SURFACES, SCALES, AND SYNTHESIS SCIENTIFIC REASONING AT THE NANOSCALE

by

Julia R. Bursten

B.A., Philosophy, Rice University, 2008M.A., Philosophy, University of Pittsburgh, 2010

Submitted to the Graduate Faculty of

the Kenneth P. Dietrich School of Arts and Sciences in partial

fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2015

UNIVERSITY OF PITTSBURGH KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Julia R. Bursten

It was defended on

April 9, 2015

and approved by

Robert Batterman, University of Pittsburgh, Philosophy

Jill Millstone, University of Pittsburgh, Chemistry

Sandra Mitchell, University of Pittsburgh, HPS

John Norton, University of Pittsburgh, HPS

Mark Wilson, University of Pittsburgh, Philosophy

James Woodward, University of Pittsburgh, HPS

Dissertation Director: Robert Batterman, University of Pittsburgh, Philosophy

Copyright © by Julia R. Bursten2015

ABSTRACT

SURFACES, SCALES, AND SYNTHESIS SCIENTIFIC REASONING AT THE NANOSCALE

Julia R. Bursten, PhD

University of Pittsburgh, 2015

Philosophers interested in scientific methodology have focused largely on physics, biology, and cognitive science. They have paid considerably less attention to sciences such as chemistry and nanoscience, where not only are the subjects distinct, but the very aims differ: chemistry and nanoscience center around synthesis. Methods associated with synthesis do not fit well with description, explanation, and prediction that so dominate aims in philosophers paradigm sciences. In order to synthesize a substance or material, scientists need different kinds of information than they need to predict, explain, or describe. Consequently, they need different kinds of models and theories. Specifically, chemists need additional models of how reactions will proceed. In practice, this means chemists must model surface structure and behavior, because reactions occur on the surfaces of materials.

Physics, and by extension much of philosophy of science, ignores the structure and behavior of surfaces, modeling surfaces only as boundary conditions with virtually no influence on material behavior. Such boundary conditions are not seen as part of the physical laws that govern material behavior, so little consideration has been given to their roles in improving scientists understanding of materials and aiding synthesis. But especially for theories that are used in synthesis, such neglect can lead to catastrophic modeling failures. In fact, as one moves down toward the nanoscale, the very concept of a material surface changes, with the consequence that nanomaterials behave differently than macroscopic materials made up of the same ele- ments. They conduct electricity differently, they appear differently colored, and they can play different roles in chemical reactions. This dissertation develops new philosophical tools to deal with these changes and give an account of theory and model use in the synthetic sciences. Particularly, it addresses the question of how models of materials at the nanoscale fit together with models of those very same materials at scales many orders of magnitude larger. To answer this and related questions, strict attention needs to be paid to the ways boundaries, surfaces, concepts, models, and even laws change as scales change.

Keywords: philosophy of science, philosophy of physics, philosophy of chemistry, explanation, nanoscience, synthesis, models, theories, kinds.

TABLE OF CONTENTS

PREFACE		
I.	SMALLER THAN A BREADBOX: SCALE AND NATURAL KINDS	1
	A. A Puzzle of Kinds	1
	B. The Methodological Problem of Kinds	4
	C. Chemical Kindhood: The Structure–Property Paradigm	10
	D. What are Nanomaterials?	14
	1. Classifying Nanomaterials	16
	E. Scale-Dependence and Nanoscale Kinds	22
	F. Conclusions	27
II.	SURFACE TENSION	30
	A. Introduction	30
	B. The Sound of an A String: Boundary Conditions, Multi-Valued Concepts,	
	& Model Interaction	32
	C. The Multi-Valuedness of <i>Surface</i>	39
	D. Modeling Macroscopic Surfaces	42
	1. Continuum Models	43
	2. Structural Models	46
	E. Nanoscale Surfaces	51
	F. Conclusions	55
III.	MAKING THINGS SMALL	56
	A. Introduction	56
	B. Limiting Growth: The Challenge of Nanosynthesis	58

	C. Models of Growth	59
	D. Model Interaction	62
	E. Growth Limitation Techniques	64
	1. Vapor–Liquid–Solid Synthesis: Limiting Growth and Directing Growth .	65
	2. Anisotropic Metal Nanoparticles	67
	F. Tuning in Nanosynthesis	71
	1. Modeling LSPR	72
	G. Model Interaction and Multi-Scale Modeling	75
IV.	HOW THEORIES (AND MODELS) SHAKE HANDS	84
	A. Introduction	84
	B. Tuning: The Synthetic Use of Theories and Models	88
	C. Modeling Synthetic Constraints	93
	D. Iterative Interpolation and Instrumental Attitudes	97
	E. Conclusions	99
V.	CONCLUSIONS 1	03
	A. Overview	03
	B. Classification, Scale-Dependence, and Model Interaction	05
	C. Implementing Model Interaction	07
	D. Synthetic Uses of Models and Theories: Tuning, Design, and Iterative Inter-	
	polation	09
	E. Coda: Use and Structure	12
BIB	LIOGRAPHY	14

LIST OF FIGURES

II.1	Violin harmonics	34
II.2	Violin structure	34
II.3	Surface bond structure in a crystalline solid	38
II.4	Surface-to-volume ratios and scale	41
II.5	Material property groupings	45
II.6	Bravais lattices	48
II.7	Metal grains	50
II.8	Stained glass	52
II.9	Localized surface plasmon resonance	53
III.1	LaMer's Model of Growth	61
III.2	Vapor–liquid–solid growth	67
III.3	Nanoparticle rainbow	72
III.4	Mie model	74
IV.1	Nanocubes	87
IV.2	"Tuning" Across Sciences	88
IV.3	An inclined-plane problem	91

PREFACE

Nanoscience research and development has the potential to reshape human understanding of physical systems and the technologies they produce. Research into nanoscale systems is currently one of the fastest-growing areas of scientific interest, and nano is finding niches across all the mathematical and natural sciences—in solid-state physics (e.g. Atwater 2007), synthetic chemistry (e.g. Klabunde and Richards 2001), biological technologies (e.g. Nussinov and Alemán 2006), diagnostic (e.g. West and Halas 2003) and therapeutic (e.g. An and Hyeon 2009) medicine, and more.

Small is getting huge, and the rapid growth of the field has focused almost exclusively on experiments and application. Most research projects in nanoscience are centered around the development of particular material systems, methods for production of those systems, and applications of those systems toward technologies that could change the way energy and information is stored and deployed. Consequently, nanoscience researchers have not so far devoted much time to creating a systematic understanding of the conceptual and theoretical foundations of nanoscale systems.

