerical stability can be expressed with the slogan: *a numerically stable method is only guaranteed to give accurate answers on a well-conditioned problem.* This expresses the idea that some notions of numerical stability do not actually guarantee accurate solutions.

To explain this seemingly contradictory situation, consider the following case. Chaotic dynamical systems, like the double pendulum system, have the property that small changes in initial conditions grow exponentially fast, leading to large differences between trajectories over time. Thus, chaotic systems are ill-conditioned. One definition of numerical stability requires that the method used gives solutions with a small *backward error* (over the range of inputs of interest). Algorithms with this property are called *backward numerically stable*, or simply *backward stable*. Well, numerical solutions of chaotic dynamical solutions often give small backward error, but the trajectory computed by the method will be useless if one wishes to use it to know precisely where in phase space a chaotic system will be at any given time. Another definition of numerical stability, however, and the more typical one, requires both a small backward error and a small forward error. This kind of numerical stability, called *mixed stability*, cannot be ensured on an ill-conditioned problem. So, the failure of a backward stable algorithm to compute accurate solutions to an ill-conditioned problem is not a problem with a method but, rather, with the application of such a method to an ill-conditioned problem.[5] The relationship between forward

[5]It should be noted, however, that there are occasions when the user is wise to do this. As Corless (1994*a*) has pointed out, a backward numerically stable algorithm applied to an ill-conditioned problem can still be useful provided that the *quantities of interest* are stable over small variations of the input. For this reason, backward stable algorithms can be useful for studying chaotic dynamical systems. We saw two examples of this in chapter 3, with computing trajectories of a NEO past the Lyapunov time and studying chaos for a double pendulum, where information can be obtained from

and backward error ensures that a backward stable algorithm on a well-conditioned problem will give accurate results. So, a backward stable algorithm applied to a well-conditioned problem will also be mixed stable.

Now, numerical stability is a necessary but not sufficient condition for an algorithm or method to solve a problem. The reason for this is that we do not simply require methods that can give us accurate answers in principle; rather, we require methods that give accurate answers in a reasonable period of time. Using the Euler method to compute solutions to differential equations could provide accurate results for most problems, provided that we could take the value of h to be sufficiently small. The problem is that as h becomes smaller, the number of steps needed to evolve the system forward in time by a given amount increases. For example, if a problem required us to reduce h to 10^{-40}, then even using our fastest supercomputers it would take longer than the age of the universe to evolve the system forward in time by one second.

The missing ingredient in making computations actually *feasible* is an algorithm that can give accurate results in a reasonable period of time given available technology. Since the performance capacity of available technology varies over time and place, and since what counts as a reasonable period of time depends on the demands of the problem, the concept of a feasible computation is a strongly context dependent notion. The branches of applied mathematics that deal with considerations of time and resources required for computing are computational complexity theory and the analysis of algorithms. Here computational problems are typically classified according to their *time complexity*, *i.e.*, the number of steps they require for a given input to the problem (see Sipser, 2012, chapter 7).[6]

In the case of computational problems, the interest is typically in the *computational complexity*, which is the scaling behaviour between the size of the input and the number of steps the best algorithm requires to obtain a solution. Suppose that the size of the input is measured by a natural number n; then, if the number of steps required by the best algorithm increases linearly with the size of the input, the problem is said to be solvable in linear time. Problems for which the increase

the model by replacing a deterministic study with a statistical one (see pp. 119 and 121).

[6]The amount of storage space, the *space complexity*, is also sometimes considered. In practical computing, this can be important, since in may cases memory access time a greater limitation than processing time (see Sipser, 2012, chapter 8).

is n^p for some natural number p are solvable in polynomial time; for exponential increase, exponential time. Since this classification is based on the scaling behaviour of the *best* algorithm to solve the problem, which is usually not known, computational complexity is a measure of the hardness of the problem.

The study of the complexity of computational problems is generally a theoretical concern, and so is studied independently of any particular algorithm used to solve the problem. In practical computing, however, it is often important to know the scaling behaviour of a particular algorithm, and often also for particular implementations of a given algorithm. Essentially the same concepts apply in the case of analysis of algorithms, except that a model is now required to specify what counts as a step. In order for there to be a reliable relationship between the number of steps and the computing time, the time per step must have a constant upper bound. A model of this sort is called a *cost model*, and it is the basis for referring to the computing time required by an algorithm as the cost of the algorithm. A common cost model measures steps simply in terms of machine operations, such as floating point operations. But for more specific applications, such as solving differential equations, the cost may be measured in terms of the number of times that the vector field must be evaluated.

Therefore, we may see that the conditions for numerical stability are quite complex and that some depend on the problem and others on the algorithm used to solve it. In terms of the problem, the two major concerns are conditioning and computational complexity. In terms of a particular algorithm and its implementation, the two major concerns are numerical stability and cost. The relationship between computational complexity and cost is similar to the relationship between numerical stability and conditioning. Solutions may be feasible with an algorithm that is applied to problems with a certain range of computational complexity, but the algorithm could cease to be capable of feasibly finding solutions for problems outside of this range because the computational cost becomes too high.

Returning now to the question of feasibly generating accurate solutions to a computational problem, we may see that the dominant concerns are the accuracy of the solutions generated by a method and its computational cost. Any given application will place constraints on acceptable tolerance for error in the solution and will also place constraints on acceptable tolerance for computing time. Depending on

the manner in which the two are important, there will be trade-offs between these two considerations, and a balance between accuracy and cost must be sought, as we will now see from the consideration of the numerical methods used in the dynamical astronomy problems we have been considering.

In any case involving the choice or construction of numerical methods to solve a mathematical problem, one is looking for the fastest algorithm possible subject to the constraints of the case. These constraints are quite varied, however, and so across cases there is no simple general criterion for what counts as the fastest method. General purpose computing systems like MATLAB include robust algorithms that can be applied stably to a wide variety of problems. For example, the MATLAB function `ode45` is a general purpose ODE solver that can be supplied with a tolerance ε, and it attempts to compute solutions efficiently with the error in the solution being below around ε. This algorithm does very well on a wide variety of ODE problems. But methods like `ode45` are not sufficiently accurate, are too slow, or are unreliable for some special problems that require numerical methods that are tailored to the specific nature of the problem.

One such special ODE problem arises from newtonian mechanics models that do not account for dissipation of energy. Most numerical methods require that such problems be reformulated as a pair of first order differential equations on phase space. But for second order ODE without dissipation (no first derivatives), the second order equation can be solved directly, which is of interest in cases where the values of the first derivatives are of little significance in the modeling and do not need to be computed (Henrici, 1962, 289*ff.*). The method of this sort that is common in dynamical astronomy is called Cowell's method. This method has been widely used in the history of computational astronomy and is used by JPL to compute planetary ephemerides and to solve the orbit determination problem (Chodas & Yeomans, 1999). An advantage of this method in terms of optimizing accuracy and cost is that it is a *multi-step* method, *viz.*, a method where previous values of the solution (and its derivative) are used to compute the next value. In comparison to the more common *Runge-Kutta* methods, which are single step methods, multi-step methods attempt to gain efficiency by computing the next step using information already computed in previous steps, meaning the computational cost per step is smaller in comparison to

Runge-Kutta methods. There are also other features of the JPL code that optimize speed and cost that we will mention below.

Hamiltonian problems form another type of special problem. As we have seen, such problems have special geometric structure, namely a symplectic form on the phase space that ensures the conservation of various areal quantities and the energy of the system over time. Most numerical methods pay no attention to this structure and as a result either accumulate or dissipate energy over time. Thus, they can obtain accurate results only over relatively short periods of time, after which the error grows too large and the method becomes unstable. It is possible, however, to construct methods that *preserve* the symplectic structure on phase space and simultaneously approximately preserve energy.[7] For these methods, the map on phase space evolving the system forward in time is a symplectic map. Not unsurprisingly, such numerical methods are called *symplectic methods*. When implemented, symplectic methods can provide accurate solutions over much longer periods of time than typically possible with other methods. Since studies of the dynamical stability of the solar system require accurate computations over millions or billions of years, these methods have been important for solar system computations.

The symplectic methods now in use in dynamical astronomy derive from a method developed by Wisdom & Holman (1991). This method obtains gains over arbitrary symplectic methods by taking advantage of the fact the the N-body problem in the solar system has a dominant central mass. The advantage of this is that in addition to the energy conservation properties of a symplectic method, they are also considerably faster than other methods for solving the N-body problem with a dominant central mass. This illustrates how the detailed features of a particular type of problem can allow the construction of algorithms that are particularly efficient for solving problems of that type.

Balancing cost and accuracy in numerical methods typically requires certain kinds of adjustments of the method *during* the computation. One very common way of

[7]Note that it can be proved that a symplectic method that preserves both symplectic structure and energy is a discrete solution of the original problem (this is proved in Zhong & Marsden, 1988). As a result, it is not feasible to simultaneously preserve symplectic structure and energy, but it is possible in general for problems with periodic solutions, which includes orbits, to preserve symplectic structure and *approximately* preserve energy over long time intervals (see Stuart & Humphries, 1996, 621-623,636).

doing this, used by `ode45`, JPL and some symplectic methods, is to use a measure of the error to adjust the size of the time step. The basic idea is to make the time step as long as possible to keep the error within the desired tolerance ε, but to decrease the time step when necessary to ensure accuracy in a way that is optimal in some sense. This can have a dramatic effect on the length of time required to compute solutions of a given accuracy. A similar kind of adjustment is also often used to optimize accuracy and cost by changing the numerical method used during the computation. This is sometimes done in celestial mechanics computations to deal with close approaches.[8] According to Chodas & Yeomans (1999), the JPL code does this by changing the version of the method used to one capable of higher precision in order to handle close approaches. A more dramatic version of this occurs in the computation of the general relativistic version of the two-body problem, *viz.*, two black holes.[9] For sufficient distances the gravitational fields are sufficiently weak for post-newtonian mechanics to be effectively valid, but as the black holes get close together, the numerical method must switch to one that solves the field equations of general relativity.

We can see from these cases that the choice of numerical method can be quite strongly dependent on the modeling problem at hand. Moreover, we see that it is often necessary to tailor a numerical method to a particular problem to solve it both accurately and efficiently. This gives us another indication that numerical methods are really part of the mathematical modeling process, rather than being tacked on at the end for merely pragmatic reasons in order to get solutions. If we cannot find a numerical method that can solve a problem accurately at sufficiently low cost, then we do not get solutions to that problem. In such cases, we cannot find out what the predictions of a theory are under the given circumstances. When it comes to the difference of being able to know the predictions of a theory and not being able to know, it is hardly a "merely pragmatic" issue.

The fact that computing is necessary to actually determine the consequences of theory highlights an important difference between feasible epistemology and classical epistemology. Feasible epistemology is geared toward understanding what the

[8] There are symplectic methods that can handle close approaches accurately without changing the method (see Chambers, 1999).

[9] Thanks to Chris Smeenk for this example.

actual limitations of feasible scientific inference are, which includes cases where no accurate, low cost numerical method exists for a problem. This indicates a clear limitation in our ability to draw inferences from theory, which means a limitation in the range of applicability of the theory in practice. Thus, the limitations of numerical methods indicate an important boundary to feasible scientific inference and an important boundary of applicability of a theory. This obviously differs markedly from the classical boundary of applicability of a theory, which is not dependent on context and does not need to deal with issues of error tolerance of computational methods or the amount of time it takes to compute things. Although a boundary of feasible applicability does not indicate that the theory fails, it does indicate a very real limitation in our possible knowledge. As we shall see, there are a variety of boundaries of applicability in feasible epistemology, each pertaining to different kinds of epistemic limitations.

The classical formulation of a theory, however, is an essential part of what allows it to have a range of feasible applicability as wide as it does. We saw in the last chapter that classically exact models are essential for allowing flexibility in terms of finding formulae that are still effectively valid in a given context but that are much easier to draw inferences from. We see now that this is a kind of balancing of accuracy and computational cost. There is another way in which classical descriptions have this advantage in the context of numerical methods. Because truncation and rounding errors are generally introduced in any numerical computation, better numerical stability and accuracy are obtained by avoiding unneccesary computations. An important way of doing this is to use a combination of analytic calculations and numerical ones, absorbing some of the computation into an analytical calculation performed before the algorithm is implemented. Such methods are called *semi-analytic*.

A clever example of this is a method that is now commonly used for chaotic problems in the solar system. Using an improved version of Wisdom and Holman's original symplectic method, Mikkola & Innanen (1999) show how to compute the Lyapunov time and the state transition matrix that is required to compute the covariance matrix needed to estimate impact probability, in a semi-analytic way as a byproduct of the main computation of the orbit. This provides a boost in both accuracy and efficiency in solar system computations involving chaotic orbits. What is clever about

their approach is that they perform the analytic calculation by computing the variational derivative of the algorithm itself. The general idea is that since the algorithm is tracing trajectories forward in time, the variational derivative of the algorithm provides information on how trajectories change with a small change of the initial condition. The power of this approach is that the expression for the behaviour of nearby orbits computed analytically can be evaluated numerically using the results of the computation of the orbit of an object, allowing efficient and accurate computation of quantities such as the Lyapunov time of the orbit. This is a good example of the effectiveness of combining analytic and numerical methods.

One last example, to mention before moving on, shows how avoiding redundant calculations can accomplish a similar result. Rather than solving the full N-body problem to find the orbit of a NEO, the method of solving the orbit determination problem uses ephemerides for the planets as well as minor planets and some asteroids. The result of the JPL ephemerides being computed to extremely high accuracy is an increase in accuracy and numerical stability and a decrease in cost. If we only need to compute the object's orbit, then this will obviously improve efficiency, and one can hardly expect to improve on the JPL ephemerides, so accuracy is also improved. Additionally, by only having to compute the object's orbit, the potential for truncation and rounding errors to accumulate is reduced significantly, which improves numerical stability as well.

5.2.2 Modeling Computational Error

Another way to see that numerical methods are part of the modeling process comes from the fact that modeling considerations are not only entangled with the choice of numerical method but also with the way in which computational error is analyzed. As we have seen, the simple epistemic advantage of backward error over forward error, is that the backward error can be computed or estimated using the computed solution, whereas the forward error cannot. Furthermore, with estimates of the condition number of a problem, the backward error can be used to estimate the forward error. Despite this advantage, many numerical methods, including most of the ODE solvers in MATLAB, use estimates of the forward error over each time-step to determine how to optimally adjust the step size of a method. There are many cases, however,

where the advantages of backward error are employed to analyze and control error for solutions to ODE.

There are three main approaches to this sort of backward error analysis of ODE: shadowing; method of modified equations; and defect analysis (and control). The kind of backward error for ODE we have been considering throughout the thesis is *defect analysis*. As we have seen, the defect is computed by interpolating a numerical solution in such a way that it can be substituted into the model equation of motion. This allows a very straightforward, feasible computation of a nearby equation of motion that the computed solution solves exactly. In defect analysis the initial condition is held (effectively) fixed.[10] Shadowing methods work by showing semi-analytically that the numerical solution is following an exact solution to the original problem with a slightly different initial condition. Thus, shadowing methods work by holding the equation of motion fixed and allowing the initial condition to vary. The advantage of such a method is that one is assured to obtain an effectively valid solution to the original equation, at least for a certain period of time. Such methods find use in various applications, including assessing error in large-scale climate models (see, *e.g.*, Smith, 2002*a*). The last approach, which is useful for analytic studies of the error and stability of solutions to a numerical method, allows both the equation of motion and the conditions on it to vary.

This last approach is of significance in the present context because it is related to symplectic numerical methods. The way that it is proved that symplectic numerical methods are nearly energy conserving is by showing that the exact solution of the numerical method is a solution to a nearby hamiltonian problem, *i.e.*, a nearby problem *of the same general type*. Thus, rather than introducing an arbitrary perturbation, the numerical method introduces a *hamiltonian perturbation* of the problem, just like the other sources of perturbation in a hamiltonian problem. This is why it is ensured that symplectic methods are energy conserving. The proof of this fact is done using the method of modified equations. And the proof of the additional fact that symplectic methods are approximately energy conserving over long time intervals, even in the presence of rounding error, also employs the method of modified equations. This establishes, then, that solutions to symplectic methods on hamiltonian problems are

[10] Fixed, that is, except perhaps for rounding error.

exact solutions to a nearby nearly hamiltonian problem. Thus, this result from error analysis shows that the error introduced by the numerical method is of the same kind as the error introduced by other perturbations in the problem that are not taken into account. This is what makes symplectic methods particularly well-suited to hamiltonian problems, and shows that even error analysis for numerical methods is entangled with the modeling process.