What this means is that key questions about the relation between scientific concepts deployed at macroscopic and molecular scales and the nanoscale analogues of those concepts have yet to be answered. Models and theories at the nanoscale are currently scattershot; they have been developed semi-empirically for use in prediction of the behaviors of individual systems, and many models make assumptions about the target systems that conflict with either assumptions made by other models or with known facts about the nature of the target system itself. For instance, many theoretical descriptions of nanosynthesis, cobbled together from conflicting models, describe the synthesized material as sometimes continuous, sometimes discrete; sometimes symmetrical and defect-free, sometimes non-symmetric and defect-ridden.

Traditional philosophical accounts of modeling and of theory structure suggest this blooming, buzzing confusion of inconsistent models and conflicting assumptions indicates poor theory design. Nonetheless, models and theories of nanoscale systems are often quite successful at providing nanoscience researchers with the information they seek about a target system. This dissertation aims to justify this success and to suggest strategies for future success, both in the case of nanoscience and in other synthetic sciences.

I begin from the premise that the central activity of nanoscience research, and of chemical research more generally, is *synthesis*, the development of substances and materials from other substances and materials. Synthesis provides a different aim for scientific theory and practice—namely, making—than standard philosophical accounts have thus far acknowledged. Most standard philosophical accounts teach that the aim of scientific theories, insofar as they have an aim, is either description of natural systems, or prediction of future behaviors of natural systems, or explanation of observed correlations or causal connections between those behaviors, or some combination thereof. Synthesis of a material or substance is neither describing nor predicting nor explaining: it is *making*, and making is a different kind of scientific activity. There is no *a priori* reason to expect that the structure of scientific theories aimed at describing or predicting or explaining should be the same as the structure of scientific theories aimed at making.

Nor is there any *a priori* reason to expect that theories aimed at making should all display a single structure. So rather than generalizing from a survey of such theories—borrowing from, e.g., mechanical engineering, biomedicine, informatics, and applied economics, among others—I study nanoscience in detail. Uncovering the character of theories in nanoscience thereby serves as a groundwork for future work on the structure and use of theories in other synthetic sciences. Nanoscience makes a particularly apt base case because it also contains a number of pressing, endemic philosophical puzzles, whose resolution could affect the development of the field itself. Among those, most prominent are methods for gaining control over surface structure and behavior through synthesis techniques, but other considerations also arise.

From this focused study of nanoscience, I draw a number of philosophical morals about

the structure of scientific reasoning in synthetic science in general and in nanoscience in particular. I demonstrate that central debates in philosophy of science, including the nature(s) of explanation and understanding, the structure(s) of theories and the referents of scientific terms, have all grown up on the premise that the central activity of science is to describe the natural world. With the rejection of that premise comes a host of new opportunities for philosophical reflection on scientific reasoning.

Philosophers interested in scientific methodology have focused largely on physics, biology, and cognitive science. They have paid considerably less attention to sciences such as chemistry and nanoscience, where not only are the subjects distinct, but the very *aims* differ: chemistry and nanoscience center around *synthesis*. Reasoning strategies associated with synthesis do not fit well with description, explanation, and prediction that so dominate aims in philosophers' paradigm sciences; in order to synthesize a substance or material, scientists need different kinds of information than they need to predict, explain, or describe. Consequently, they need different kinds of models and theories. Specifically, chemists need additional models of how reactions will proceed.

In practice, this means chemists must model surface structure and behavior, because reactions occur on the surfaces of materials. Much of physics, and by extension much of philosophy of science, ignores the structure and behavior of surfaces, modeling surfaces only as boundaries on a system with virtually no influence on material behavior. Such boundaries are not seen as part of the physical laws that govern material behavior, so little consideration has been given to their roles in improving scientists' understanding of materials and aiding synthesis. But especially for theories that are used in synthesis, such neglect can lead to catastrophic modeling failures.

In fact, as one moves down toward the nanoscale, the very concept of a material surface changes, with the consequence that nanomaterials behave differently than macroscopic materials made up of the same elements. They conduct electricity differently, they appear differently colored, and they can play different roles in chemical reactions. These differences are the result of their sheer smallness: when materials shrink down to the nanoscale, atoms in a material stand in different relations to one another than in macroscopic materials. A larger proportion of atoms lie on the surface of nanomaterials than macroscopic materials, and this difference is responsible for most of these scale-dependent changes in material behaviors. This dissertation develops new philosophical tools to deal with these changes and construct an account of theory and model use in the synthetic sciences. Particularly, it addresses the question of how models of materials at the nanoscale fit together with models of those very same materials at scales many orders of magnitude larger.

The size of the material, and the proportion of atoms on the surface of a material, are nothing like the laws and mechanisms that philosophers usually cite as *explanantia*. This difference in size is instead a difference in the role of a *boundary condition*, that is, an independently-specified parameter that constrains the behavior of a system. Boundary conditions are a type of modeling parameter usually associated with mathematical models containing differential equations. Traditional accounts of theories and models in philosophy of science have mostly ignored the role of boundary conditions in generating explanations, predictions, and descriptions of systems of scientific interest. Traditional accounts have also ignored the function of theories and models in synthetic scientific activities, such as the design and production of new substances and materials. This dissertation explains why these oversights are problematic for accounts of theories, models, explanations, and concepts in nanoscience and in the physical sciences more generally, and it constructs an alternative account of theories and models that rectifies these oversights.

Models of nanomaterials have developed by combining and adapting theories about macroscopic material behavior and theories about atomic and molecular behavior in response to empirical observation. For instance, some models of nanomaterials have evolved from continuum models of macroscopic materials. Continuum models ignore molecular structure, which makes them effective or phenomenological rather than "foundational," and they model surfaces as infinitesimal boundary conditions. These effective models produce accurate predictions and useful explanations of the behavior of macroscopic materials while ignoring both the atomic constitution of matter and the structure and behavior of material surfaces.

Treating surfaces as boundary conditions is an effective strategy for macroscopic materials modeling because surfaces make up a small proportion of the material as a whole. However, when materials shrink down to the nanoscale, this is no longer the case, and continuum models of materials must be modified and supplemented with structural information about inter-atomic relations in the material. This challenge to traditional materials modeling strategies demonstrates that *surface* and other central concepts in materials modeling are scale-dependent; that is, their role in describing and constraining the behavior of materials changes as a function of scale.

The modeling challenges posed by the scale-dependence of *surface* are not resolved by standard accounts of inter-theory and inter-model relations, such as reduction and emergence. Instead, scientists employ strategies such as recruiting models of surface structure to supplement continuum models and empirical data, and using these supplemental conceptual resources to produce hybrid, multiscale models. This strategy, which I term *model interac-tion*, produces models that are especially well-suited not only to explanation and prediction in nanoscience, but also to designing and synthesizing nanomaterials. I use model interaction as the foundation for a new view of inter-theory and inter-model relations that focuses away from "vertical" relations like reduction and emergence and toward the use of multiple theoretical frameworks to generate novel predictions, explanations, and synthetic protocols.

In designing and synthesizing nanomaterials, and materials more generally, scientists' primary concern is determining input parameters that will lead to a specified outcome. Input parameters are modified by manipulating the boundary conditions of the proposed system, rather than by modifying the laws or mechanisms in play in the system. So modeling strategies that put boundary conditions center stage, as the model interaction account does, serve scientists' design and synthesis needs better than law- and mechanism-centered accounts of theories and models.

One of the upshots of re-conceiving contemporary philosophy of science in light of synthetic science is a new relationship between science and its philosophy. Synthetic sciences tend to focus on experimental outcomes more than on the organization of theoretical knowledge, and here philosophers may find themselves in a position to gather information about the science and make suggestions for improvement of theories that could guide future synthetic research. Rather than acting as a handmaiden to, or watchdog for, science, the philosophy of science I carry out in this dissertation acts as a partner in the project of making materials better.

To that end, this dissertation has been developed by working closely with the Millstone

Nanochemistry Research Group, an active laboratory at the University of Pittsburgh led by Dr. Jill Millstone. The Millstone group, like many nanosynthesis communities, frequently encounters challenges that are not resolvable by gathering more experimental data alone. For instance, in a recent experiment to find a new method of synthesizing a particular kind of nanoparticle, the group struggled over whether to conceive of their new method as a form of homogeneous nucleation or to relax their definition to the more permissive *in situ* nucleation. Cooperative critical reflection on what hung in the balance of using one or the other term brought to light the fact that the different definitions carried with them different assumptions about how future syntheses would be modeled and carried out. While I will not make general recommendations about how to carry out future projects with this sort of partnership-based relationship between philosophy and science, I hope that the cases in this dissertation prove illuminating both to the possibility of such projects and to their benefits to both science and philosophy.

The dissertation proceeds in five chapters. In Chapter I I study the problem of classifying nanomaterials in order to accomplish three projects. First, I demonstrate the relevance of nanoscience to extant philosophical problems in metaphysics, philosophy of language, and philosophy of science. I present a first result for philosophy of attending to nanoscience, namely, I urge the need to reframe the natural kinds debate in philosophy of science in light of classification practices in nanoscience. Second, I motivate the two central philosophical themes of the dissertation, namely (1) the scale-dependence of scientific concepts, models, and theories, and (2) the possible differences in reasoning strategies between descriptive and synthetic scientific projects. Additionally, in accomplishing these goals, I introduce readers to fundamental concepts of nanoscience, which will be revisited throughout the remainder of the dissertation.

In Chapter II, I provide an extended case study of nanoscale surfaces, arguing that *surface* is a scale-dependent concept whose role within synthetic practice changes significantly at the nanoscale. I show how scale constrains material behavior, modeling practices, and concept formation, both in nanoscience and in the physical sciences more generally. I argue that this pattern of constraint is not captured by existing philosophical frameworks for laws and inter-theory relations. I provide an extended case study of nanoscale surfaces, arguing that *surface*

is a scale-dependent concept whose role within synthetic practice changes significantly at the nanoscale. This descriptive result is coupled with suggestions for philosophy of science to the effect that traditional views of inter-theoretic relations do not hold for nanomaterials, and suggestions for nanoscience about how to use multiple conflicting models of a material effectively. My alternative *model interaction* account of inter-theoretic relations is proposed.

Chapter III turns to the activity of synthesis at the nanoscale, where I reinforce the role of scale-dependence in nanoscience by showing how constraining growth to the nanoscale poses novel synthetic challenges for nanoscientists. I argue that the ways in which multiscale models are used to solve problems in nanosynthesis are distinct from the uses of models philosophers have identified. I show how nanosynthesis poses a certain kind of design problem, which I identify as a *tuning* problem, and that this problem is distinct from problems of prediction, description, and explanation. In nanosynthesis in particular, the main tuning problem lies in constraining the growth of surfaces such that materials do not exceed the nanoscale. To solve this problem, chemists must draw on multiple models across a variety of scales, and these models are used to determine what inputs will lead to a target output, rather than what outputs are expected of a given set of inputs. This difference in problem type leads to a difference between the ways in which models are used in nanosynthesis and the ways they are used in non-synthetic sciences, and this difference explains why models and inter-theory relations in nanoscience look different from models and inter-theory relations in non-synthetic sciences.

Chapter IV builds on the results in Chapter III to articulate the structure and use of models and theories in synthetic sciences. I address the question of what it means for scientific modeling to be effective in the context of synthetic sciences more generally, where a central goal of the practice is to make new substances, materials, or objects. I build on the previous chapter's discussion of tuning to articulate an account of model and theory use in synthetic science. I argue that the ways in which models are used influences the structure of those models, which paves the way for a new view of models and inter-theory relations grown from the soil of synthetic science. Drawing from materials science, engineering, and design, I develop an account of theory use that explains how model interaction arises as a response to tuning problems in the synthetic sciences, why it differs from standard views of the roles of theories and models in scientific practice, and how this story of model formation and use is necessary for a proper understanding of how knowledge is obtained and confirmed in the synthetic sciences.

In Chapter V I conclude by addressing the broader implications of this research in philosophy of science, focusing especially on the implication of my research that study of scientific reasoning tools should begin not by articulating the logical structure of theories and models, but by assessing the means by which these tools are used to solve scientific problems.

* * * *

There are many metaphors that people commonly invoke to express sentiments about the process of writing a dissertation: "a labor of love," "a Sisyphean trial," "like giving birth," "like pulling teeth," "like walking on hot coals," "like that time when Khaleesi was standing at the gates of Qarth asking to be let into the Greatest City that Ever Was or Will Be with her dying horde of Dothraki followers and she threatened the guardians of the city instead of begging for mercy and then you had to wait to see what they did, but way more boring and it took 4 years of actual life instead of 20 seconds of screen time." They are all true; they are all inadequate; they are all the best that we writers of dissertations can do in the end.

This project is, ultimately, the work of many inspirational and editorial minds. In my own terminology, I have tuned my dissertation—iteratively, interpolatively—using input from many models of what philosophy of science should look like. John Norton first inspired my investigation into the reasoning strategies that support synthesis when he asked me the question, "What are the interesting philosophical problems native to chemistry?" Sandra Mitchell guided the early stages of my research by sharpening my focus on how models and theories interact during a seminar on Emergence that she led, and her theory of integrative pluralism makes many an appearance in the ensuing pages. Jim Woodward has provided a keen editorial eye and ear, honing my arguments especially in the first few chapters of the thesis and forcing me into clarity despite my best efforts to resist it. And no text has inspired me as a philosopher more than Mark Wilson's Big Pink Book, *Wandering Significance*. Mark's ability to connect seemingly disparate pieces of scientific reasoning through clever examples and sheer deep conceptual analysis has saved this dissertation more than once.