The first approach, defect analysis, is of significance in the present context because it demonstrates the simplicity of one feasible means of error analysis for ODE. As part of the research for this study, I constructed a symplectic numerical method for the purposes of computing solutions to the double pendulum system. This allows the computation of a solution given any initial conditions. Once this is done, a continuously differentiable solution can be computed using a simple interpolation code. Using the theory for this kind of interpolation, the defect can then be computed semi-analytically using the interpolant of the computed solution. After that, the defect ($\boldsymbol{\delta}_{\mathbf{x}}(t)$ from equation 1.13) can be interpreted as a perturbation to the original equation of motion, showing that the computed solution to the problem is the exact solution of the resulting modified equation of motion; recall the right-hand size of equation 1.13, which is reproduced here:

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{h}(\tilde{\mathbf{x}}) + \boldsymbol{\delta}_{\mathbf{x}}(t).$$

The results of such a computation on the double pendulum system are shown in figure 5.1.

The epistemic value of computing the defect in this way is that it shows a simple, low cost method for finding the size of a time-dependent perturbation of the equation of motion required to find an exact solution to the original problem. This demonstrates that the solution we have computed using a computer processor is the exact solution to a problem that is only a very small perturbation of the original problem. Moreover, it is a perturbation that we can control the size of by reducing the time step. And since most modeling contexts will involve small time-dependent perturbations of the target phenomenon, such as vibrations do to thermal motion or passing vehicles, it is reasonable to interpret the defect as *modeling* these small

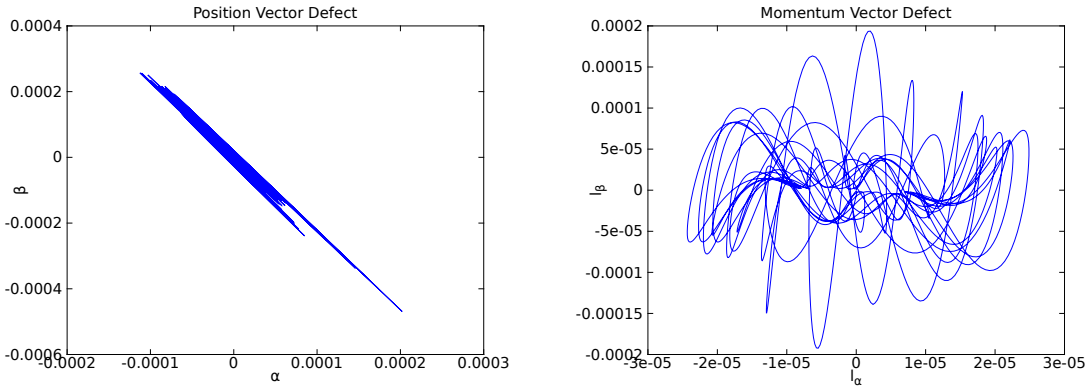

Figure 5.1: The position $(\alpha, \beta)(t)$ and (angular) momentum $(l_\alpha, l_\beta)(t)$ components of the defect $\boldsymbol{\delta}_{\mathbf{x}}(t)$ of a solution of the ideal double pendulum system using the Störmer-Verlet method with a step size of $h = 0.01$. Since the defect is a function of time, it has a phase trajectory just like the solution does, which is what the plots show. The maximum value of the defect is seen to be of the order 10^{-4}, the square of the step size h, which is an indication of the method being a second order method, *i.e.*, the forward error decreases proportionally to the square of the step size. This is explained further below.

errors that are not taken into account. In this way, rather than being a vice, error in the computation becomes a virtue, since it gives evidence that the model equation is stable under such perturbations.

These considerations only apply provided that the problem is well-conditioned, *i.e.*, not chaotic. If the problem is chaotic, then a small defect does not imply an accurate solution in the sense of a small forward error. In such a case, however, any error introduced in the modeling process will ensure that the exact solution to the model problem is not accurate (small forward error) as a description of the dynamical evolution of the target phenomenon either. The defect still has epistemic value in such cases, however, because if the model is to be of any use at all, then *some* quantities of interest must be stable under perturbations. An example of a quantity that we would expect to be stable under perturbation of the model is the Lyapunov exponent, or Lyapunov time. Thus, ensuring that the computed solution is the exact solution to a problem with a very small perturbation, computing the Lyapunov exponent from the computed solution gives us evidence that the Lyapunov

exponent of the model problem is very similar, and if the model is descriptively stable, also for the target phenomenon. Thus, in this way, the interpretation of a computational solution in terms of backward error analysis allows us to regard numerical solutions as providing evidence for properties of the model problem the numerical method was applied to, in much the same way that data provides evidence for properties of the target phenomenon.

The other significance of this kind of backward error analysis is that it shows how numerical methods are simply a continuation of the analytic modeling process. Recall that the whole reason that numerical methods are part of the modeling process is that they are taking over a key task of modeling that cannot be handled analytically. Recall also from chapter 3 that the feasible construction of a model equation of motion typically involves a whole host of approximations concerning the structure and behaviour of the phenomenon, introducing both reduction and abstraction error. Now, being able to interpret the computed solution as the exact solution of a perturbed problem means the problem we can feasibly solve is a nearby problem in a control space of the model problem. Thus, feasibly computing a solution forces us to modify the problem slightly to get an exact solution. But a tiny shift in control space is typically a much less significant change than the kind of projections that mathematical model reduction methods typically introduce, and such a shift is typically much smaller than the effects of error introduced by the selection of constitutive relations or equations of state. Thus, defect analysis allows us to show that the error introduced by computation is of the same order as or smaller than the other sources of error in the problem. Consequently, defect analysis allows a very low cost means of showing us that if the solution to the model equation is an effectively valid representation of the dynamics, then so also is a computed solution.

The result of this is that modeling approximations and numerical methods play essentially the same role in two different stages of the modeling process. The problem for which modeling approximations are employed is "find a valid equation of motion". Typically, we cannot feasibly do this with the full accuracy that the theory allows, so we use approximation methods in order to find an equation of motion that is effectively valid. The problem for which numerical methods are employed is "find a valid solution to the effective equation of motion". Again, we typically cannot

feasibly do this with the full accuracy the effective equation of motion allows, *i.e.*, by finding closed-form analytic solutions, so we use numerical methods to find a solution that is effectively valid. Therefore, it is evident that, far from being "merely pragmatic", numerical methods are just as much a part of the modeling process as the approximation methods used to construct models and play the same role in the second stage of the modeling process as model reduction methods do in the first.

5.2.3 Implementing Algorithms

If numerical methods are "merely pragmatic" tools for finding solutions to equations, their implementations on physical machines are *really* merely pragmatic. Consequently, it would be easy to ignore the details of this part of the process. But this would be a mistake from the point of view of understanding feasible epistemology, and not just because of considerations of error in machine arithmetic. In fact, implementations of algorithms are an essential part of what makes inference possible in real science, much like the reduction of raw data to a data model does.

Let us consider now the computational process from an algebro-geometric point of view. We have already seen the effects of feasible computation on the structure of ODE problems, *viz.*, numerical methods (generally) force time to be discrete and machine arithmetic forces time, phase space, and any other continuous space to be both discrete and bounded, and, hence, of finite and bounded precision and size. And these are the terms in which a semantic approach would analyze the process of computation. This is evidenced by the fact that, according to Suppes (2011), the most significant concern for extending the semantic view to considerations of practical computation is to axiomatize the floating point number system. This could indeed provide useful insights into the structural properties of floating point arithmetic, but I argue that this will give us very little insight into either the limitations imposed on our knowledge by machine arithmetic or what inferences machine arithmetic makes possible.

Among the aspects that such an approach is blind to is the epistemological significance of programming languages. The shift from a differential equation to a numerical method is significant, since it involves constructing a difference equation from the differential equation that can generate solutions that, when suitably interpolated, give

us effectively valid solutions of the differential equation. But this shift is reasonably simple and can be described in this simple way. The shift from a numerical method to its machine implementation however, is extraordinarily complex. It requires not only the task of writing code to implement the numerical method in some high level programming language, but also the complex procedure by which this code is converted into a form that can be run on a microprocessor, and the result of the microprocessor's computation registered in a form that can be interpreted in terms of the data types of a high level programming language. This involves, "algebraically", a complex sequence of translations between code in different computer languages, and, "geometrically", the representation of the meaning of the code in different forms, from high level data types, to machine language, and back.

The syntactic and semantic transformation that occurs from numerical methods to their machine interpretation is a major shift. Both the model differential equation and the difference equation of the numerical method lie within the ordinary framework of classical mathematics. This is why the time translation map of symplectic method is capable of being a symplectomorphism, along with the areal and energy conservation that this implies, because it is a map on exactly the same phase space as that of the differential equation. And it is in this framework of the numerical method that we interpret the results of a machine computation. The machine computation itself, however, does not lie within the ordinary framework of classical mathematics. The language in which machine implementations are described is lines of code written in a programming language. And though this code may include versions of the numbers, functions, and other quantities appearing in the difference equation problem, running the code requires that algorithm written in computer code be translated in various ways, ultimately into machine code.

Even the *theories* of all of the constructions involved in the process of running a numerical method on a machine are hardly theories in classical mathematics. Of course they could be formulated in such a way, but with these concerns we are firmly in the territory of applied mathematics and computer science. An important example is the classical theory of computation, which considers issues of computational complexity, but it does so without imposed constraints on time or resources. Accordingly, potentially infinite computations are entirely welcome in that framework.

The code running on a machine and the problems it is being used to solve, however, allow nothing of the sort. The accuracy and cost of implemented algorithms depends on much different considerations than what are accounted for in the classical theory of computational complexity. Included in this is the need to translate from mathematical descriptions (target phenomenon model, numerical method) to descriptions in computer code that can be run on physical machines. In this way, we see that the translation of a numerical method to its machine implementation involves a major shift in terms of both the language and methods used. This shift can be regarded as crossing a boundary between classical epistemology, which works with exact and ideal descriptions, to feasible epistemology, which works with real descriptions and is fundamentally computational in nature.

The need to be able to efficiently write code that can be used to compute accurate results brings out another aspect of feasible inference that is tied to the trade off between accuracy and cost discussed above. In order to write efficient and accurate code to implement a numerical method it is usually necessary to write the code in some high level programming language. Even writing and testing code directly in C or C++, however, can be extraordinarily time consuming, which is why systems like MATLAB, and its open source equivalent OCTAVE, are so useful—they simplify the programming task for general purpose computations. Although the coding task in C or C++ can be costly in human hours, and the task in MATLAB or OCTAVE much less for certain kinds of calculations, the choice of the latter sacrifices generality and flexibility, and potentially accuracy, for efficiency and ease of use. On the other hand, the choice of the former sacrifices coding speed and simplicity for certain calculations, for almost universal applicability and ability to optimize for special tasks. Thus, what approach is optimal depends very much on what the code is needed to accomplish.

There is another aspect of machine computation that shows its importance for feasible epistemology. We may see that the microprocessor and its associated memory play a similar role in feasible computation to the telescope and CCD camera in optical astrometry. In much the same way that the CCD camera provides a record of where the object is at a particular time, the computer memory provides a "record" of where the solution is at a time. In both cases, these "records" can be interpreted as a point in phase space under appropriate constraints. Moreover, in both cases it is the

interpretation of this "record" in the framework of the mathematical model of the target phenomenon that allows us, under appropriate constraints, to *know* where the system is in phase space at a time. Thus, the microprocessor plays a role in feasible computation that is similar to the role the telescope plays in feasible measurement.

Before moving on to the next section, it is worth noting that this analogy extends further in to the methods of design of microprocessors. Similar constraint approaches, such as that used by Moore (1996), are used in microprocessor design, but the difference here is that microprocessor design has to overcome a large difference in scale between the designer and the object being designed. The result is the use of constraints at different levels of abstraction. The implementation of the design at a high level of abstraction, which deals with the functional relations between different kinds of components, provides additional constraints at a lower level, but need not entirely pick out a particular design. Consequently, it requires a series of nested design constraints to result in an acceptable design of a physical processor. And at the lowest, physical level, there are similar concerns with respect to ensuring that the results of computations are not affected by physical perturbations from the environment. In terms of the analogy to telescope design, the design of a CCD device would be more comparable.

5.3 Sources of Computing Error and Their Structural Variation and Stability

5.3.1 Truncation Error and Its Structural Variation

The structural significance of truncation error is related to its effect on the mathematical problem being solved by a numerical method. The standard approach to constructing a numerical method to solve an ODE involves forming a difference equation from the differential equation. As pointed out above, this is often an equation of the form

$$y(t_{n+1}) = y(t_n) + h\Phi(y(t_n), h).$$

A standard means of constructing a difference equation from a differential equation of the form $\dot{y} = f(y)$ starts by expressing the solution $y(t)$ to the difference equation

as a Taylor series expanded around some initial time t_0:

$$y(t_0 + h) = y(t_0) + h\dot{y}(t_0) + \frac{h^2}{2}\ddot{y}(t_0) + \cdots \tag{5.1}$$

We are given $y(t_0)$ as an initial condition, and we can compute $\dot{y}(t_0)$ because we know that $\dot{y} = f(y)$. Thus, we can compute the first two terms of the Taylor series exactly. Indeed, Euler's method exploits this fact to provide a simple numerical approximation of the solution. The limitation of Euler's method, however, is that the increase in accuracy by choosing a smaller time step is linear, the practical upshot of which is that the accuracy gained by decreasing the step size comes at a large computational cost. For this reason Euler's method is called a *first order method*. For second, third and higher order methods, the increase in accuracy with a decrease in step size is quadratic, cubic, *etc.* Thus, the higher order the method the more significant the increase in accuracy by reducing the step size. The trouble, however, is that the higher order the method the greater the computational cost per time step.

The reason for this is that higher order methods are constructed by matching more and more terms in the Taylor series expansion of the solution (5.1). In order to do this a method must be devised that expresses higher order terms in the Taylor series in terms of the vector field $f(y)$. The complexity of each term increases with each increase in order, which means that a more complex arithmetical calculation is required and increasingly many evaluations of the function $f(y)$. Accordingly, there is a trade-off between the order of the method and its computational cost. But in any case, except under special circumstances only finitely many of the terms in the Taylor series can be computed, which effectively means that we compute the solution to the differential equation by computing finitely many terms of its Taylor expansion. This truncation of the Taylor series thus introduces error into the calculation, which is not surprisingly called *truncation error*.

The behaviour of the truncation error can be seen to be determined largely by the order of the method. Since Taylor's theorem provides an estimate of the truncation error as a term of one order higher than the highest order term of the truncated series, this means that the higher order the method, the faster the rate at which the truncation error shrinks when the step size is reduced. This is expressed in terms of

the asymptotic convergence of the computed solution to the exact solution as the step size is reduced to zero. For a method of order n, then, the estimate of the truncation error provided by Taylor's theorem shows that the truncation error goes to zero at the rate of h^{n+1}. But this is only the error introduced over a single step, so the asymptotic behaviour of the truncation error for the entire path is not guaranteed to be the same. It can be shown, however, that the error introduced over the entire path goes to zero more slowly, at the rate of h^n (see, *e.g.*, Henrici, 1962, 177). This is the reason for the definition of the order of a numerical method, *viz.*, the order of the method is the order of the asymptotic rate of convergence of the computed path to the exact path as the step size goes to zero.