I will never forget the trepidation I felt when Jill Millstone first requested a meeting with me before admitting me to her introductory class on nanoscience, in large part because it contrasts so strongly with the camaraderie and passion that have marked the rest of our interactions over the past three years. In Jill, I found a kindred spirit. She and the members of her lab—especially Patrick Straney, Michael Hartmann, Chris Andolina, Ashley Smith, and Zachary Eddinger—have always met my questions with open minds and the genuine, intense curiosity that philosophers of science can too easily overlook as the primary motivation of many practicing scientists. I have learned so much just by showing up to listen to them puzzle through their research together, and I am so grateful to have been afforded the opportunity. All the rest of our work together—teaching philosophy of science to scientists, learning the details of microscope operation and the travails of instrument upkeep, developing the Nanosynthesis Wiki—has been gravy.

As for Bob Batterman, his tireless efforts to get me to put the science first and make the arguments I want to make, rather than framing my work as a response to existing debates, have shaped the dissertation I've put forth in these pages and allowed me to become the kind of philosopher I want to be. His own ability to make a career out of thinking hard about interesting things has convinced me not to set down problems I find just because no one else has already written about them yet. By getting interested in the conundra I brought to our meetings, he convinced me they might just be interesting. He has been a better mentor, and a better friend, than I ever expected to find in graduate school. Long ago, I lost count of the number of tears I shed in my frustration that he wouldn't just let this be another survey of philosophical theories of explanation—and longer ago, I lost count of the number of beers I owe him. I've run out of fancy ways to try to say it, so Bob, from the bottom of my heart, thanks. I shudder to think where I'd be without you.

Finally, a huge thanks to the people, places, and things that got me through the research and writing process: my friends and family, the extraordinarily supportive Pitt HPS graduate community, hordes of helpful colleagues—especially Michael Weisberg, Robin Hendry, and Hasok Chang—and constructive conference audiences, the National Science Foundation's Graduation Research Fellowship Program (Grant #1247842), the intrepid editors of the LATEX wiki, my inspiring students at Pitt and TakeYoga, the folks at Caustelot Farms, my neighborhood purveyors of caffeine and vegetables, the producers of Apple electronics, the existence of mountaintop views, and my teachers of composition, philosophy, science, yoga, and patience. Very special thanks for my continued (relative) sanity to Jeffrey Sykora, Katie Tabb, Samantha Finkelstein, Bruce Bursten and Carol Edelstein, Elizabeth Gibson, Karen Jones, Jessica Kalhorn Williams, Leigh Ann Handel, the members of STARS, the baristas and my fellow dissertation writers at Commonplace Voluto, and to Linus and Junior. I love you all.

I. SMALLER THAN A BREADBOX: SCALE AND NATURAL KINDS

I argue that nanoscience, as a science that is both *scale-dependent* and *synthetic*, challenges a variety of traditional philosophical account of how science works. I introduce readers to the classification of nanomaterials, which both poses problems for traditional theories of classification and familiarizes readers with the basic structures of nanoscience. I show that solving the problem of nanomaterials classification requires borrowing concepts, models, and theories from a variety of larger- and smaller-scale sciences. This chapter provides a foundation for the dissertation's central argument, which is that philosophical views of scientific methodology need to change in a variety of ways to accommodate (1) scale-dependence and (2) the structure of theories and models in nanoscience and other synthetic sciences.

A. A PUZZLE OF KINDS

Suppose I have a collection of a few billion billion gold atoms. If I group all the atoms in the collection together into one massive lump, the result looks yellow and shiny. I can heat this lump to 1064° C, and it will melt. With the appropriate metalworking tools, I can roll the lump into the shape of a ring or smash it into the shape of a coin. The electrons in the atoms in this lump are delocalized, flitting among many nuclear centers. I can expect this lump to resist corrosion and rusting, and to be more chemically inert than similarly-proportioned lumps of metal made up of different metal atoms.

However, if I rearrange the atoms in my lump into clusters of a few hundred atoms each, all these properties change. In order to preserve this arrangement of atoms in my collection I will have had to have suspended the clusters in a carefully-prepared solution, designed to fight the forces of surface tension, chemical kinetics, and thermodynamics so that the clusters will not coagulate back into one massive chunk. The individual clusters in this solution have physical properties, such as malleability and ductility, that differ from the physical properties of my lump and which change depending on the size, shape, and surface chemistry of the clusters. As a group, the clusters appear differently-colored than my shiny, yellow lump: out of solution the group of clusters has the dull, reddish-black appearance of river mud. In solution, the clusters turn a previously-clear liquid red or brown, a color brought about not by normal pigmentation but instead by a rare optical phenomenon known as localized surface plasmon resonance (LSPR¹). And finally, unlike my big lump of gold, these clusters are far from inert. Instead, they are catalytically active—that is, they can be used to speed up other reactions that are performed in my solution.

At least one further case is possible, namely that my atoms are neither a lump nor a collection of clusters, but are, simply, atoms in a vacuum, with all and only their atomic properties: 79 protons, 79 neutrons and electrons, some electronic symmetries and characteristic dispositions toward bonding, but without any macroscopic properties whatsoever—no color, no temperature, no malleability or ductility. All this is to say that the macroscopic kindhood of my collection of gold atoms is *underdetermined* by solely specifying the identity of the atoms involved.

This underdetermination lies at the intersection of a variety of debates in philosophy of science and contemporary metaphysics, including debates about natural kinds, reduction, emergence, and the structure of scientific theories and explanations. It appears to pose an immediate threat to naïve reductionist views about substance identity, such as those in the early microstructural essentialist writings of Kripke and Putnam.² Roughly put, the dilemma

¹This phenomenon is rare because it only occurs when my clusters of atoms are small enough that their diameter is smaller than the wavelength of visible light.Neither smaller-scale nor larger-scale samples of gold can produce the LSPR response necessary to make gold red. The LSPR response disappears around 2-3nm for many preparations of gold nanoparticles, and again around 300 nm. But not only do these particles look different; the difference in their optical response is one of many ways in which gold nanoparticles can perform functions that neither atomic gold nor macroscale gold can. LSPR responses can be *tuned* to respond to a particular frequency of incoming light, and it can be inserted, as a sort of switching mechanism, into other chemical and physical machinery on the surface of a nanoparticle. Similarly, while macroscopic gold is not very chemically reactive, gold nanoparticles can be used as catalysts for other reactions. This catalytic behavior arises as a consequence of the greater percentage of atoms on the surface of gold nanoparticles, relative to macroscale gold. For more details on the mechanics of LSPR, see III.F

²The microstructural essentialist view asserts that kind terms refer as a result of proper uses of the term

is this: either each of my three arrangements is a member of the same kind, in which case that kind class fails to pick out the macroscopic properties of any of the samples, or each of the arrangements is a member of a different kind, in which case atomic number is not the essential identity-conferring property of the samples. And on either horn of the dilemma, the question of the connection between microstructural similarity and macrostructural difference still remains.

My three samples in this case are not the only examples of microstructural similarity underdetermining macroscopic properties. For instance, Batterman (2012) has recently introduced the literature to a case of a composite material made of a mixture of conductive and non-conductive substances. When "islands" of non-conductive material are mixed into a "sea" of conductive material, the result is conductive, and *vice versa*. In Batterman's composite, as in my three arrangements of gold atoms, further information about higher-level structural features of a sample are needed to determine the macroscopic properties of the material. And the question remains: where in the move from microstructure to macrostructure should one draw the borders of kindhood?

In this chapter, I present a somewhat deflationary answer to this question, arguing that kindhood should be indexed to particular theoretical contexts and that changes in the scale of the objects of a scientific investigation indicate changes in theoretical context. Further, this indexing is not a result of mere pragmatic convention, but instead a consequence of a rather unassuming intuition that the structure of scientific concepts and theories should follow the structure of the systems they aim to represent, reference, and understand. That is, it is a scale-dependence in the systems themselves that leads to scale-dependent changes in theoretical context. For my three samples of gold, the upshot is that I permit each of the samples to persist as a kind unto itself, but each is a different kind of kind than the next. Each kind is useful in different theoretical settings because each picks out a kind whose properties of interest exist *at a different scale* from the others.

Physicists already have some semantic infrastructure to discuss how properties can group

picking out a set of things in the world with a shared, identity-conferring essence (essentialism), and that what that essence is is shared microstructure. "Gold is the thing with atomic number 79," and "Water is H_2O ," are both statements of microstructural identities of macroscopic kind terms, made famous by Kripke and Putnam, respectively.

across different scales: the *characteristic length*, time, and energy scales of a system or object of theory or experiment pick out precisely this notion (cf. e.g. Phillips 2001). And recently, philosophers of physics have picked up on this infrastructure, building an intersection between physical discussions of characteristic scales and philosophical discussions of reduction, emergence, idealization, and explanation (cf. e.g. Batterman 2003; Wilson 2012; Butterfield and Bouatta 2014). But their discussions have yet to result in an application of the concept of a characteristic length or time scale to debates about kindhood and substance identity; that fork of the road simply has yet to be taken. This chapter is a first step down that path. In Section I.B. I sketch the problem of natural kinds within the context of contemporary philosophy of science, showing that while explicit discussions of kindhood have largely been confined to individual sciences (notably chemistry, biology, and psychiatry), the worries that underlie the natural kinds debate are in the background of a number of recent and celebrated accounts of the structure of scientific theories and the shape of scientific progress. I focus in particular on Chang's Inventing Temperature (2004) and Wilson's Wandering Significance (2008). In Section I.C, I review an account of chemical kinds I have given in (Bursten 2014) in order to pick up some relevant basic notions of the connection between microstructure, macrostructure, properties, and kindhood among chemical kinds. Section I.E adapts these notions to the problem of nanoscale kinds and shows why standard accounts of chemical kindhood either fail to account for the difference between the lump of gold and the collection of nanoclusters, or fail to pick out their relevant similarities. I argue that in order to carry out the projects of nanoscience, scientists need both to be able to distinguish the former case from the latter case and to draw a connection between the two cases. I then present a sketch of a scale-dependent theory of kinds that permits and explains these two complementary relations. Section I.F contains brief summary remarks.

B. THE METHODOLOGICAL PROBLEM OF KINDS

A useful if oversimplified way of understanding the contemporary natural kinds literature is to divide worries about kindhood into two related but distinct categories: the metaphysical problem and the semantic problem. Metaphysical worries about kinds concern questions about naturalness, ontology, and essences: Do natural kinds have essences? Can the possession of essences, or other criteria, be used to separate natural from conventional kinds? What sorts of things are kinds, if anything at all? Can kinds or essences change over time? Semantic worries about kinds, on the other hand, primarily concern questions about kind terms: How is the reference of a kind term fixed? What happens in cases where a kind term's extension is underdetermined by its use history? Can the extension of a kind term change over time? Often, these two sorts of worries are combined: Are essences what fixes the reference of kind terms, or is it something else?

The aim of this section is to show that this division of worries about kindhood into metaphysical and semantic worries is inadequate in philosophy of science. In philosophy of science, the metaphysical problem is often answered with a quietistic rejoinder, by listing cases where classification continues to occur successfully in the absence of a specification of essence, or even in the absence of an answer to the essentialism question. This response can be seen especially in projects in the philosophy of biology, where kindhood is argued to be conferred by the history of a species or phylum, rather than by a shared, eternal essence (e.g. by Hull 2010). Even in cases of stereotypically essence-based kinds, such as *qold* and *electron*, kind essences are not indisputably accepted.³ Likewise, the semantic question is often reframed in terms of the ability of scientists to refer to individual substances or species over time, coupled with stories of how realization of failures of reference (e.g. "phlogiston," "brontosaurus") led to scientific advance. So while there are many interesting responses to metaphysical and semantic questions in philosophy of science, a third category of worry about kinds underlies most writing about kinds in philosophy of science. I call this worry "the methodological problem of kinds." It is the endemic problem of kinds for philosophy of science.

Surveying discussions of kinds in philosophy of science reveals a recent and somewhat tentative consensus: whether or not kinds are natural, and whether or not they have essences, and whether or not the extensions of kind terms change on Pluto or Twin Earth, what kinds

³For instance, see (Boyd 1999) or (Needham 2011) for discussions of the absence of microstructural essences in chemical and other scientific kinds.

are are tools that scientists use to carry out the practice of science. From the problem of species classification in biology to the development of the Standard Model of Particle Physics, kinds are useful to science insofar as they can help scientists to determine how to proceed in carrying out an experiment, confirming a theory, or explaining a set of natural phenomena. The basic motivation behind theories as diverse as Boyd's homeostatic property clusters and Hacking's pragmatic theory of kindhood is a shared one, namely to determine how classification aids the practice of science.