We may see from this that an infinite order method would provide the exact solution to the original problem, in the sense that all of the points computed by the method would lie on the same solution of the original problem. This shows that the discretization step that renders time discrete is not what is responsible for the error. Truncating the Taylor series, however, does introduce error and is the source of the error introduced by the numerical method, hence why it is called truncation error. The approach of matching terms in the Taylor series does, however, ensure that the error introduced by the method over a time step goes asymptotically to zero at a rate depending on how many terms of the Taylor series are matched. The result of this, however, is that in general a numerical method can only be accurate for a limited period of time before the error grows too large (fixed points and periodic solutions are exceptions to this). Accordingly, numerical methods generally give effectively valid solutions, but they do so only over a limited interval of time before becoming unstable. Thus, just as the trajectories of chaotic problems in the presence of small error can only be tracked for a limited time, even in principle, numerical methods applied to nonlinear, but not necessarily chaotic, problems generally allow trajectories to be tracked only for a limited period of time. Thus, in both cases we are restricted to locally effectively valid solutions.

This understanding of the asymptotic behaviour is important for understanding the effect of reducing the step size to increase accuracy, *viz.*, a linear change in step size has a polynomial effect on accuracy. Accordingly, the error benefits of a reduction in time step can be significant for a method with a reasonably high order. And when

short time steps are not required, the time step can be increased to reduce cost. To accomplish this, a measure of the error in the generated solution is required, and a common method for doing is this is to use a predictor-corrector approach, which computes a time-step of the method in two steps. The first computes a rough approximation to the time step and the second computes a refinement to the rough approximation. A common approach is to use two multi-step methods on the same set of values at previous time-steps, where one method computes a less reliable time-step at lower cost (predictor method) and the other computes a more reliable time-step at higher cost (corrector method). The use of the same prior information lowers the cost per step compared to using a pair of Runge-Kutta methods. The difference between the two values can measure the (local) forward error, which can then be used to control the size of the time step. The version of the symplectic tangent map developed by Mikkola & Innanen (1999), described above, provides another means for tracking (local) forward error, since it allows semi-analytic computation of the behaviour of nearby solutions. As a result, monitoring the rate at which nearby solutions diverge could provide a way of controlling error for symplectic methods in dynamical astronomy.

As useful as knowing the asymptotic behaviour of a numerical method is, in general this just tells us that if we reduce the time step sufficiently, the error in the solution can be kept in check. Whether this can be feasibly done in general depends on the problem. As a result, it is helpful to know more about the properties of the numerical method in relation to error. One way of accomplishing this is to study the behaviour of the numerical method on certain classes of problems, such as dissipative problems of a given sort, which can be very useful. Another approach is to construct the method in such a way that it has nice error behaviour properties. This is the approach that geometric numerical methods (*i.e.*, methods, such as symplectic methods that preserve certain geometric structures of a problem) are useful. Using such a method on a hamiltonian problem, we know that the method, in the absence of rounding error, will rigorously nearly preserve energy. And as we have seen in the case of the symplectic tangent map, such approaches not only guarantee certain nice error behaviour, they can also significantly increase the efficiency.

5.3.2 Floating Point Error and Its Structural Variation

Since the difference equations of a numerical method generally cannot be solved exactly, since the methods we have, and the range of functions we can feasibly write down, are limited, it is necessary to use machine arithmetic to compute solutions to problems. For models involving real numbers, this generally requires the use of floating point arithmetic to solve problems. Floating point numbers are designed to allow a wide range of values and high-precision while only requiring a small number of bits to specify a value. This essentially involves the representation of real numbers in fixed precision scientific notation, *viz.*, as

$$fraction \times base^{exponent},$$

where the $fraction$ is a signed fraction with a certain number of significant digits. The most common implementation of this is a binary representation using 64 bits. 53 bits are used to specify the fraction, one of which specifies the sign of the number, which gives a precision of about 16 significant digits. The remaining 11 bits are used to specify the exponent. The 16 digits of precision that the standard implementation provides shows why the simple example of floating point instability considered above occurs, *viz.*, when you add two numbers that differ in precision by more than 16 significant digits, all the information about the smaller number is lost. Thus, if you subtract the larger number after this addition, you get zero rather than the original small number. This limitation of floating point arithmetic is an example of the effects of *rounding error*. Rounding error results from a computation where the exact value is properly between the two closest floating point values. The given implementation of the floating point system will provide an algorithm to decide how to round the results of computations to the nearest floating point digit. The ability to design algorithms that will be numerically stable when computed using floating point arithmetic requires understanding the properties of the floating point system that is used on the machine that performs the calculations.

The effects of floating point error on hamiltonian problems show that the effectively valid solutions to hamiltonian problems that we can feasibly access are generally exact solutions of non-hamiltonian problems. So, we end up with an exact solution,

but at the expense of having to consider perturbations that take us outside of the theoretical framework within which the model is constructed. If we consider numerical perturbations to model physical perturbations, however, it is natural that the non-hamiltonian problems we can feasibly solve exactly are still valid models of the target phenomenon since physical phenomena do not follow classical hamiltonian dynamics exactly. So we have no reason to expect that models would have to be hamiltonian in order to be descriptively stable. The fact that the problems we can feasibly solve are outside of the theoretical framework in which the model is constructed is problematic, however, for typical logical approaches.

First of all this case provides a very succinct argument that scientific methods are not deductive. Since feasibly computing solutions to hamiltonian problems forces us out of the framework of hamiltonian mechanics, there is no way that scientific methods can possibly be deductive. Someone might argue that it is only the merely pragmatically necessary computation that requires leaving the framework of hamiltonian mechanics, and the result is thus an approximate solution that is very close to the solution of a problem of the theory. First of all, as I have already argued, feasible computation is not merely pragmatically necessary, it is necessary to get *any* solution, and feasible computation typically requires some form of machine computation in contemporary science. And second of all, the result can just as well be considered an exact solution to a non-hamiltonian problem, which should be the default position for a logical view since exact solutions are epistemically on firmer ground than approximate ones. So, from this perspective, obtaining a feasibly computable solution does not force us to accept approximate solutions, but it does force us to accept exact solutions to problems outside of the theoretical framework.

This matter of floating point arithmetic forcing exact solutions to problems outside of the theoretical framework has implications for the stability of physical theories in applications. Since all known theories make abstractions of some kind, to be useful they must be stable under small perturbations out of their theoretical framework. This has two implications for floating point arithmetic. First of all, continuing the theme of allowing computational error to model physical perturbations, floating point perturbations could be construed as modeling a certain kind of physical perturbation that cannot be described within the framework of the theory. Second of all, along

with this, the fact that numerical methods are able to produce accurate solutions for modeling problems despite the effects of rounding error provides evidence, albeit quite limited evidence, that theoretical frameworks are stable under physical perturbations that cannot be modeled in the theory. This shows another way that backward error analysis can convert computational error from an epistemic vice into an epistemic virtue.

It is important to note that there are means of overcoming the computational limitations of floating point arithmetic without analyzing things as we have here. Using techniques of interval arithmetic, it is possible to use a floating point calculation to find an interval around the computed value within which the exact value can be rigorously proved to lie. This approach underlies a number of methods that prove rigorous mathematical results on the basis of machine computations, including shadowing methods for ODE. A famous example of this approach is used in the computationally aided rigorous proof by Tucker (2002) that the Lorenz equations are indeed chaotic for certain parameters and this chaos is stable under small perturbations of the parameters. Up until this proof was devised, the evidence that the Lorenz equations are chaotic was inconclusive, because it was not known whether the chaos observed in numerical computations was introduced spuriously due to rounding errors in floating point arithmetic.

5.3.3 Structural Stability of Computational Inference

Up to this point we have noted certain similarities between feasible data handling in astronomy and feasible computing. Specifically, the manner in which both cases involve a physical device that records information that can be used to find the effective state of a system (model system or target phenomenon) at a time, and the manner in which both use systems of constraints in the design process for the physical device with similar consideration of insulating the "record" from physical perturbations in the environment. The analogy between the two cases, however, can be extended much further than this, as we will now see.

In the case of data handling, the process was broken up into three main frameworks, indicating the most significant transformations required in the process of measurement: the theoretical model; the measurement data model; and the telescope/raw

data model. In the case of feasible computation, the process may also broken up into three main frameworks, also indicating the most significant transformations required: the theoretical model; the "numerical model", *i.e.*, difference equation; and the processor/register data model, *i.e.*, the low level of the physical processor and memory it used to output its results. We have already noted the similarities at the level of the physical device model. The theoretical models actually coincide, in the astronomy case, since the mathematical model of the orbit of an object and the measured orbit of the object are, if not identical, then two different solutions to the same model equation. And the intermediate models are also functionally analogous, since the framework of the numerical method provides a meaningful interpretation of the raw results of the microprocessor computation recorded in a register in terms of a discrete sequence of states in phase space. This is similar to the way that the measurement data model provides a meaningful interpretation of the raw results of the telescope image recorded on a CCD—in both cases the intermediate models serve to interpret accessible raw information (image plates, register states) into a form that is conceptually meaningful in relation to the theoretical model (celestial locations, phase space states).

There are also similar reasons for making the transformations to these different frameworks. In astronomy, we attempt to constrain the orbit of an object using information that we can know for certain. This first transformation requires relating the position of an object to the model of data used to measure its location on the celestial sphere. And the second requires relating a measured location on the celestial sphere to the actual image recorded from a telescope. In the computational case, the transformation to the framework of the numerical method relates the position of a solution in phase space to a model of a phase space position determined by a microprocessor. And the transformation to the framework of the processor register requires relating a phase space position to the actual result of the computation recorded in the register. Both the CCD image and the register state are effectively *positive* information, known with effective certainty. The ability to draw inferences from them then requires that their results can be interpreted back in the data model framework and theoretical framework. Unlike the measurement case, the computer has an exact means of translating the register result into a floating point (vector)

value that can then be interpreted as a point in phase space in the framework of the numerical method. And rather than requiring the solution of a complex problem, the translation from the framework of the numerical method to the theoretical framework involves a relatively straightforward computation of an interpolant for the numerical solution.

Lest it seem like these similarities are somewhat trivial, consider the interpretation of the transformations down toward the device as information transmission models, as we did in the measurement case. It becomes evident that the two transformations (mathematical model to numerical method, numerical method to machine arithmetic) are effected in order to convert the problem into a similar one of lower computational complexity. In order for the resulting computation to be numerically stable, however, the simplified problem must still contain enough information about the dynamical behaviour in phase space in order to permit the computation, locally, of an effectively valid solution to the problem. Thus, numerical stability relies upon the transformations *effectively preserving the information needed to (locally effectively) determine the information about the dynamical behaviour of the model.* Stated more clearly, then, we may see that feasible computation of dynamical systems is the effectively valid transmission of information about dynamical behaviour from a given framework to one that is more easily computable, with the final transformation ensuring feasibility. On the same token, we can see the measurement case transformations as what is required to make the dynamical behaviour of the target phenomenon more easily accessible, with the final transformation ensuring feasibility.

It is then not difficult to see that numerical stability is a kind of descriptive stability in the technical sense of (local effective) soundness and completeness. This is simply because the transformations from the mathematical model to the numerical method and from the numerical method to the device are implementations that determine what states would obtain in the lower framework under given conditions on the original problem. And for the computation to be numerically stable, it is necessary that the state at a given time determined by the mathematical model must be effectively the state determined by the numerical method or machine computation at that time. Thus, one kind of numerical stability, namely, having a small forward error, is equivalent to effective soundness. Given the definition of numerical stability

as mixed stability, however, the results of the computation interpreted back in higher frameworks must result in a small backward error, which we find when we show that the defect is small. Thus, backward numerical stability is equivalent to effective completeness. In which case numerical stability, in the mixed sense, is equivalent to descriptive stability in the technical sense introduced above (see figure 5.2, *cf.*, figure 4.2).

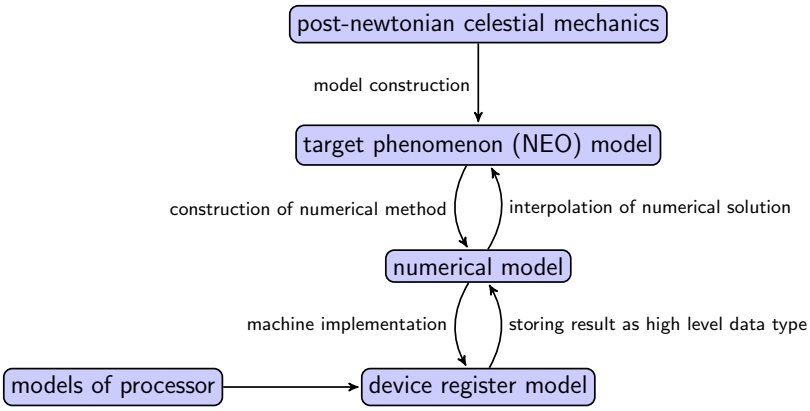

Figure 5.2: A representation of a model of the theory-computing device relation in the case of near-Earth object modeling using post-newtonian celestial mechanics. This model shows the similarity to the process of using data to compute the orbit of a NEO, as shown in figure 4.2, and how numerical stability is a variety of descriptive stability, *i.e.*, (local) effective soundness and completeness.

One interesting feature of this equivalence is that it completes a well-known analogy between numerical analysis and dynamical systems. As was discussed above, the concept of conditioning in numerical analysis is analogous to the concept of dynamical stability in the theory of dynamical systems. The theory of dynamical systems, however, has no equivalent of numerical stability. We may now see that the equivalent only appears when we consider the stable description of dynamical phenomena using dynamical systems, in which case the analogous concept is descriptive stability.[11]

[11]There is also a more subtle connection that comes out here. The notion of descriptive stability could be seen to emerge in the theory of dynamical systems if we consider the stability of characteristics or quantities across nearby systems or trajectories in terms of changes of descriptive framework. This view emerges quite naturally in the local constraint logic we are using in this study. On such a view, descriptive stability obtains provided that the relevant quantity or characteristic is effectively unchanged under the change to the nearby system or trajectory. This shows that descriptive stability is connected to dynamical stability. A similar phenomenon occurs in numerical analysis as is illustrated by the fact that the logistic map becomes chaotic as the time step is increased into the range of values greater than 2. In this case numerical instability is related to ill-conditioning. Thus, because the notion of descriptive stability allows the consideration of the stability of arbitrary characteristics of quantities across changes in descriptive framework, descriptive stability (local effective soundness and completeness) provides a logical means of interpreting the notion of well-enough-conditioning introduced by Corless (1994*a*).

Before concluding our discussion of the case examples, there is one last point about numerical stability in feasible computing that has not come out of the above analysis. A full analysis of the information transmission required in computing estimations of impact probabilities demands that we consider how information is transmitted between observation stations and a central computing station, a process that involves sending data over the internet. Because the accuracy of data transmission over the internet is ensured by error control technologies, the concern here is time lag over long distances or slow connections. Since the JPL ephemerides are important for many solar system dynamics problems, and the ephemerides made available online are very large, there is a virtue in being able to compute a high accuracy ephemeris for oneself for the purposes of computing of trajectories of solar system objects. For this reason Sitarski (2002) developed a version of the most recent version of the JPL empherides that does not require an internet connection, by preparing special initial data for integrating the orbits of the planets in order to obtain accuracy comparable to the JPL planetary ephemeris. Interestingly, the JPL ephemeridies are used as measurement data to test the accuracy of Sitarski's method. Considering the seriousness of the consequences of an object potentially colliding with Earth, it is wise to have an accurate and efficient means of computing orbits that does not rely on an active internet connection. This shows not only the very different kinds of stability questions that arise in feasible epistemology, it also shows yet another reason why tracking information transmission is central to understanding descriptive stability.

Part III

Putting It All Together

Chapter 6

Stably Capturing Variation: Feasible Local Models of Feasible Theory Application

6.1 From Classical To Feasible Epistemology

6.1.1 Limitations of Classical Epistemology

I have argued that classical epistemology of science, reflected in its reliance on classical mathematics and metamathematics in the formulation and investigation of epistemological theories, is unable to give adequate insight into the complexity of method in science and unable to account for the reliability and robustness of those methods in generating knowledge about phenomena in real epistemic contexts. One of the central lines of evidence for this has been how the classical picture of the theory-world relation in terms of a "derivation from theory" fails to account for how it is actually possible for us to connect our theories up to the world in anything other than quite highly idealized circumstances. This evidence came in three main forms, which we can understand in connection with the classical picture of a derivation of laws from theory using valid rules of inference and the further derivation of a description, prediction or explanation from laws and conditions.