This motivation for studying kindhood through a philosophy-of-science lens also appears in philosophical discussions of kinds in contemporary metaphysics and epistemology more generally, such as Millikan's *On Clear and Confused Ideas* (2000). While it is clear from the quote below that Millikan does not endorse a pragmatic approach to kindhood as a path to understanding substance concepts, the acknowledgement of this function of kinds in science is nonetheless apparent:

When confidence is lost in the reality of a substance or in the univocity of a substance term, it may begin to be used in a strictly classificatory way. For example, terms for many mental disorders have vacillated over the years between being understood as capturing substances, naming single diseases for which single etiologies and therapies might eventually be discovered, and as being merely classificatory, defining useful groupings of symptoms for efficient transfer of information. (Millikan 2000, p. 39)⁴

This motivation for a philosophical investigation is what I mean by the methodological problem of kinds. In much of philosophy of science, the semantic question is reframed as part of a larger project to provide a responsible descriptive characterization of the apparent success (and failure) of science to develop theories, models, concepts, and explanations. Consider, for instance, the following two examples. First, the account developed in Chang's *Inventing Temperature* (2004). The book traces the evolution of the concept *temperature* during the revolutions in scientific measurement capabilities of the 18th and 19th centuries. Chang demonstrates how precisifying what was meant by "temperature" and associated terms (e.g. "melting point," "boiling," and, interestingly, "comparability") was an iterative process of meeting conceptual ambiguities with changes in measurement technology and

⁴It is worth noting that some recent work on kinds in the philosophy of psychiatry (e.g. by Tabb 2015) has addressed just this question of the value of classifying mental disease and symptom categories.

meeting increased measurement capabilities with new conceptual queries. By way of analyzing this historical process, Chang introduces his theory of *conceptual extension* (Chang 2004, pp. 141–158). His theory offers a mechanics of conceptual evolution and refinement that relies on multiple instrumentation techniques agreeing on a value ("mutual grounding," Chang 2004, pp. 156–158) as a predominant method of conceptual extension. In this account, the uses of "temperature" and its affiliates are not considered for the sake of semantics, but instead in order to understand how clarifying the referents of such terms aids the development of scientific practices and contributes to the structure of scientific theories. Chang has broadened this conceptual-extension project to include accounts of "phlogiston" (2010), "acid" (2012a), and "water" (2012b), among other concepts. These projects *contain* responses to the semantic problem, but they *are* solutions to the methodological problem: mutual grounding and Chang's accounts of the usage-histories of kind terms are responses to the question of how scientists use kind terms to accomplish their projects.

Second, consider Wilson's account of concepts from *Wandering Significance* (2008). One of the central positive projects of the book is to articulate a theory of how science influences the development of concepts and *vice versa*. Wilson develops an account of scientific theories in which the conceptual architecture of a theory is much less cohesive than it might, on first glance, appear. He calls this view the "theory façade" account. In this account, theories are made up of "patches" in which a concept is well-defined over a given domain. Domains are individuated by reference to various uses to which the concept can be put and various parts of the world that the concept reliably picks out (if questions are asked of the concept in the right way). To understand Wilson's account as a response to the methodological problem, it helps to focus on one feature of the ways in which Wilson defines patches, which can be seen most clearly in his discussion of hardness:

[[]O]ur usage of the predicate "is hard" displays a fine-grained structure that we are unlikely to have noticed, for our everyday usage is built from local patches of evaluation subtly strung together by natural links of prolongation. More specifically, in everyday contexts we adjudicate the "hardnesses" of various materials, both comparatively and absolutely, through a wide variety of comparatively easy to apply tests—we might squeeze the material or indent it with a hammer; attempt to scratch it or rap upon it; and so on. In most cases, we will be scarcely aware of the exact technique we will have employed for this appraisal: "Did I rap, squeeze or scratch that piece of wood? I can't really remember." In fact, our choice of tests is likely to have been suggested by the material in question: we instinctively

appraise a wood by rapping upon it, a rubber by squeezing, a metal by attempting to make a small imprint; a glass or ceramic by rapping lightly or scratching (not by trying to make a small imprint!). In fact, we are normally interested in comparing hardnesses mainly within *natural groupings of stuffs of generally allied characteristics*, generally metals with metals, ceramics with ceramics and so forth, although interesting crossover cases also arise. (Wilson 2008, p. 336, emphasis added)

Wilson goes on to examine the variety of measures for hardness across various material domains, concluding that,

our employment of "hardness" silently distributes itself into a patchwork of sheets, locally distinguished by a certain vein of probing (scratching, tapping, etc.), that sit over various varieties of material stuffs and continue smoothly into one another. ... As an inferential tool, most of us employ the term "hardness" informally as a generic, single-scale "folk physics" parameter that we consult in selecting a material for, e.g., manufacturing purposes. Left at this general level, the term requires further specialization before it can carry much data usefully. Accordingly, over each localized domain of related stuffs, the "hardness" parameter is likely to silently specialize on the forms of probing that prove most informative with respect to the materials and purposes locally at issue. (Wilson 2008, p. 338)

This discussion suggests that the "patches" that comprise the conceptual structure of scientific theories often arise as a consequence of the tendency of members of certain higherlevel kinds (e.g. metals, as opposed to gold) to exhibit similar properties or behaviors—that metals are all susceptible to scratch tests is both a fact about how the concept *hardness* behaves around that kind of materials, as well as a fact about the class of materials themselves. To put the point in slightly different terms, the "natural groupings of stuffs of generally allied characteristics" that make up kinds of (at least) scientific interest are kinds because they exhibit shared properties of interest to science.⁵ In order to explain how classification aids the practice of science, it becomes necessary to formulate an account of how these groupings of properties manifest, how they are measured, whether they can be controlled, and what explains them. These are the sorts of questions the methodological problem of kinds addresses.

Reconsider the puzzle that began this essay in light of the methodological problem. The three cases of gold atoms—the macroscopic lump of gold, the collection of gold nanoclus-

⁵In this light, then, the groupings of properties to which Batterman, Butterfield, and others refer in their discussions of "universality classes" are really just a special case of scientific kinds.

ters, and the collection of gold atoms—exhibit different groups of properties. In fact, the properties of the different cases are so different that it is difficult even to develop a basis of comparison for properties they share: the macroscopic lump is shiny, yellow, malleable, ductile, inert, has a melting point, not to mention certain secondary economic properties, and so forth. On the other hand, the gold nanoparticles are catalytically active, exhibits an LSPR response, appear in collections as a reddish-brown-hued solution (or, when dried out, as a nondescript dark stain in the bottom of a test tube), and are necessarily locked into relationships with various stabilizing chemicals. Finally, the collection of individual atoms, which is admittedly a wholly theoretical but conceptually useful potential arrangement, exhibits only the properties of individual gold atoms, namely electronic configuration, orbital symmetries, quantity of neutrons and protons, etc. It is not only that properties such as color change their value (from yellow to brown) in the move from one case to another, but also that some properties, such as temperature, LSPR response, and orbital symmetry, *only exist at certain scales*.

I have shown that the methodological problem of kinds, namely the question of how classification aids the practice of science, is answered via discussions of the groupings of properties that scientists use to classify, that is, discussions of the measurement, organization, explanatory power, and apparent universality of properties across a kind or class of kinds. In this light, the differences in properties among the three cases of gold atoms leads to two distinct conclusions: (1) the three cases are different kinds under the auspices of scientific investigation, and (2) scale plays a role in delimiting the groupings of properties that demarcate kinds. This means that the methodological problem of kinds requires, as part of its solution, an understanding of the role of scale in classification.