The first is the role played by approximation throughout the process of applying theories to the world, which pertains to the "rules of inference" used in theory application. Approximation methods, including systematic model reduction methods, are essential for determining what the consequences of a theory are in real scientific cases.

This is a general fact about modeling practice, since such methods are important not just in the construction of the primary model, because many of the vast number of models involved in data handling also rely on similar forms of approximation. Furthermore, in increasingly many contemporary cases, computation must be used to find approximate solutions to the equations of motion, themselves being the result of approximation. Thus, the path from theory to laws given in the form of equations of motion and then to descriptions relies heavily on "rules of inference" that are not deductive.

The second form of evidence is the role played by data in the process of scientific inference. In part, this pertains to the "derivation" of the equations of motion, since we have seen the importance of symmetry assumptions and behavioural assumptions, such as constitutive relations and equations of state, in model construction. And, as I have emphasized, this is not a "merely pragmatic" issue, because this kind of input is necessary to feasibly construct a model and theories themselves often have little or no distinct empirical content. Indeed, as Smith (2002*b*) has pointed out, because many theories in science require constitutive relations to describe the behaviour of any system, such theories are simply too abstract to describe any physical system. Thus, constitutive relations and equations of state are *adding content* to the theory, which means that equations of motion or modeling frameworks are not "derived" from theory by logical deduction using valid rules of inference.

The role of data in inference also pertains to the specification of "conditions" on the laws. This can be seen from the manner in which data is entangled with the inferential process in astronomy. For example, one aspect of this entanglement can be seen in the fact that the specification of the initial condition of a NEO's orbit requires the solution of a complex problem that is heavily reliant on data—the orbit determination problem. Solving this problem is a substantive part of the process of estimating an impact probability, so this example shows that a significant amount of data and modeling can go into the specification of the conditions on a law. Another aspect pertains to the "derivation" of the description from "laws" and "conditions". Aside from the fact that computation is necessary to solve the orbit determination problem, the importance of the measurement reduction function in the NEO problem shows that data modeling is a key part of the "derivation" of the

orbit. No matter how someone might try to massage the classical picture to account for these cases, it remains the case that all of these facts, central to drawing reliable inferences in astronomy, are not accounted for in a logical derivation model of theory application. And even if the "derivation" model could account for all this, it still does not represent the epistemic importance of the fact that real scientific inference can require an iterative correction process between theory and data, as is required to obtain an orbit precise enough to rule out a collision with a NEO that will make a close approach to Earth, or to determine that one will occur.

The third form of evidence comes from the role played by computation in real scientific inference. This pertains in part to the specification of "conditions", since computation is required to solve the orbit determination problem in order to specify the initial condition for the impact probability problem. But, more significantly, this pertains to the "derivation" of descriptions, predictions or explanations from the "laws" and "conditions". In increasingly many contemporary cases in science, theory cannot be applied to a phenomenon of interest without the use of computing devices, meaning that, in many cases, we cannot actually determine what the consequences of theory are without computation. This shows that, far from being "merely pragmatic", computation actually extends our epistemic capacities, allowing application of theory in cases where it would not be possible otherwise. In addition, real computing systems are discrete and bounded, making them incompatible with the classical mathematical frameworks within which most theories are formulated. Since determining consequences of theory often requires this radical semantic transformation of the spaces of the model, the "derivation" of descriptions in science is often not deductive. And when we understand that a computer solution is an exact solution to a nearby problem, we see that, where machine arithmetic cannot be made to respect the geometric structure of a theory, the only solutions we can actually obtain are solutions to models outside of the mathematical framework of the theory.

These lines of evidence show decidedly that the classical model of the theory-world relation in terms of descriptions, predictions and explanations being deduced from theory by valid logical laws, which are then compared with data, is invalid as a representation of how we actually gain knowledge about the world. If such a model is to be regarded as a methodology avoidance strategy and not as a statement about

actual scientific practice, then it is not necessarily invalid, but then it is unclear what sort of insight it is supposed to give us into scientific inference in real science. Consequently, it is unclear what value such a general theory of the theory-world relation can be, when it cannot account for actual scientific method. Accordingly, I argued that we are in need of new systematic tools that are capable of providing faithful descriptions of real scientific method and knowledge.

Another issue with this picture, which has been mostly an implicit part of our considerations up to this point, concerns its explanation of the reliability of scientific inference. The reliability of scientific inference in this picture is supposed to come from the fact that the laws and conditions are true and the rules of inference used to derive descriptions, predictions and explanations are valid logical laws, meaning what is described, predicted or explained is also true. This makes falsification of a prediction from theory a catastrophic event for a theory, not only because it contradicts the claim made by theory but also because it is the truth of the laws that was explaining the reliability of the theoretical descriptions. The importance of this trivial fact that a (logical) theory can be undermined by a single recalcitrant observation, *viz.*, one that cannot be accounted for by adjusting anything other than the theory, is that the explanation of the reliability of a theory is unstable—a single recalcitrant observation undermines the explanation. Of course, the natural move is from strict truth to approximate truth.

Approximate truth provides a much more robust explanation of the reliability of inference because it allows for reliable descriptions in the presence of error. And error is not only introduced in applications, it is intrinsic to our scientific theories. Although some hypothetical exact theory of everything could in principle be exactly true, a theory with a limited domain of applicability could not. Since all of our actual theories have limited domains of applicability, it is natural to explain reliability in terms of approximate truth, or some weaker notion of similarity to the phenomena that Giere (2010*a*) advocates. A theory is problematic then, when its actual domain of applicability does not cover the range of phenomena over which the theory is intended to apply. This is a satisfactory way of accounting for the reliability of theoretical descriptions and the features that make theories unreliable, but aside from the problematic manner of giving a satisfactory account of approximate truth

or pinning down a notion of similarity, there are still significant problems with simply shifting to approximate truth or similarity in order to obtain a general account of the reliability of scientific inference. However, what is entirely missing from an account of this sort, an element that reflects the classical nature of these accounts, is the consideration of what makes the actual methods used in real science reliable.

This is a question that is difficult to address from the point of view of a logical derivation model. There is a natural classical response, however from the point of view of a semantic view that takes into account some well-known facts about differential equations. The story from this point of view is that real methods just provide useful ways of finding descriptions of phenomena that are approximations of what the theory actually predicts, accurate to within experimental error. And since for theoretical descriptions to be applicable the equations have to be well-posed (*i.e.*, have unique solutions depending continuously on the data), the reliability of real methods results from the theoretical descriptions being stable under error in the data coming from measurement and approximation. This is certainly an important part of the explanation of the reliability of scientific inference, but there are reasons why it is problematic that are demonstrated clearly in the case examples considered in this study.

The first problem with this explanation is that it is too simplistic. A clear way of demonstrating this comes from considering the theoretical study of chaotic phenomena. As we have seen, trajectories for chaotic systems can be accurate within the Lyapunov time of the system. This means that for such systems the theoretical description can only be accurate over a limited period of time. Accounting for our descriptions being limited in this way is not a problem for a classical account, however, because the true theoretical account can still be exact, but the descriptions we can obtain in practice contain error, which means that they can only be accurate locally to a range of time. What makes the methods of generating descriptions reliable, then, is that the exact theoretical descriptions are (approximately) true, or at least yield results that are similar in behaviour to the phenomena. Given that practical methods of computing solutions introduce error means that the trajectories we can compute will diverge from the classically exact solution after the Lyapunov time.

So far so good, but the successful use of the ideal double pendulum model under

chaotic conditions is more difficult to explain. In this case, the nonlinearity is more severe and the model can no longer accurately track the phase trajectory of the real system, which is the typical case for chaotic phenomena. In this case the description breaks down, which would make it useless on a simple classical account. This does not prevent the model from continuing to be useful, however, since it can still provide valuable information, such as the conditions (*i.e.*, regions of phase space) under which the pendulum becomes chaotic, and be useful by providing an estimate of the Lyapunov exponent. This is explained as a result of the fact that *certain* quantities are stable under variation of the model even though the trajectory is not. So a model of a chaotic system continues to be useful provided that we are interested in gaining information about quantities that are stable under shifts to nearby models. What clearly follows in this case, then, is that the model is only *partially* descriptive of the real system, so it is not appropriate to take it literally. This is troublesome for an account that appeals to approximate truth to explain reliability. But partial descriptiveness is all we require in this case, since we cannot expect to accurately track the system's behaviour over time anyway. Indeed, this is also shown in the case of orbit models that continue to provide useful information about the orbit beyond the Lyapunov time when the orbit itself is stable, even though the predicted location on the orbit is not. In this case the model is also only partially descriptive because it no longer accurately tracks the system's trajectory in phase space, even though it can accurately place the orbit in a certain region of physical/configuration space. And, much like the double pendulum case, the model in such a case is useful for the determination of quantities provided that they are stable across nearby orbits.

Accounts that explain reliability in terms of similarity to the phenomena, however, *e.g.*, Giere (2010*a*), are perfectly compatible with partial descriptiveness. But the explanation is still appealing to well-posedness (existence of a unique solution that varies continuously under changes of the data), since this is what is accounting for how the introduction of error leads to approximations. But is too simplistic as an explanation because reliability depends not only on continuity under variations but also on how significant the changes are for a given size of variation. Stated more simply, it is the *rate of change*, not only continuity, that matters for reliability. Thus, the *conditioning* of the problem is also important. To a certain extent this is actually

accounted for by the natural classical story in terms of well-posedness. Proofs of well-posedness for differential equations often provide a measure of how large a change in the solution results from a given change in the data. But, in the case we are considering here, the models are strictly speaking ill-conditioned because they are chaotic. Nevertheless, the model we construct is still useful because *certain* quantities are well-conditioned even though the trajectory is not. Those quantities of a chaotic system that are well-conditioned are the ones that we can use to obtain reliable information from models we can construct. Thus, models of chaotic systems are useful because they are *partially*-well-conditioned. This behaviour was pointed out by Corless (1994*a*), who called problems that have this partial-well-conditioning property *well-enough-conditioned*. This shows clearly that in general it is well-conditioning, not simply well-posedness, that accounts for the reliability of differential equation models. But it also shows that what matters even more when it comes to nonlinear systems is the more nuanced idea of partial-well-conditioning, or well-enough-conditioning.

This explanation of reliability in terms of partial-well-conditioning could be seen as an elaboration of the above semantic account of reliability. There is a reason, however, why incorporating partial-well-conditioning into the explanation of the reliability of theoretical descriptions pushes beyond a typical sort of classical account. The justification for the classical account above relies on general proofs of well-posedness for general classes of differential equations. And, as mentioned, these proofs often also provide a measure of the conditioning of the problem, in the ordinary sense of conditioning. These general proofs contain no information, however, about when systems have this partial-well-conditioning property. This requires proofs that are more specific to special classes of problem or particular problems, in particular because it is not clear *a priori* which quantities will be stable for a given chaotic dynamical system. So even though it may be a partial-well-conditioning property that is explaining the reliability of a theoretical description, it is often difficult, impractical or impossible to give a classical justification for this explanation. Indeed, the justification typically relies on numerical solutions, which give not proofs but *evidence* that models are partially-well-conditioned. The evidence comes from the fact that assuming that models are partially-well-conditioned in a required sense often produces successful results. So, *the usefulness of computer solutions in practice provides evidence of sta-*

bility properties that can explain the reliability of theoretical descriptions. Thus, the justification for a classical account is now seen to rely in many practical cases on methods of feasible computation. This is the first clear sign that there is a great deal more to the reliability of scientific inference than the stability properties of differential equations.

One way to motivate the need to look beyond the framework of a given theory for explanations of reliability comes from cases where there is no way to regard computational solutions as exact solutions to problems from the framework of the theory, as is the case in general for hamiltonian problems. On the semantic view, then, the only problems we can actually solve are ones that are outside the set of $\mathfrak{models}\ \mathfrak{T}$ of the theory. Of course, a simple response is that the results of computation are reliable because they are good approximations to what the theory actually predicts. An appeal of this sort to approximation or approximate truth is missing a more significant point. The reliability of solutions to models outside of the theoretical framework shows that the theoretical framework itself is stable under variations, *viz.*, small variations outside of the framework of the theory still produce accurate models. This is evidence for a much stronger robustness property than usual appeals to well-posedness or well-conditioning provide. Indeed, it is evidence for the *well-conditioning of theoretical frameworks* under small perturbations that are not modeled within the theoretical framework. This is the interpretation we reach when we recognize that it is generally not possible to feasibly obtain an exact solution to a hamiltonian model but that we can and do obtain an exact solution to a nearby nearly hamiltonian problem—an exact solution to a non-hamiltonian model that is nevertheless reliable. This shows that the tendency to focus on approximation, approximate truth, or even similarity to phenomena as an explanation of reliability misses important features of what explains the reliability of scientific methods.

Since the error implied by approximate truth corresponds to forward error in numerical analysis, and the error in the data associated with well-posedness and conditioning corresponds to backward error, the evidence numerical solutions provide that theoretical frameworks are themselves well-conditioned shows that classsical explanations of the reliability of scientific methods have not paid enough attention to the analysis of backward error. Moreover, classical accounts do not consider the

importance of the evidence from numerical solutions for stability properties that we cannot prove analytically (partial-well-conditioning properties, stability of theoretical frameworks under perturbations). The consideration of the importance of conditioning in explaining the reliability of scientific methods shows that what is explaining the reliability of methods is the ability of a theory to support descriptions that are approximately true or sufficiently similar to the phenomena, but *also* the fact that the descriptive models are, in general, partially-well-conditioned. But there is more to the overall picture of reliable scientific methods that is missing from this sort of explanation of their reliability: we need not only to explain why theories are reliable sources of information about the world, but we also need to explain why our methods of getting information out of theories are reliable.

This takes us to a way in which the classical semantic account of the reliability of scientific methods is actually incomplete. The example of reliable numerical solutions shows not only evidence for well-conditioning properties of methods, it also shows the reliability of the combination of numerical methods and machine computation to generate accurate solutions of equations. So, an essential part of the explanation of the reliability of scientific descriptions is the explanation of why numerical computation is reliable. Since we have seen that the difference equations of numerical methods are actually kinds of continuous mathematical models (continuous system space), it is not surprising that the methods used to explain their reliability are similar to those used to explain the reliability of differential equations. The more complex problem of how to explain the reliability of scientific methods, however, comes from the explanation of the reliability of machine computation. Given the considerations of the last chapter, we have seen that this involves explaining why the reduction of a theoretical description into a form that can be computed by a microprocessor preserves enough information about the solutions of the differential equation to get accurate results. Since such an explanation must be given in terms of the methods used to program a computer and those methods' relation to physical facts about the computer, the explanation of the reliability of scientific methods takes us out of the bounds of classical epistemology and into the realm of feasible epistemology.

Classical epistemology of science tends to explain the reliability of methods almost, or entirely, exclusively in terms of theories, so it fails to explain how we reliably

get information out of theories. This requires explaining why the methods used to construct models are reliable, why the methods used to model, gather and process data are reliable, and why the methods used in numerical computation are reliable. So all of the questions about reliability that are external to considerations of theoretical frameworks, the questions about how we *actually* are able to get reliable information about the behaviour of phenomena from a theory when the full theoretical descriptions are only accessible in the simplest of circumstances, are ignored. It is these questions that are specifically the focus of feasible epistemology. Since the classical explanation in terms of approximate truth and well-conditioning are still an important part of the explanation of the reliability of scientific methods, classical and feasible epistemology are not actually at odds with one another in this context, they are complementary. Classical epistemology gives us insight into what features of scientific theories explain their reliability and feasible epistemology gives us insight into those features of the methods we have available to us that explain why they are able to reliably extract accurate information from theories.

6.1.2 Limitations of Classical Methods

Now let us turn to a different matter, namely the "meta-methodological" question of what methods we are to use in scientific epistemology to study feasibility in science. In terms of classical versus feasible approaches, this is not a question of whether classical accounts are able to capture the structure and behaviour of scientific methods and knowledge in principle, but what they are capable of elucidating as tools for the study of feasible scientific methods. There is, of course, no unique answer to the question of what methods are optimal for studying feasibility, but the concern here is to identify methods that are suited to the detection of relatively general structure and behaviour of method in science. One of the main problems here is that the closer one gets to what actually makes scientific inference feasible, the more and more complex the methods used and the details explaining their reliability actually become. So, what might be a trivially simple matter in principle, such as evaluating a function over a limited range of values, becomes more complex and non-trivial in practice and forces one to consider issues such as numerical stability and conditioning. Feasible epistemology needs methods that are able to cut through the complexity of scientific method in

order to focus in on what is important in order to address or answer epistemological questions about science. And I have argued that typical approaches of the syntactic and semantic views of theories are ill suited to this task because they are geared toward abstraction not investigation.

The syntactic and semantic approaches to theories are generally geared toward the most general descriptions and accounts of scientific method possible. This is quite explicit in the case of the received view since it claimed to be a fully general account of theories as axiomatic systems that acquired meaning through the partial interpretation of theoretical language in experience. And an important feature of the account was that by specifying what gives theories cognitive significance by explaining how they are given content by interpretation in experience, a solution could be given to the demarcation problem. In the case of the semantic approaches, they have been less ambitious, focusing only on giving an account of scientific theories and their relation to the world in terms of some mapping relationship between the theory and certain phenomena. The main semantic approaches, *viz.*, the set-theoretic predicate approach of Suppes, the state-space approach of Beth and van Fraassen, and the relational systems approach of Suppe, all share the property that they attempt to provide a general account of theories in terms of how their formulations are given a semantic interpretation in the universe of sets. Indeed, Suppe (1989) points out that they also share the property that no matter what type of mathematical entity theories are identified with, they all function to specify the admissible behaviour of state transition systems (*op. cit.*, 4).

But there is more to saying that the logical approaches are geared toward abstraction not investigation than saying simply that they seek fully general accounts of scientific theories. Some would point to Suppes's work, such as his work with McKinsey *et al.* (1953) to axiomatize simple point-particle newtonian mechanics or his work on the foundations of the theory of measurement, *e.g.*, (Scott & Suppes, 1958) and the series of *Foundations of Measurement* volumes, (Krantz *et al.*, 1971; Suppes *et al.*, 1989; Luce *et al.*, 1990), as clear evidence of the usefulness and fecundity of the use of a semantic approach in the investigation of science. And some would point to the general defence Suppe (1989) provides for his semantic conception, by showing how it clarifies issues in the foundations of biological taxonomy, biomedical

sciences and even cases of public policy disputes, as evidence of the same. Independently of how successful these efforts are, however, these efforts show precisely the manner in which the logical approaches are geared toward abstraction. In each of these cases the approach is to formulate a general account and then argue for it by showing how it applies to specific cases and clarifies philosophical issues. The problem with this as a reconstruction of epistemology of science, however, is that the general account is formulated *a priori* according to certain preconceived notions of the structure of theories and their mapping relations to phenomena. The formulation does not result from careful and detailed investigation of method and knowledge in real scientific cases in order to generalize to an account that is known to be consistent with those details; it results from some limited level of engagement with scientific practice followed by an abstraction of a fully general account. What follows then is rarely a rigorous checking of the account for consistency with the details of scientific method and knowledge in a variety of cases, which is necessary to ensure that the account applies to scientific practice. What usually follows is an articulation of a general view of scientific theories, or some aspect of them, from the point of view of that formulation, leaving it to other philosophers to point out limitations of the account.

In this way, then, the general semantic approaches are not scientific in the very basic sense of their claims not being tested against the phenomenon, which in this case is actual scientific method and knowledge. The approaches are not suited to the investigation of scientific method and knowledge in a way that leads to the discovery of new knowledge about scientific methodology and epistemology. They are suited to an abstract representation of epistemology of science that is not guaranteed with any kind of reliability to make solid contact with actual science in any given case. A good example of this is the ease with which the details of scientific practice are dismissed by claiming that is is possible "in principle" to incorporate the details. This results in an implicit or explicit belief that the details of scientific practice are mostly irrelevant to epistemology of science, or at least do not need to be considered carefully for the purposes of building epistemological theories. Given that these approaches are supposed to be theories of epistemology of science, it is a curious matter indeed that such theories should be developed independently of the facts. Of course, there is also

the belief that epistemological theories of this sort are dealing with issues that are entirely independent of fine details, since they are supposed to be general theories. The issue with this, however, is that there are lots of ways in which general theories can be independent of fine detail, but that they are independent of fine detail does not guarantee that they are consistent with it. The only way to know that the theories are consistent with the facts is to check, and, as the arguments of the last subsection show, the picture of the theory-world relation in terms of a deduction from theory and conditions to a description that is compared with data is not consistent with how theories are actually applied in applied mathematics.

I believe that there is a deeper reason, however, why the logical approaches are inclined toward regarding most of the details of how knowledge is actually gained as irrelevant to epistemology of science. This is because the scientific methods that they are modeled upon, namely pure mathematics and metamathematics, are methods that do not require engagement with facts. Theories in pure mathematics must be logically consistent (to ensure that they have models), but this is where the fundamental consistency constraints end. Pure mathematicians are free, in principle, to formulate any theories they like provided that they are logically consistent and bear some interesting relation to other pure mathematical theories. It is applied mathematicians that are forced to guarantee not only logical consistency but also consistency with the phenomena their theories are intended to describe. Unlike pure mathematicians, applied mathematicians are not free to formulate any logically consistent theories they like, not even in principle, since consistency with the facts is essential.

Additionally, the methods of applied mathematics reflect the need for consistency with the facts, since they are designed to generate descriptions of phenomena in such a way that the reasons underlying the dominant behaviour of phenomena become clear, yet the descriptions are still sufficiently accurate and can be generated at a low computational cost. That the methods of applied mathematics are able to do this so reliably reflects the fact that the methods were developed in order to meet constraints of consistency, accuracy, low cost and insight. Although principled mathematical arguments are important, it is methods of this sort that actually allow us to gain knowledge of the structure and behaviour of natural phenomena.

Therefore, since we are interested here in what methods are best to use to study

feasibility in epistemology of science, we should look to applied, not pure, mathematics as the paradigm example, since it exemplifies our most sophisticated and reliable methods of feasibility gaining knowledge about natural phenomena. The phenomenon in the case of epistemology of science is the structure and behaviour of scientific method and knowledge. Accordingly, the methods used to study this phenomenon should allow us to feasibly gain insight into the structure and behaviour of actual science. Thus, we may look for ways to design methods that meet the constraints of consistency, accuracy, low cost and insight.

First of all, to meet the consistency requirement we need to ensure logical consistency. One way of ensuring this is to use a "semantic" modeling approach, where by 'semantic' I mean that any epistemological "*theory*" would have the descriptions developed by the modeling approach as its "𝔪𝔬𝔡𝔢𝔩𝔰". This, of course, is the approach that the semantic view uses, since all of its descriptions are based in set theoretic structures. But to meet the requirement of consistency with actual scientific method we need something more. First of all, as I have argued, we need an approach that represents both the syntax and the semantics of scientific theories and how the two covary. This can be done by using a logico-structural or algebro-geometric approach in a manner I have outlined in chapter 2 and used as the basis for the epistemological model constructed through this study. The other key thing that consistency with actual scientific method requires is that epistemological descriptions are ensured to capture the structure or behaviour of method or knowledge they are supposed to.

This relates to the next constraint of accuracy. There is no unique way to approach this, but the strategy argued for and adopted in this study is to consider the subject matter in full detail and then abstract away from the details that are seen to not be relevant to the aims of the epistemological investigation. This approach ensures accuracy because one effectively starts with a fully descriptive account and then generates a more abstract one that is consistent with it by removing detail that is *recognized* to not be relevant in the given case. This approach also helps to meet the constraint of insight because abstracting away irrelevant detail helps to focus in on important features concerning the structure and behaviour of method and knowledge in the given case. Indeed, it is not a surprise that this approach accomplishes consistency, accuracy and insight in this way, because this approach is modeled on

the kinds of model reduction and physics avoidance strategies used in applied mathematics to convert a complex, detailed description into a simpler, more manageable one by abstracting out irrelevant detail in a way that focuses in on dominant behaviour. This is why I referred to this approach as a methodology or epistemology avoidance strategy, since it too abstracts out irrelevant detail in a way that focuses in on dominant epistemological behaviour.

The last constraint concerns computational cost. This is the trickiest constraint to meet and the one that pertains the most to the feasibility of an epistemological method. I have argued that the vastness and complexity of method in science implies that no *a priori* account will be able to track its structure and behaviour adequately. And purely informal investigations are important and useful, but to be able to develop systematic methods to gain insight into scientific method and knowledge we require a stable means of abstraction from this complexity in a way that remains consistent with the facts. This demands a modeling approach that has some kind of formal grounding in order to guarantee precision of language and meaning. But the complexity of science also means that it will not be clear *a priori* what kind of modeling approach will be needed to bring out the dominant structure or behaviour in any given case. Thus, a feasible means of generating epistemological knowledge requires a flexibility that will allow the construction of an appropriate and effective model in a given case.

The *a priori* character of the logical approaches reveals an important reason why its methods are not suited to feasible epistemology. As I have argued, this causes problems for both accuracy and even consistency with actual scientific method and knowledge. They do give insight into the structure and behaviour of scientific method and knowledge, but they do so in a way that is resistant to the growth of knowledge about epistemology of science, forcing it, in all its complexity, through the lens of a particular semantic or syntactic account.

There is another issue with logical approaches, however, which concerns computational cost. If we were to regard logical approaches as methods to "converge on the truth of scientific method and knowledge", then they are incredibly slow to converge. They tend to start with an "initial guess" in the form of a particular semantic or syntactic account, and then undergo an iterative correction process through a combination of investigation from the point of view of the account and criticism from other

researchers. This results in an abstract general account creeping toward consistency with the relevant details of scientific practice, in the expectation of ultimate convergence on the correct general account. Setting aside the unlikelihood of convergence given the complexity of actual scientific method, the steps of this iterative correction process can take an extraordinarily long time, which shows the infeasibility of a general epistemology of science approached in this way. This shows how the logical approaches approach the problem from the wrong direction—they start from general descriptions and move toward the phenomenon rather than starting from the phenomenon and moving toward general descriptions. The latter approach is superior to the former in terms of consistency, accuracy, computational cost, and epistemological insight.

The quasi-formal algebro-geometric approach developed through the process of this study is a first attempt at developing flexible modeling methods suited for feasible epistemology that meet the four constraints of consistency, accuracy, low cost and insight in the manner just described. In order to accomplish this, the guarantees of consistency and accuracy had to be quite local, since I am only considering two examples of theory application in detail. The approach I adopted is grounded as a syntacto-semantic approach developed to describe method in the limited domain of finite dimensional ordinary differential equations and applied to phenomena using newtontian and hamiltonian mechanics. This involved treating equations as constraints on a variable interpretation in a fixed framework and treating descriptions as (partial) translations (implementations) of constraint systems in one framework to another.

It became evident through the process of development that a considerably clearer presentation of structural and behavioural features of theory application in applied mathematics is made using a high (abstraction) level conceptual language based on the framework of algebro-geometric constraint systems. This is a local constraint logic that allows important features of the application process to be captured in terms of local generalizations of the familiar concepts of truth, validity, soundness and completeness from mathematical logic. The second part of this study involved motivating the application of these concepts to the clarification of structural and behavioural patterns throughout the application process. The following section will draw together

the results of our use of these concepts throughout this study to clarify structural features of the application process in applied mathematics. This will provide evidence of the feasibility and fruitfulness of this approach to elucidating dominant structural and behavioural features of method and knowledge in real science.

6.2 Feasibly Modeling Feasible Theory Application

6.2.1 Feasible Inference Using Data, Computing and Modeling

In the previous chapter I showed how there is a distinct structural similarity between the process of using data to draw inferences about a dynamical phenomenon in astronomy and the process of using computation to draw inferences about a dynamical system, as depicted in figure 6.1. The inital step of constructing the model of the target phenomenon, an NEO, is identical in the two cases, since the same modeling framework is used in both cases. From that point on the inferential tasks are quite different in nature, but share a very similar structure.

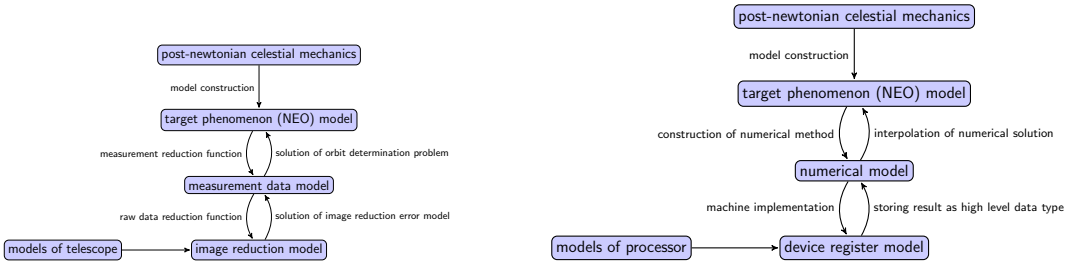

Figure 6.1: Structural similarity of the process of using measurement data collected from a telescope to compute an orbit of a solar system object and the process of using a numerical method implemented on a microprocessor to compute a solution of an initial value problem.

In the case of the problem of determining a high accuracy orbit for a NEO, the inferential task is to use measurement data collected from the object over a long period of time to determine its orbit to within a very small uncertainty. This effectively involves the measurement of a solution in the modeling framework that accurately tracks the orbit of the NEO within its Lyapunov time. In order to accomplish this inferential task it is necessary to model the transmission of information from the object to data in increasingly more feasibly accessible modeling frameworks until one in which recordable observational data is reached, an image on a CCD in the case we

considered, and then "back-interpreting" by interpreting the feasible data through a data model into the modeling framework for the NEO. We saw that the steps of increasing feasibility ("downward" in figure 6.1) required very careful modeling of how light information is transmitted from the object through the telescope to the CCD, carefully correcting for the effects of using non-inertial frames of reference. This involves representing enough of the detail of the information transfer process from the object to the light processed by the telescope that the back-interpreted results provide a highly accurate measurement of the orbit of the object. The back-interpretation steps then involved solving an overdetermined mathematical problem using a least squares method to accurately measure the structured data that was recored and to accurately "measure" the initial state of the object that the data picks out in the modeling framework. At this point, then, methods can be applied to compute an estimate of the impact probability using the computed high precision initial condition.

In the case of the problem of finding a high accuracy solution to an initial value problem, the inferential task is to use a discrete solution to a numerical method recorded from a machine computation to find a solution with a very small forward error. This effectively involves the computation of a solution in the modeling framework that accurately tracks the solution to the initial value problem within the Lyapunov time. In order to accomplish this task it is necessary to construct a dynamical system model in increasingly more feasibly computable modeling frameworks until one that feasibly allows an accurate computation of the solution trajectory, which is the framework of a microprocessor register in the case we considered, and then back-interpreting the feasible solution into the modeling framework of the initial value problem, which passes through the difference equation framework of a numerical method in the case we considered. We saw that the steps of increasing feasibility ("downward" in figure 6.1) required careful modeling of the error introduced in the transformation to a dynamical system that can be computed by a microprocessor and the results stored in a device register. This involves, in a manner similar to compression algorithms in computer science, preserving enough of the information about the dynamical behaviour of the initial value problem in the transformation to the dynamical system computed by the microprocessor that the back-interpreted results provide a highly accurate solu-

tion to the initial value problem. The back-interpretation steps involved structuring the raw result stored in a register into an accurate solution to the difference equations of the numerical method and then interpolating the numerical trajectory to find an accurate solution to the initial value problem.