In the remainder of this chapter, I develop an account of how scale influences classification, concentrating on the scientific problem of developing a classification scheme for nanoscale materials. While the role of scale is particularly vivid in nanoscience, my remarks generalize to other areas of scientific enterprise. The reason for this generalizability is itself part of a response to the question of how properties group. In the vast majority of sciences, phenomena, objects, and properties group around characteristic length, time, and energy scales, from the vast expanses of time studied by geology to the sub-picometer regions of space studied by particle physics and the hadron collisions of high-energy physics. Scale is a constraint on the manifestation of groupings of properties, and it is also a rationale for observed interactions between phenomena. It is (part of) the reason that things can be grouped together along property lines at all; because properties are scale-dependent, comparisons of property groups across widely divergent length, time, and energy scales are, in a certain sense, conceptually incommensurable. For instance, consider the project of comparing a mosquito to a mountain: other that being made of chemical elements native to Earth, it is not clear what sorts of properties either object has in common.

The title of this chapter draws from the common car-trip game 20 Questions, in which one player thinks of an object and other players take turns asking yes-or-no questions (up to 20 in all) to figure out what the object is. Whoever guesses correctly first wins, and if no one gets it, the player who thought of the object wins. One of the most common and, it is generally agreed, most effective starting questions for the game is, "Is it smaller than a bread box?" The utility of this question at partitioning the possible objects in a 20-Questions game illustrates just how powerful a role scale plays in underwriting classification.⁶

C. CHEMICAL KINDHOOD: THE STRUCTURE–PROPERTY PARADIGM

Chemistry is a science that exists across multiple characteristic length scales, develops conceptual infrastructure native to each scale, and yet can connect phenomena at one length scale to properties at another length scale. Chemical reactions, for instance, are generally taken to be macroscopic phenomena—the result of pouring the contents of a test tube into a beaker, to caricature—but they are frequently explained in terms of microscopic changes in the molecular structure of the reactants. So it is natural to wonder whether chemical kinds, both infimic kinds such as the individual chemical elements and 'higher-level' kinds such as acids and metals, are microscopic or macroscopic.

⁶Of course, the bread-box question is only actually effective when one knows about how big a bread box is. For those whose childhoods included neither 20 Questions nor bread boxes, perhaps a better partition here would be "Is it smaller than the 10-volume *Routledge Encyclopedia of Philosophy*?"

Chemical reactions occur or fail to occur because of structural features of the substances involved—a reaction is a pattern of breaking and formation of chemical bonds, which defines the structure (molecular geometry) of a chemical substance. The gold nanoparticles we have been considering here participate in different chemical reactions than macroscopic lumps of gold, such as the former's ability to act as catalysts for other reactions. Again, where macroscale samples of gold are known for their non-reactivity, gold nanoparticles can be used to speed up or increase the yield of chemical reactions. For instance, it has recently been shown by Zhu et al. (2011) that small, gas-phase gold nanoparticles can be used to catalyze the oxidation of styrene, the molecule that forms the basis of styrofoam cups.

I have previously argued (Bursten 2014) that chemical kinds are individuated macroscopically according to a criterion of similarity of chemical reactivity, but that this similarity is explained by reference to sameness of microstructural change. My discussion of why this reference was successful was ontologically minimal, focusing primarily on individual examples of successful application of microstructural rationale to observed changes in the macroscale substance. I showed that changes in the arrangement of atoms in a substance or compound, which I denoted molecular geometric properties, were necessary and sufficient conditions for changes in chemical reactivity.

The aim of that study was to preserve the intuitive appeal and explanatory power of appeals to microstructure in chemical settings while eliminating the need to justify the connection between macrostructure and microstructure through a microessentialist metaphysics. The bulk of the account was an analysis of examples that philosophers have used in the past—including jade, allotropism, and water—to challenge various essentialist and anti-essentialist theories of kinds, with commentary on how the microstructural rationalization of macrostructural features makes sense of the apparent threat each example posed. What the account revealed is in fact a methodological operating principle known in chemistry as the *structure-property paradigm*. The structure-property paradigm is an heuristic principle under which much of chemistry operates. Its central premise is that changes in microstructure correlate to changes in properties of interest, on both the microscopic and macroscopic scale. That is, changes in macroscopic properties can be explained in terms of changes in microstructure, as when annealing a metal to change its crystal structure leads to increased hardness, or as when adding an electron to the outer valence shell of a halogen element (e.g. fluorine) changes its electronic symmetry (and, consequently, its reactivity). The structure–property paradigm guides not only the design of individual experiments in chemistry, but also, as I show in the next section, the development of new classification schemes for never-before-synthesized materials.

Philosophers such as Hendry (2006) and Goodwin (2012) have previously argued from chemical examples to versions of the structure–property paradigm as the appropriate framework for understanding chemical theory and chemical explanations, so I will rehearse only a few brief examples here:

- Metals are electrically conductive because their electrons are delocalized, a feature that is both itself microstructural and which has a further microstructural explanation: metals have "empty spaces" in their outer orbital shells, which creates the space for delocalized, or free-flowing, electrons.
- Acids react with bases to form salts and water because of microstructural features of acid and base molecules: acids contain donor protons (in Brønsted-Lowry acid-base theory) or the appropriate microstructure to receive lone-pairs of electrons (in Lewis acid-base theory), and *mutatis mutandis* for bases.⁷
- The compound thalidomide was used as a sedative and anti-nausea drug for pregnant women in the late 1950s, but was later pulled from the market for causing birth defects in children born to thalidomide-using mothers. Scientists later determined that the drug was made of two different *stereoisomers*, that is, two different structures where one structure is a non-superimposable mirror image of the other (for quick reference, your right and left hands are stereoisomers of one another). It was later determined that one isomer was responsible for the sedative and anti-nausea effects of the drug, and the other isomer was responsible for the children's birth defects.

In each of these examples, some aspect of chemical microstructure is used to explain either microstructural or macrostructural properties of the substance or compound in question. This is all the structure–property paradigm is. It is not meant to be mysterious or counter-

⁷For a richer perspective on the definitions of acids and bases, see (Chang 2012a).

intuitive; in fact, it is nearly so obvious as to be overlooked. In its simplest formulation, the structure–property paradigm is just the belief in the reliability of explanations of (changes in) both macroscopic and microscopic properties in terms of (changes in) microstructure. By itself, it does not require any particular ontological or epistemological grounding, though depending on one's views about essences or warrants for belief, one may care to back it up with a more robust ontological or epistemological theory.