In both cases we see the importance of accurate preservation of information about behaviour in the shift to more feasibly accessible frameworks. In the computation case it is accurate preservation of the dynamical behaviour of trajectories in the target phenomenon modeling framework and in the data handling case it is accurate preservation of the information about the dynamical behaviour of the trajectory of the object in physical space. And in both cases we see the importance of being able to accurately back interpret the results of feasible information processing in the original modeling framework, a process that requires solving a data constrained problem of the state or path in the next less feasibly accessible framework. This is clear in the data handling case, but back-interpretation in the computation case also involves solving a problem of determining which numerical path is picked out by the result on the register and the solution of an interpolation problem to find an accurate continuous solution picked out by the numerical path. The data constraining the problem is the register result of the computation in the former case and the discrete phase space trajectory in the latter. Accordingly, we can see both cases as a simple recursive method involving a sequence of faithful transformations of the problem from a less feasible framework to a more feasible framework until accessible data can be found to solve the original problem by using the data to repeatedly compute the solution in the next less feasible framework (in the initial sequence of transformations) until a solution to the problem in the original framework is found.

I pointed out in the last chapter that the reliability of the recursive processes in both these cases can be understood in terms of the concept of descriptive stability. And since the reliability of the process in the computation case is explained in terms of numerical stability in numerical analysis, this shows that numerical stability is a special case of the concept of descriptive stability. From the more detailed considerations we just made, then, it may be seen that the "downward" steps involve the accurate specification of a data constrained problem in a more feasible framework. Then, the "upward" steps (in figure 6.1) involve the accurate solution of a

data-constrained problem in a less feasible framework. Therefore, we see that the concept of soundness of local constraint logic is associated with the specification of a data-constrained problem and the concept of completeness with the solution of a data-constrained problem. The feasibility of this recursive solution method, then, requires the ability to access data that can initiate the solution cascade to solve the original problem. This shows that feasible inference in science relies on a low cost recursive method to solve modeling problems that terminates in accessible data followed by a solution cascade that provides an accurate conclusion to the inference, *i.e.*, solution to the modeling problem.

The *explanation* of descriptive stability in a given case is a more difficult problem that comes from an understanding of how error is introduced in the transformation steps required to feasibly find an accurate solution to the original problem. The detailed consideration of the feasibility of these explanations is beyond the scope of the current study. It is clear, however, that descriptive stability is explained in terms of effects of the error introduced in each transformation of the original problem required to find feasible solutions to model problems. For instance, ensuring that the two "downward" transformations in the data handling case result in the specification of a problem that carries enough information about the original problem relies on the sufficient accuracy of the network of models, considered in chapter 4, used to model the transformation of information from the object to a telescope. And ensuring that the two "downward" transformations in the computation case result in a problem that carries enough information about the original problem relies on the numerical stability of algorithms despite the introduction of truncation (first step down) and rounding (second step down) error. Since the "upward" transformations in the data handling case involve the solution of mathematical problems, we may see that the conditioning of the problem together with the stability of the method used to solve it will be central features of concern for understanding why the two "upward" transformations result in accurate solutions to the problems that could not be solved directly. Since interpolation requires the solution of a mathematical problem, a similar thing is true for the computing case also.

One interesting structural feature to note from figure 6.1 is that the solution of the orbit determination problem, itself an inference of the data type, actually

involves a feasible inference of the computational type. And the whole process of solving the orbit determination problem is embedded in the process of specification of the impact probability determination problem. Thus, we see that this pattern of feasible inference in terms of descriptively stable transformations from less feasible to more feasible frameworks occurs on different scales of inferential processes in theory application. What is interesting about this is that it suggests that such procedures are a quite general part of feasible inference in scientific practice. Accordingly, the problem of revealing and elucidating general patterns in what ensures descriptive stability could be a quite important one in feasible epistemology.

Interestingly, when we consider the process of model construction in terms of descriptively stable implementations we find once again a similar structure, but we also find important differences (see figure 6.2). In a similar manner to data and computing, the "downward" transformations effect information preserving implementations of a less feasible framework in a more feasible one. The transformation of the target phenomenon, a near-Earth object in this case, into the theoretical framework is an abstraction of those features of the phenomenon that can be accurately represented by the theory at a given physical scale in the given theoretical framework. This picks out the classical base model, the most detailed model of the phenomenon at a given scale. Since this model is not feasibly accessible, in order to access a feasible base model, this case required imposing reasonable structural and behavioural constraints for the NEO. This resulted in a feasibly accessible model of the orbit around the sun including perturbations from the Moon, planets and other significant small objects in the solar system. Such a problem would be extremely time consuming, and less numerically stable, if it were necessary to compute the orbits of all of those bodies, so ephemerides were used for the positions of the Moon, planets and other solar system bodies taken into account. Given the stringent procedures used to generate ephemerides at JPL, this method is stable provided that the perturbations of the NEO on the other bodies are negligible. So, once again, we see that the information preserving transformations serve the purpose of specifying a model in which information is more feasibly accessed on the basis of one with less feasible access.

The "upward" transformations, however, serve a somewhat different purpose. These transformations are back-interpretations, but they have a different function.

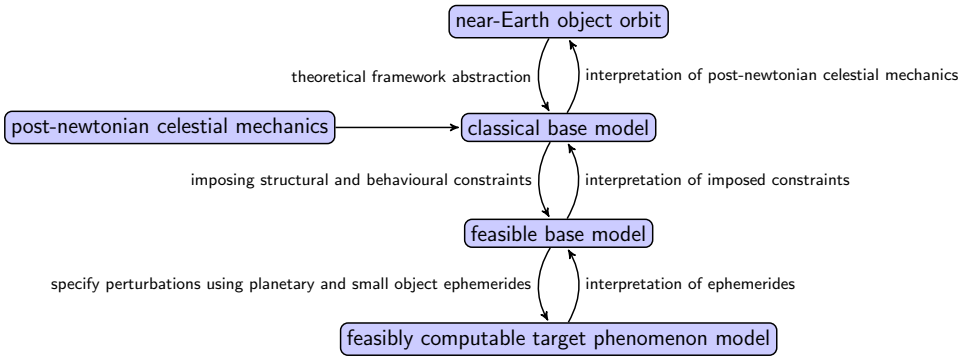

Figure 6.2: A representation of the process of construction of a reduced model of a NEO suitable for use in the process of orbit determination and impact probability estimation.

Rather than being required as part of the process of solving the principal modeling problem, they are required to interpret the the meaning of the solution obtained from this process in relation to the target phenomenon. We may see that the back-interpretation process in modeling involves the conversion of the solution to a mathematical problem into an answer to a modeling question. The original question in the NEO modeling case was to be answered by predicting a precise time of a close approach and giving a reliable estimate of the likelihood that an impact would occur. To answer this question, we converted it into a series of mathematical problems, in particular the orbit determination problem and the impact probability estimation problem. Solving this problem in a reliable way, then, allows us to make a reliable inference concerning whether a NEO will make a close approach to Earth in the future, and if so when and with what likelihood of impact. Thus, we see that the back-interpretation process in the case of modeling is an important part of completing the inference in relation to the phenomenon.

It is more than this, however, since it is also important for assessing the reliability of the inference that has been made. In the case of the back-interpretation step from the reduced model to the feasible base model, the interpretation of the ephemerides tells us that the results of the orbit computation will be unreliable if the object is sufficiently large to perturb Earth's orbit as a result of a close approach to Earth. For most cases, however, this is not a concern. In the step from the feasible base model to the classical one, the interpretation of the imposed constraints requires understanding the manner in which the imposed constraints are approximations. This may be relevant in this case, because the Yarkovsky effect, *i.e.*, perturbation from

anisotropic thermal emission of a rotating object, can have an effect on its orbit. So, assuming that the object is spherically symmetric will ignore this effect unless it has already been added as a perturbation of the equation of motion. Finally, the interpretation of what all this tells us about the actual phenomenon requires interpreting the theoretical framework at the scale of the modeling. The common feature of all of these back-interpretation steps is that they provide us with some knowledge of the conditions on the validity of the solution feasibly obtained in the reduced model. Stated in other terms, they provide some knowledge of the conditions under which the solution will no longer accurately describe the behaviour of the phenomenon. So, in the case of model construction, the back-interpretation transformations are not required for the solution of the problem *per se*, but they are essential for allowing the solution to answer the modeling question originally posed and for placing constraints on the validity of the solution as a description of the phenomenon.

The consideration of the model construction process therefore completes the overall inferential task, since the back-interpretation of this process actually makes the connection to the target phenomenon itself. This shows, therefore, that the entire application process can illuminatingly be modeled in terms of the local effective descriptive stability of transformations from a modeling question about the target phenomenon formulated as a model problem through to the process of finding a descriptively accurate solution to the model problem that is also descriptively accurate in terms of what we need to know about the phenomenon. And we gain insight into what the overall descriptive stability of the inferential process depends upon for its reliability by understanding the effects of the error introduced in each implementation transformation that is required to find a feasible solution to a modeling problem. Thus, we see clearly that, at least in cases of scientific inference similar to the case examples considered in this study, the reliability of feasible inference relies on the descriptive stability of a complex system of transformations between descriptive frameworks. More importantly, we also see clearly exactly which error considerations matter at each stage of the process and how they contribute to the overall reliability or instability, as the case may be, of the inferential process used. Given that this pattern of descriptively stable implementations is found throughout the process of feasible scientific inference, this provides evidence that it may be a pattern seen even

in cases of scientific inference that are quite different from the details of the cases we have considered.

6.2.2 Theory Application as a Computational Modeling Process

The concepts of the state mode of presentation of local constraint logic have proved to be quite illuminating, revealing a common recursive pattern of feasible inference throughout the application process. Considering local constraint logic as an epistemological model raises the question as to whether the path mode of presentation will also provide useful insights. As we will now see, it does, and allows us to see the the process of theory application as a computational modeling process in an even more robust sense than was revealed in the last subsection. As an example in this direction, consider the model construction process. This process involves the construction of an equation of motion to model the behaviour of some phenomenon. Suppose that we are working within a given theoretical framework. Then the construction of the model involves imposing constraints that are structural (*e.g.*, symmetry) and behavioural (*e.g.*, constitutive equations), followed by applying methods to reduce the complexity of the model in a way that focuses in on the dominant behaviour, which ensures that we obtain a model that is useful and/or insightful. Now, we can regard the process of imposing constraints on the theoretical framework and the methods used to simplify constraints as "rules of inference" used in the process of model construction. And considering the possible constraints in a theoretical framework as "states", we can then regard the application of these rules as generating a path through the space of possible constraints to feasibly constructible constraint systems that are capable of describing certain behaviour of a natural phenomenon.

If we consider the use of model reduction methods to convert certain constraints into less complex ones, then we can see that the model reduction methods are locally effectively valid rules of inference that transform a given model to one from which information is more accessible. Since these rules of inference are local and effective, they are not deductively valid, but they do ensure that inferences made on the basis of the more feasible model will be reliable under the conditions in which the approximation is (effectively) valid. In case that the methods used are actually deductive, *i.e.*, are exact and have no conditions, then the rules of inference are simply valid.

Thus, we see that the path mode of presentation of local constraint logic recovers the usual sense of a valid rule of inference, but generalizes it to cases where the validity is only approximate and accurate only over a certain range of conditions. This shows that the concepts from local constraint logic are able to accurately describe the methods of feasible scientific inference but the logical concepts of (classical) mathematical logic are not. Thus, rather than simply being a distortion, classical logical concepts simply fail to describe methods of inference in feasible scientific reasoning.

The case in which constraints are imposed to generate a feasible base model in a classical theoretical framework shows that "rules of inference" in feasible epistemology are not simple generalizations of classical logical concepts. This is because the initial construction of a feasibly accessible model on the basis of imposed constraints effects a transformation from a classical framework of ideally accessible constraints to a feasible framework of actually accessible constraints. This transformation bears a distinct similarity to the transformation, noted in the last chapter, from the classical framework of a numerical method to its feasible implementation on a computing device. Thus, rather than regarding such transformations as locally reliable feasible rules of inference within a constraint framework, it is more appropriate to regard them as locally reliable transformations from a classical constraint framework to a feasible constraint framework. Thus, in both the cases of constructing a feasible base model and implementing a numerical method on a machine there is a transformation between a classical framework and a feasible one.

We may start to see from this, then, that the path mode is suited to treating the recursive procedures involved in the application process, outlined in figures 6.1 and 6.2, in terms of the application of locally effectively valid rules of inference. Since the function of a locally effectively valid rule of inference is to convert a constraint problem in a source framework to a target framework where solutions provide approximately accurate solutions under certain conditions in the source framework, the considerations of the last subsection show that the steps of the application process do precisely this. This has the quite interesting consequence that the classical logical picture of theory application as a process of logical derivation does appear to apply in the case of feasible inference at the expense of generalizing the concept of deductive validity to that of local effective validity.

This enables us to see the entire application process in terms of solutions to contraint problems by finding locally effectively valid paths to solutions. The cost of increasing feasibility is generally a loss in precision or generality, but with the benefit of a description that is more amenable to analysis or computation. Indeed, this is precisely the pattern we saw in terms of the importance that classical descriptions have in making inference in science feasible. For example, we saw that the advantage of *exact* models used in the definition of the measurement reduction function is that they make it easier to find low computational cost ways of computing the information we need given the level of accuracy at which we are working. Effectively valid inference made on the basis of classically exact models is thus seen to be an essential component of the development of methods that have low enough computational cost to make them feasible. And the price for this is often that the results of the inference are valid over a smaller domain or range of circumstances than the more accurate higher cost framework or method would provide.

This way of modeling the application process also applies to the "solution" of design problems that are formulated in terms of constraints. These methods require that, given certain locally imposed constraints, such as certain functional properties the design must have, limitations of financial resources and quality of materials, it is possible to find a design that effectively meets all of the required constraints to within acceptable tolerances for error. This requires imposing more and more constraints until a particular design is picked out that effectively meets all of the required constraints. And the more complex case of layers of functional design that are implemented in lower level design frameworks, as we saw in the processor design case, there is a structure reminiscent of the more complex recursive process of computing solutions to constraint systems involved in the application of theory to the world.

To clarify the sense in which the application process is a computational modeling process, consider the fact that, from the path perspective of the application process, the "states" are the constraint frameworks within which model problems are formulated. These constraint frameworks include the models at various stages of the model construction process, the models used in the various stages of handing data and the models used in the process of computing numerical solutions. Thus, the methods used in the recursive application procedure described in the last subsection

are "state transition maps" between the "states" of the application process. And the recursive application procedure involves finding a feasible path to an accessible data-constrained problem that can be back-interpreted to provide a solution in the original modeling framework. Thus, the analogy between computation and logical inference observed in the introduction to the last chapter also obtains in the case of feasible inference. The most significant difference for our purposes is that, in the case of theory application, rather than being deductively valid, the "rules of inference" are locally effectively valid.

Now, we have seen that there are differences in the nature of the recursive processes depending on whether one is working within the framework of a model, with data handling or with computation. It turns out that the concept of a state transition map allows for a nice way of distinguishing between these types of framework as well as for distinguishing the case of an actual description of a phenomenon. In simple terms, we can consider there to be four main possibilities for a conceptual framework with a state transition map internal to the framework, based on whether the state transitions are continuous or discrete and whether the map is knowable or unknowable.

The case of a modeling framework for a continuous-time dynamical system is picked out by a continuous state transition map that is "knowable" in the sense that it is classically determined, though typically not feasibly accessible. The case of a numerical computation is picked out by a discrete state transition map that is knowable in the strong sense of being feasibly given. Thus, to draw the categorical lines here, we are mixing classical and feasible senses of 'knowable'. The case of measurement of states of dynamical phenomena is picked out by a discrete state transition, but the map generating the data points is unknowable because the measured values are obtained by observation, though in many cases the map is estimable with the aid of reliable models. And finally, the case of a phenomenon is picked out by a continuous state transition, where the map generating the transition is typically unknowable in a stronger sense, *viz.*, that the sequence of states is at best estimable. This gives rise to the four main kinds of descriptive framework shown in table 6.1.

The final matter to consider from the point of view of our model of theory application as a computational modeling process is the perspective it provides on the stability of models and theories under changes in the phenomenon being considered.

	T_t knowable	T_t unknowable
t continuous	modeling	phenomenal
t discrete	computing	measuring

Table 6.1: Four main kinds of descriptive framework distinguished in terms of the properties of the state transition map T_t and time t.

Let us restrict attention to the construction of models from theory. In the application of theory to phenomena we require not simply that a theory or model provide us with accurate descriptions in a single case, we generally require it to provide accurate descriptions of all the phenomena that the model or theory is supposed to describe. Inevitably a model will fail to be valid under certain changes of epistemic context. But the result, of course, is not necessarily the failure of the applicability of the modeling or theoretical framework, it just means that we need to rethink or adapt our modeling or theoretical strategy. The manner in which this is described in our epistemological model is in terms of the concept of *backtracking* from artificial intelligence.

Considering the model construction process on the path view of local constraint logic, we follow a sequence of locally reliable inferences to construct a model that is accurate in a given context. When we encounter a context where the model fails, we then "backtrack" by reversing the last effectively valid inference rule we made in an attempt to find a different inference rule that can result in a model that is accurate in the new context. For example, if our model of the orbit of a NEO fails, we do not immediately conclude that post-newtonian mechanics is inapplicable. Rather, we work backward along the sequence of inferences we made to use data or computation to solve a problem until we find the source of the failure of the model. At this stage we can then move forward again along the sequence of inferences until we find a feasible solution. If this solution is accurate then we have restored applicability. If it is not accurate, then we repeat this backtracking process until we find a valid model. And given that we understand the effect of the approximations used in the inference rules we use, we can have some insight into why the model might have failed and how to restore validity. In some cases, will have to backtrack quite far along the model construction path in order to find a valid model, potentially to the basic equation of a modeling or theoretical framework. This provides a picture of the application of

theory as a tree, with the basic equations of a theory as the roots. The branches are then the various ways that methods of inference can be applied to construct and solve a model of a phenomenon. The concern of classical epistemology could be construed as the entire tree of possible inferences. The concern of feasible epistemology, however, is simply those branches of the tree that use feasible methods of inference.

If the backtracking process takes us all the way back to the basic equations of the theory and there is still no way to construct a valid model, then we have reached a phenomenon that the theory does not apply to. Such a phenomenon marks a *boundary* of applicability of the theoretical framework. A classic example of this kind of case is the failure of newtonian celestial mechanics to account for the 43″/century of precession of the perihelion of Mercury. There was no way that a newtonian model could account for this in terms of observable celestial bodies. Accordingly, a broader theoretical framework was required in order to allow a "backtracking" to a framework that could support a valid model of the phenomenon: general relativity. The subtlety of the failure here, however, is that the theory of newtonian celestial mechanics, as construed in the sense of the semantic view, contains plenty of models that do accurately describe the 43″/century. This is why it is possible to formulate post-newtonian models as perturbations of a newtonian N-body equation of motion, and why the framework of newtonian celestial mechanics can still provide valid descriptions even though the theory has been invalidated. This shows that there is a problem with understanding a theory in terms of its set of **models** and saying that the theory is applicable if one of its **models** is satisfied in a given context. There is a newtonian model that is satisfied in the case of the motion of the perihelion of Mercury, it is just that this model cannot be constructed on the basis of the principles of newtonian celestial mechanics, *i.e.*, Newton's laws of motion and universal gravitation. This reveals that the syntactic methods that a theory provides for the construction of models are very much a part of the theory, in contradiction to the claims of the semantic approaches we considered in this study.

This shows one way that the concept of a domain of applicability of a theory is more subtle and complex in feasible epistemology. There is also another way, one more subtle and complex, which is that there are actually multiple distinct domains of applicability that result from considerations of feasibility. One such domain is

the classical notion of a domain of application defined in terms of the existence of a model of the theory that is valid in the given context. This domain of applicability is central in considerations of applicability in classical epistemology, since classical epistemology is concerned with the most general questions about a theory and its relation to the world. But there are also other domains of feasible application defined in terms of the accessibility of knowledge.

One of the key constraints on feasibility is the ability to compute solutions to mathematical models, as we have seen. Provided that the phenomena one wishes to describe are simple, linear and well-behaved, then solutions are typically quite easy to come by. In the case, however, of phenomena that are complex, nonlinear or chaotic, solutions typically require numerical computation and sometimes the statistical interpretation of the information carried by solutions. So, the more complex the phenomenon, or the less "compressible" the phenomenon, in the sense that the dominant behaviour can be captured in a simple description, the more important numerical computation is, and the harder feasible application is. Thus, the complexity of phenomena in the domain of a given theoretical framework defines an effective boundary of applicability given our computational capacities at a given time. Since this is a boundary of feasibly computable application, it changes over time depending on the methods and technology available. Indeed, as was noted in chapter 1, the introduction of machine computation led to a vast expansion in the range of phenomena to which theories could be feasibly applied.

But machine computation is not the only factor determining the boundary of feasibly computable application. The limitations of accessible closed form analytic solutions in a given theoretical framework also have a significant impact on the boundary of feasibly computable application. An excellent example of this is general relativity. There are precious few cases in general relativity where exact solutions are possible, notable examples being certain forms of black holes and the FLRW cosmological model. Accordingly, perturbation and approximation methods are very important in general relativity itself, and post-newtonian approximations are very important in many applications. Quantum field theory is another excellent example. The computational constraints in this theory are demonstrated clearly by the fact that underlying the common scheme of using Feynman diagrams to compute decay amplitudes and

scattering cross-sections is a perturbation method. This approach allows for high accuracy, but the computations are generally extremely difficult. And where analytic methods fail, machine computation comes back to extend the feasible domain of applicability. So, the range of phenomena for which these theories can feasibly provide accurate descriptions within the theoretical framework is actually very small. Indeed, the proportion of the classical domain of applicability of a theory that is feasibly computable appears to be inversely proportional to the generality of the theory.

Another important boundary of feasible application is determined by our limitations to collect data that can carry information about features of the phenomena we are interested in. This is exemplified by a theories in high energy physics. Here we are fundamentally constrained by our ability to generate data that carry information about interactions at very high energies. This places a strong limitation on our ability to test speculative theories in high energy physics, but it also places a strong limitation on the kinds of knowledge we can have about high energy phenomena. Thus, both data and computational limits define important boundaries in our ability to feasibly apply theories to gain knowledge about phenomena.

A final example of a boundary of feasible applicability is determined by our limitations in constructing analytical models within the context of a given theory. Again, general relativity and quantum field theory provide examples where analytical construction of models can be a formidable task. So it is not only limitations on computability and data, it is also analytic limitations that determine boundaries of feasible applicability.[1] This relates to the matter discussed in the previous chapter of the role that computer algebra stands to play in extending the range of cases where theories can be applied. Developing more sophisticated techniques for the analytic construction of models, to which reduction methods and computation can then be applied, will allow us to gain insight into and knowledge about phenomena that are now out of the reach of theory.

Therefore, in addition to providing a picture of the theory application process as a computational modeling process, our epistemological model also provides provisions for considering matters of the stability of the ability of models and theories to generate

[1] This can be regarded as another case of limitations of feasible computable application, since we may regard the model construction process as a form of computation, as it is in terms of the four kinds of descriptive framework described in the last subsection.

knowledge over variations of epistemic context. This is understood in terms of the capability of a back-tracking process in a modeling or theoretical framework to permit the framework to adapt to new contexts. There are a number of reasons why such adaptability might fail in addition to reaching a boundary of the classical domain of applicability of a theory, *i.e.*, where the theory itself becomes inapplicable. The other reasons for failure of adaptability have to do with the inaccessibility of data that carries the information we need to be able to know features of the phenomenon and the inability to compute solutions to models at sufficiently high accuracy and low cost. So, a backtracking process may enable a theory to adapt to a new context by some kind of method of search for a valid model provided that the context does not lie at or beyond a boundary of feasible applicability.

6.2.3 Descriptive Validity and Adaptation for Local Models

We will now briefly consider a few important features of the structure and behaviour of our epistemological model before concluding the study. It is important to recognize that this model has only been ensured to be sufficiently accurate in a quite limited scientific context. So, the suggestions that it provides about more general structural and behavioural patterns in epistemology of sciences are hypotheses that require testing in order to be validated in a more rigorous manner. Thus, the model has a limited accuracy, as, for example, it is not describing all of the details underlying the descriptive stability of implementations in any given case. And the model has a limited domain of applicability, because, for example, those features of the model that depend essentially on the consideration of finite-dimensional ordinary differential equations will fail in contexts of infinite dimensional equations. Thus, the model is understood to be both effective and local. But it is assured to be consistent with scientific practice in the more limited context from which it was generated because it has been designed to do so. Despite the fact that the domain of applicability of the epistemological model is limited, the advantage of the approach we have taken, however, is that the model can be considered in different pieces or modules. For example, the model construction part can be separated from the numerical computation part, the numerical computation part from the data handling part, *etc.* In this way, those modules that continue to be valid in a new context can be retained, and those

modules that do not can be adapted or discarded and modeled anew. In this way, the model is designed to be local and effective, but adaptable to new contexts of scientific practice.

6.2.4 Feasible Methods in Scientific Epistemology

One of the central aims of this study was to determine characteristics that an epistemological method must have in order to be able to gain insight into the complexity of method in real science. The numerous ways in which typical logical approaches to modeling theory application fail to adequately or accurately represent the actual ways of gaining knowledge in scientific practice show that typical logical approaches are ill-suited to the task of investigating feasible scientific inference and knowledge. And the fact that feasible methods of applying theory require potentially quite complex sequences of locally valid approximations of the models of a theory shows that a model of the theory-world relation in terms of logical deduction from laws and conditions, themselves deduced from theory, is false as an account of how theories are related to phenomena in actual scientific practice. Such a deductive model can, however, still provide a valid model of inference within the theoretical framework itself. It is just incapable of accounting for the feasible methods required to actually apply theory to gain knowledge about the behaviour of phenomena. Thus, we have showed that traditional logical concepts are inadequate for the purposes of the study of feasible methods in science.

The detailed consideration of the feasible modeling of double pendulum systems and the computation of impact probabilities for near-Earth objects showed that feasible scientific inference requires the systematic application of methods to simplify descriptions of phenomena in such a way that simultaneously preserves the relevant information about behaviour and makes this information more accessible. As we have seen in this chapter, these procedures are used throughout the application process because they are essential for feasible inference in model construction, data handling and in numerical computing.

The ability to clearly bring out the patterns we have observed in feasible inference throughout the process of applying a theory relied on the epistemological modeling strategy that we developed through the course of the study. This involved the for-

mulation of an algebro-geometric constraint framework that was capable of faithfully representing the structure and behaviour of the methods of model construction in applied mathematics, followed by the formulation of a high level conceptual framework, which we called local constraint logic, in order to provide a feasible means of clarifying large scale structural patterns in theory application that are reflected in the representation of the application process in the algebro-geometric constraint framework.

The result of the application of local constraint logic to the detailed consideration of the application process in the case examples, we have seen in this section that the application process as a whole, at least in cases similar to those considered in part II, can be represented as a computational modeling process. This involves a recursive process that is capable of accessing data to constrain a modeling problem that cannot be solved directly in order to initiate a solution cascade that solves the original problem. In cases where this approach fails to result in a valid model of the target phenomenon, a backtracking process can be used to reformulate the model, vary or change modeling frameworks and, if necessary, change theoretical frameworks, in order to find a valid model of the phenomenon. We also saw how this model clarifies the manner in which error is introduced in the application process, which gives insight into what a given application depends on for its reliability. This information can then be useful in both designing reliable methods and in locating the failure when a given method fails.

Since this model has only been developed on the basis of a very limited study of scientific method, it is not expected to have wide applicability throughout science in its complete form. It does, however, provide distinct hypotheses concerning the structure and behaviour of feasible reasoning in science that can be tested by applying it to cases for which it was not designed to apply. An example of a test that is likely to be illuminating is the application of the epistemological model to continuum mechanics. This will provide a means of determining what features of the model are applicable in different cases and a means of extending the model to incorporate cases of feasible scientific reasoning it does not currently accurately represent.

Another central aim of this study was to argue that for epistemology of science to develop fully as a scientific discipline it must look to the methods of applied

mathematics, not only pure mathematics or metamathematics, since this is essential for developing effective tools for the investigation of scientific method. One line of argument for this thesis is that the semantic and syntactic views of theories are ill-suited to this task, as has been indicated. Another line of argument for this thesis derives from showing how epistemological methods that are modeled on the methods of "physics avoidance" and model reduction used to construct mathematical models of phenomena can provide feasible descriptions of scientific method that are both accurate in their representation of science and provide insight into the structure and behaviour of method in real science. Therefore, the success that our epistemological model has in elucidating new large-scale patterns in the application process in applied mathematics provides strong evidence for the fruitfulness of epistemological methods modeled on the methods of applied mathematics.

Bibliography

Arnold, V. (1989). *Mathematical methods of classical mechanics*, vol. 60. Springer.

Azzouni, J. (2000). Applying mathematics: an attempt to design a philosophical problem. *Monis*, 83(2): 209–227.

Bacon, F. (1620). *Novum organum*. Clarendon press. Published 1878.

Batterman, R. (1998). Why equilibrium statistical mechanics works: Universality and the renormalization group. *Philosophy of Science*: 183–208.

——— (2002*a*). Asymptotics and the role of minimal models. *The British journal for the philosophy of science*, 53(1): 21–38.

——— (2002*b*). *The devil in the details: Asymptotic reasoning in explanation, reduction, and emergence*. Oxford University Press, USA.

——— (2009). Idealization and modeling. *Synthese*, 169(3): 427–446.

——— (2010). On the explanatory role of mathematics in empirical science. *The British Journal for the Philosophy of Science*, 61(1): 1–25.

——— (2011). The tyranny of scales.

Bell, J. (1986). From absolute to local mathematics. *Synthese*, 69(3): 409–426.

——— (1988). *Toposes and local set theories: an introduction*. Clarendon Press.

Bell, J. & Machover, M. (1977). *A course in mathematical logic*. North-Holland Pub. Co.

Benettin, G. & Giorgilli, A. (1994). On the Hamiltonian interpolation of near-to-the identity symplectic mappings with application to symplectic integration algorithms. *Journal of Statistical Physics*, 74(5): 1117–1143.

Beth, E.W. (1948). Analyse sémantique des théories physiques. *Synthese*, 7(3): 206–207.

——— (1949). Towards an up-to-date philosophy of the natural sciences. *Methodos*, 1: 178–185.

Beth, E. (1961). Semantics of physical theories. In: *The concept and the role of the model in mathematics and natural and social sciences*. Springer, pp. 48–51.

Blanes, S. & Budd, C. (2005). Adaptive geometric integrators for Hamiltonian problems with approximate scale invariance. *SIAM Journal on Scientific Computing*, 26(4): 1089–1113.

Bogen, J. & Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97(3): 303–352.

Bourbaki, N. (1950). The architecture of mathematics. *The American Mathematical Monthly*, 57(4): 221–232.

Bueno, O. & da Costa, N. (2007). Quasi-truth, paraconsistency, and the foundations of science. *Synthese*, 154(3): 383–399.

Bueno, O., French, S., & Ladyman, J. (2002). On representing the relationship between the mathematical and the empirical*. *Philosophy of Science*, 69(3): 452–473.

Chambers, J. (1999). A hybrid symplectic integrator that permits close encounters between massive bodies. *Monthly Notices of the Royal Astronomical Society*, 304(4): 793–799.

Chodas, P. & Yeomans, D. (1999). Orbit determination and estimation of impact probability for near earth objects.

Corless, R. (1994*a*). Error backward. In: *Chaotic numerics: an International Workshop on the Approximation and Computation of Complicated Dynamical Behavior, Deakin University, Geelong, Australia, July 12-16, 1993*, vol. 172. American Mathematical Society, p. 31.

——— (1994*b*). What good are numerical simulations of chaotic dynamical systems? *Computers & Mathematics with Applications*, 28(10-12): 107–121.

Da Costa, N. & French, S. (2003). *Science and partial truth: A unitary approach to models and scientific reasoning*. Oxford University Press, USA.

Darrigol, O. (2008). The modular structure of physical theories. *Synthese*, 162(2): 195–223.

Davis, M., Sigal, R., & Weyuker, E. (1994). *Computability, complexity, and languages: fundamentals of theoretical computer science*. Morgan Kaufmann.

Enright, W. (2000). Continuous numerical methods for ODEs with defect control. *Journal of computational and applied mathematics*, 125(1-2): 159–170.

French, S. & Ladyman, J. (2011). In defence of ontic structural realism. *Scientific structuralism*: 25–42.

Geddes, K., Czapor, S., & Labahn, G. (1992). *Algorithms for computer algebra*. Springer.

Giere, R. (2010*a*). An agent-based conception of models and scientific representation. *Synthese*, 172(2): 269–281.

Giere, R.N. (2010*b*). *Scientific perspectivism*. University of Chicago Press.

Goldstein, H., Poole, C., & Safko, J. (2002). *Classical Mechanics*. Addison Wesley.

Goldstine, H.H. (1977). A history of numerical analysis from the 16th through the 19th century. *Sources in the History of Mathematics and Physical Sciences*, 2.

Green, R.M. (1985). *Spherical astronomy*. Cambridge University Press.

Hairer, E. (1994). Backward analysis of numerical integrators and symplectic methods. *Ann. Numer. Math.*, 1(1-4): 107–132.

Hairer, E., Lubich, C., & Wanner, G. (2006). *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer Verlag.

Hamming, R. (1962). *Numerical methods for scientists and engineers*. Columbus: McGraw-Hill.

Hand, L. & Finch, J. (1998). *Analytical mechanics*. Cambridge Univ Pr.

Harper, W. (2011). *Isaac Newton's Scientific Method*. Oxford University Press.

Henrici, P. (1962). *Discrete variable methods in ordinary differential equations*. Wiley & Sons.

Holmes, M. (2009). *Introduction to the foundations of applied mathematics*, vol. 56. Springer.

Horstemeyer, M. (2010). Multiscale modeling: a review. In: *Practical aspects of computational chemistry*. Springer, pp. 87–135.

Humphreys, P. (2009). The philosophical novelty of computer simulation methods. *Synthese*, 169(3): 615–626.

Jackson, J.D. & Fox, R.F. (1999). Classical electrodynamics. *American Journal of Physics*, 67: 841.

Knuth, D.E. (1981). The art of programming, vol. 2, semi-numerical algorithms.

Knuuttila, T. (2004). Models, representation, and mediation.

Kopeikin, S., Efroimsky, M., & Kaplan, G. (2011). *Relativistic celestial mechanics of the solar system*. Wiley-VCH.

Krantz, D., Suppes, P., & Luce, R. (1971). *Foundations of measurement*, vol. 1. academic press New York.

Ladyzhenskaya, O.A. (1967). New equations for the description of the motions of viscous incompressible fluids, and global solvability for their boundary value problems. *Trudy Matematicheskogo Instituta im. VA Steklova*, 102: 85–104.

Lawden, D.F. (1989). *Elliptic functions and applications*. Springer New York.

Lichtenberg, A. & Lieberman, M. (1992). *Regular and chaotic dynamics*, vol. 38. Springer Verlag.

Luce, R., Krantz, D., Suppes, P., & Tversky, A. (1990). *Foundations of measurement, Vol. 3: Representation, axiomatization, and invariance*. Academic Press.

Mac Lane, S. (1986). *Mathematics, form and function*. Springer-Verlag New York.

Massimi, M. (2011). From data to phenomena: a kantian stance. *Synthese*, 182(1): 101–116.

McKinsey, J.C., Sugar, A., & Suppes, P. (1953). Axiomatic foundations of classical particle mechanics. *Journal of Rational Mechanics and Analysis*, 2(2): 253–272.

Mikkola, S. & Innanen, K. (1999). Symplectic tangent map for planetary motions. *Celestial Mechanics and Dynamical Astronomy*, 74(1): 59–67.

Milani, A., Chesley, S., Chodas, P., & Valsecchi, G. (2002). Asteroid close approaches: analysis and potential impact detection. *Asteroids III*, 1: 89–101.

Milani, A. & Gronchi, G. (2010). *Theory of orbit determination*. Cambridge University Press.

Moir, R. (2010). *Reconsidering backward error analysis for ordinary differential equations*. Master's thesis, The University of Western Ontario.

Moore, G.E. *et al.* (1965). Cramming more components onto integrated circuits.

Moore, S.W. (1996). *Multithreaded processor design*. Kluwer Academic Publishers.

Morgan, M. & Morrison, M. (1999). *Models as mediators: Perspectives on natural and social science*, vol. 52. Cambridge Univ Pr.

Moulines, C. (1996). Structuralism: The basic ideas. *enW. Balzer y CU Moulines, Structuralist Theory of Science. Focal Issues, New Results, Berlin, Walter de Gruyter*.

Muller, F. (2011). Reflections on the revolution at stanford. *Synthese*, 183(1): 87–114.

Naylor, R.H. (1989). Galileo's experimental discourse. In: *The Uses of Experiment: Studies in the Natural Sciences*. Cambridge University Press, pp. 117–34.

Norton, J. (2011). Challenges to bayesian confirmation theory. *Philosophy of Statistics*, 7: 391.

Pincock, C. (2009). Towards a philosophy of applied mathematics. *New Waves in the Philosophy of Mathematics', Palgrave Macmillan*: 173–194.

——— (2004). A new perspective on the problem of applying mathematics. *Philosophia Mathematica*, 12(2): 135–161.

Rossi, F., Van Beek, P., & Walsh, T. (2006). *Handbook of constraint programming*. Access Online via Elsevier.

Roth, G.D., Augensen, H.J., & Heintz, W.D. (1994). *Compendium of practical astronomy*, vol. 1. Springer-Verlag Berlin Heidelberg New York.

Roy, A.E. (2010). *Orbital motion*. CRC Press.

Sanz-Serna, J. (1992). Symplectic integrators for Hamiltonian problems: an overview. *Acta Numerica*, 1: 243–286.

Scott, D. & Suppes, P. (1958). Foundational aspects of theories of measurement. *The journal of symbolic logic*, 23(2): 113–128.

Segre, M. (1980). The role of experiment in galileo's physics. *Archive for History of Exact Sciences*, 23(3): 227–252.

Sipser, M. (2012). *Introduction to the Theory of Computation*, vol. 3. Cengage Learning.

Sitarski, G. (2002). Warsaw ephemeris of the solar system: De405/waw. *Acta Astronomica*, 52: 471–486.

Skeel, R. (1999). Symplectic integration with floating-point arithmetic and other approximations. *Applied numerical mathematics*, 29(1): 3–18.

Smith, L.A. (2002*a*). What might we learn from climate forecasts? *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl 1): 2487–2492.

Smith, S. (2002*b*). Violated laws, ceteris paribus clauses, and capacities. *Synthese*, 130(2): 235–264.

Stetter, H. (2004). *Numerical polynomial algebra*. Society for Industrial Mathematics.

Stuart, A. & Humphries, A. (1996). Dynamical systems and numerical analysis. 1996. *Cambridge Monographs on Applied and Computational Mathematics*.

Suppe, F. (1974). *The structure of scientific theories*. University of Illinois Press.

——— (1989). *The semantic conception of theories and scientific realism*. University of Illinois Press.

Suppes, P. (1962). Models of data. In: *Logic, methodology and philosophy of science: Proceedings of the 1960 International Congress*. Stanford University Press: Stanford, pp. 252–261.

Suppes, P., Krantz, D., Luce, R., & Tversky, A. (1989). *Foundations of measurement, Vol. 2: Geometrical, threshold, and probabilistic representations*. Academic Press.

Suppes, P. (1961). *A comparison of the meaning and uses of models in mathematics and the empirical sciences*. Springer.

——— (1969). *Models of data*. Springer.

——— (2011). Future development of scientific structures closer to experiments: Response to fa muller. *Synthese*, 183(1): 115–126.

Taff, L.G. (1985). Celestial mechanics: a computational guide for the practitioner. *New York, Wiley-Interscience, 1985, 540 p.*, 1.

Taff, L. (1991). *Computational spherical astronomy*. Krieger Pub Co.

Teller, P. (2010). "saving the phenomena" today. *Philosophy of Science*, 77(5): 815–826.

Trefethen, L.N. & Bau III, D. (1997). *Numerical linear algebra*. 50. Siam.

Tucker, W. (2002). A rigorous ode solver and smale's 14th problem. *Foundations of Computational Mathematics*, 2(1): 53–117.

van Fraassen, B.C. (1970). On the extension of beth's semantics of physical theories. *Philosophy of science*: 325–339.

van Fraassen, B. (1980). *The scientific image*. Oxford University Press, USA.

——— (2010). Scientific representation: Paradoxes of perspective. *Analysis*, 70(3): 511–514.

Weyl, H. (1949). *Philosophy of mathematics and natural science*. Princeton University Press.

Wiegert, P.A., Innanen, K.A., & Mikkola, S. (2007). The orbital evolution of near-earth asteroid 3753. *The Astronomical Journal*, 115(6): 2604.

Wilson, M. (1998). Mechanics, classical. *Routledge Encyclopedia of Philosophy*.

——— (2000). The unreasonable uncooperativeness of mathematics in the natural sciences. *Monis*, 83(2).

——— (2006). *Wandering significance: An essay on conceptual behavior*. Oxford University Press, USA.

Wimsatt, W. (2007). *Re-engineering philosophy for limited beings: piecewise approximations to reality*. Harvard Univ Press.

Winsberg, E. (2001). Simulations, models, and theories: Complex physical systems and their representations. *Philosophy of Science*: 442–454.

——— (2010). *Science in the age of computer simulation*. University of Chicago Press.

Wisdom, J. (1987). Urey prize lecture: Chaotic dynamics in the solar system. *Icarus*, 72(2): 241–275.

Wisdom, J. & Holman, M. (1991). Symplectic maps for the n-body problem. *The Astronomical Journal*, 102: 1528–1538.

Worrall, J. (1989). Structural realism: The best of both worlds?*. *Dialectica*, 43(1-2): 99–124.

Zeeman, E. (1976). Euler buckling. In: *Structural stability, the theory of catastrophes, and applications in the sciences*. Springer, pp. 373–395.

Zeigler, B. (1976). *Theory of Modelling and Simulation*. John Wiley.

Zhong, G. & Marsden, J.E. (1988). Lie-poisson hamilton-jacobi theory and lie-poisson integrators. *Physics Letters A*, 133(3): 134–139.

Curriculum Vitae

Robert Hugh Caldwell Moir

Education

PhD, Philosophy 2004-2013
Western University, London, Canada
Thesis: *Structures in Real Theory Application: A Study in Feasible Epistemology*
Supervisors: Robert Batterman and John Bell

MSc, Applied Mathematics 2009-2010
Western University, London, Canada
Thesis: *Reconsidering Backward Error Analysis for Ordinary Differential Equations*
Supervisor: Robert Corless

MA, Philosophy 2003-2004
Western University, London, Canada

BA, Mathematics and Philosophy (*First Class Joint Honours*) 2001-2003
McGill University, Montréal, Canada
Honours Thesis: *Infinity and Physical Theory*
Supervisor: Michael Hallett

BSc, Physics (minor Chemistry) (*First Class Honours*) 1995-2001
McGill University, Montréal, Canada

Areas of Specialization

Philosophy of Physics, Philosophy of Applied Mathematics, Applied Mathematics

Areas of Competence

Logic, Philosophy of Mathematics, Philosophy of Science

Awards and Distinctions

Research Awards

University of Pittsburgh

- Visiting Scholar ($26,000), 2011-2012

Western University

- Western Graduate Research Scholarship ($8,000), 2009-2010

Social Sciences and Humanities Research Council of Canada

- Doctoral Fellowship ($40,000), 2007-2009

Western University

- Western Graduate Research Scholarship ($8,000), 2005-2006

Western University

- Special University Scholarship ($13,000), 2003-2005

Academic Awards

Chemical Institute of Canada National High School Chemistry Examination

- Toronto District Winner, 1995

Publications

Proceedings

Moir, R. (2009). The Conversion of Phenomena to Theory: Lessons on Applicability from the Early Development of Electromagnetism. In: A. Cupillari (Ed.), *Proceedings of the Canadian Society for History and Philosophy of Mathematics*, St John's NL, June 2009, pp. 68-91.

Posters

Batterman, R.W., Fillion, N., Moir, R., Overton, J. (2010). "Idealization in Scientific Explanation", Western Research Day, Western University, March 24, 2010.

Talks

**Peer-reviewed *Abstract Submission †Invited

1. **"Rational Discovery of the Natural World: An Algebro-Geometric Response to Steiner" MathFest 2013/Canadian Society for the History and Philosophy of Mathematics (MathFest 2013/CSHPM) Conference, Hartford CT, 1-3 August 2013.
2. **"Computation for Confirmation." Conference on The Plurality of Numerical Methods and their Philosophical Analysis. Université Paris 1—Panthéon-Sorbonne, Paris, November 3-4 2011.
3. † "Computation in Scientific Explanation," in the course Contemporary Philosophy of Science, Department of Philosophy, University of Guelph, 17 November 2010.
4. ** with Nicolas Fillion, "Explanation and Abstraction: The Case of Backward Error Analysis" Philosophy of Science Association Biennial Meeting, Montréal, Québec, 4-6 November 2010.
5. * with Nicolas Fillion, "Modeling and Explanation: Lessons from Modern Error Theory." Canadian Society for the History and Philosophy of Science (CSHPS) Conference, Concordia University, Montréal, Québec, 28-30 May 2010.
6. with Nicolas Fillion, "A Step Forward with Backward Error," PGSA Colloquium Series, Department of Philosophy, The University of Western Ontario, 12 March 2010.
7. * "The Conversion of Phenomena to Theory: Lessons on Applicability from the Development of Electromagnetism." Canadian Mathematical Society/Canadian Society for the History and Philosophy of Mathematics (CMS/CSHPM) Conference, Memorial University, St. John's, Newfoundland, 6-8 June 2009.
8. "From the World to Mathematics and Back Again: What We Can Understand About Applicability from the Development of Electromagnetism." PGSA Colloquium Series, Department of Philosophy, The University of Western Ontario, 25 March 2009.
9. * "Theories, Models and Representation: Lessons from Solid State Physics." Canadian Society for the History and Philosophy of Science (CSHPS) Conference, University of British Columbia, Vancouver, British Columbia, 3-5 June 2008.
10. "Theories, Models and Representation: Lessons from Solid State Physics." PGSA Colloquium Series, Department of Philosophy, University of Western Ontario, 12 March 2008.

Academic Experience

Instructor

Western University 2010-2013

- Metaphysics and Epistemology of Witchcraft (Full-Year Course), 2012-2013
- Introduction to Logic (12-week Accelerated Two-Semester Course), 2011
- Critical Thinking (Full-Year Course), 2010–2011

Teaching Assistant

Western University 2003-2010

- Linear Algebra for Engineers (Half-Year Course), 2010
- Calculus (Half-Year Course), 2009
- Introduction to Philosophy (Full-Year Course), 2005–2006, 2007–2008
- Critical Thinking and Reasoning (Full-Year Course), 2003–2005

Workshop Leader

Rotman Institute of Philosophy 2008-2009

- Modern Mathematics for Philosophers (Algebra and Analysis), 2008-2009

Research Assistant

Western University 2007-2010

- Robert Corless, Department of Applied Mathematics, 2009–2010
- Rotman Canada Research Chair in Philosophy of Science, 2009–2010
- The Joseph L. Rotman Institute for Science and Values, 2008-2009

Service

Conference Organization

- Logic, Mathematics, and Physics Graduate Philosophy Conference
  Department of Philosophy, Western University
  2006: Co-organizer with J. Noland and D. MacDonald.
  Keynote Speaker: Michael Hallett (McGill University)