The account I developed to capture the structure–property paradigm, while true to the minimalist ontology and epistemology of the framework, did not go far enough in drawing attention to the fact that the relationship between structure and property is one of the things chemists *aim to explain* when they sort the products of their reactions into kinds. What I mean is this: chemists work with kinds at a variety of scales, and each of these kinds is a grouping of properties of interest to chemists. At the macroscopic scale are chemical substances, both individual substances and compounds (e.g. gold, table salt) that may be reidentified by their reactive properties (e.g. inertness, solubility in water). At the microscopic scale are molecular geometries, which are architectural arrangements of atoms, such as the familiar mickey-mouse shape of a water molecule, the double helix of DNA, or the zig-zag chain of carbon atoms in the hydrocarbons that make up fossil fuels and natural gases. At still smaller scales are the symmetries of atomic electrons, which dictate constraints on the ways in which atoms can combine to form molecules.

Each of these characteristic scales of chemistry is accompanied by its own conceptual infrastructure; that is, its own kinds. "Gold" is not a useful kind when articulating the properties of orbital symmetry groups (i.e. presence or absence of an inversion point, number of permitted rotations and reflections, etc.), and neither is "point group" for articulating the properties of macroscopic gold (i.e. color, malleability, ductility, etc.). When I argued that microstructure is an explanatory basis for chemical reactivity, and that this licenses references to microstructure without requiring essentialism, what was missing from the account was the following: one of the central *purposes* of classification in chemistry is to develop explanations of just such a "scale-bridging" sort, and to connect conceptual infrastructures from different scales by means of experiment, observation, and control of the manifestation of properties at multiple characteristic length scales. This aim of chemical classification in fact helps to explain and rationalize the effectiveness of the structure-property paradigm.

It is possible to illustrate the influence of scale on chemical concepts and on the role of the structure–property paradigm in chemical explanation through common chemical examples of the sort itemized above. However, the role of scale is all the more striking when one considers the changes in structure and properties that arise when materials shrink from the bulk scale to the nanoscale. The three cases of gold illustrate that there are times when there is more to structure than simply specifying composition. Even when specifying the composition of a chemical substance or compound encodes information about structure, such as the association between "water" and the polar, mickey-mouse molecules of H_2O , nanoscience reveals that there is more to structure than composition and geometry. In the next section I argue that to preserve the effectiveness of the structure–property paradigm in nanoscience, one need consider not only the chemical composition, but also the size, shape, and surface chemistry of materials under consideration. Only once these features are all specified is it possible to justify classifying nanomaterials according to differences in properties of interest. This shows that classification of nanomaterials differs from classification in other areas of chemistry as a result of differences in the scale of the materials being classified.

D. WHAT ARE NANOMATERIALS?

Nanomaterials⁸ are small collections of regularly arrayed atoms that display material properties not seen at either larger (macroscopic) or smaller (molecular) scales. They can be defined structurally, in terms of the number of atoms or size of the particle; or they can be defined functionally, in terms of *scale-dependent behaviors* that they display. Structurally, nanomaterials are not large enough to be counted as macroscopic objects or instantiate many properties associated with macroscopic materials, nor are they small enough to be strictly molecular compounds. Nanomaterials can contain anywhere from a few dozen to a few hundred thousand atoms. Some small nanomaterials contain only 13 atoms⁹—fewer

 $^{^{8}}$ I follow the convention of the nanoscience community here by using "nanomaterial" to refer to any object or system with nanoscale features, and "nanoparticle" to refer only to 0D or short 1D nanomaterials.

⁹Such as the Au_{13} nanoparticles first reported in (Schmid et al. 1999).

than many molecular materials—but display definitively nanoscale properties, and so are counted as nanomaterials. In other words, while structure and size are important, ultimately what makes something a nanomaterial is functional: namely, its ability to exhibit a scaledependent difference in behavior, relative to macroscopic materials or molecular substances of similar elemental composition.

Many molecular materials exist at the same orders of magnitude as nanomaterials, but the properties of molecular materials differ significantly from the properties of nanomaterials. Most notably, the bonding in molecular materials such as proteins is highly directed, and there are few repeated unit-cell structures within the material, whereas most nanomaterials are constructed out of regularly repeating geometric arrays. Additionally, many molecular materials have complex material compositions—that is, they are composed of highly specific ratios of a wide variety of chemical elements. Most nanomaterials, on the other hand, are made of one or two elements, and the ratios of these elements can vary. In fact, varying the ratios of different elements in a nanomaterial is one way that synthetic chemists aim to control the behavior of the materials.

The relation between size and behavior will become a central theme in later discussions. Nanoscale material behaviors can be broadly understood as falling into one of two categories: 1) modifications of known macroscopic material properties, and 2) properties that exist only at the nanoscale. Examples of the former include scale-dependent changes in the strength, color, and conductivity of a material, and examples of the latter include quantum confinement and LSPR. Both these sorts of differences are discussed further in Chapter II.

There are two distinct but related kinds of differences between nanoscale and macroscale systems. The first difference is that *physical and chemical phenomena change at the nanoscale* relative to the macro-scale. The second difference is that *different tools are needed to model, explain, and predict the behavior of nanoscale systems* than are needed for macroscale systems. One of the aims of this dissertation is to clarify the relationship between these two differences through detailed investigations of the theories and models that scientists use to understand and make nanomaterials.

1. Classifying Nanomaterials

One of the central problems nanoscientists face is the classification of their materials. Nanomaterials are not necessarily molecular, so two samples of "the same" material could contain different numbers of atoms, different arrangements of chemical bonds, or different surface chemistry. In other words, there is no equivalent to the periodic table of elements, or the standard model of particle physics, for nanomaterials. Figuring out how to sort materials into useful categories is still an active project in nanoscience, and it is one that must be solved by attending first to the aims of classification.

Ultimately, the goal of synthesizing new nanomaterials is to produce new material structures and behaviors that can be controlled, reliably reproduced, and put to use in various scientific and non-scientific projects. The classification of nanomaterials must reflect these goals. To that end, I have worked with the Millstone Group to develop the following classification scheme, which individuates materials in a manner that provides information both about the behaviors the materials can be expected to exhibit, as well as about the methods by which the materials can be synthesized and their existence confirmed by analysis.

This classification scheme can be seen as a set of conceptual distinctions that are unique to nanoscale materials. While it may sometimes prove useful to know all of the following information about macroscopic materials, it is often unnecessary for the identification of a material or for determining whether it is usable in a given environment. Conversely, without specifying each of the following attributes of a nanomaterial, a scientist cannot determine whether the material is compatible with a given environment or problem setting.

The four aspects of a nanomaterial that require specification under this classification scheme are:

- a. composition
- b. dimension
- c. shape
- d. surface chemistry

Controlled changes to one or more of these four aspects of a nanomaterial can produce all the interesting changes in the material's behavior that the nano scientist may seek. While each of these characteristics of a nanomaterial is structural, changes in a particular characteristic are systematically related to changes in the properties or behaviors that the material exhibits.

a. Composition Composition answers the question *What element or elements make up the nanomaterial?*. Here it is useful to borrow terminology and categories from standard chemical literature. Many nanomaterials are *inorganic*, meaning that they do not contain carbon–hydrogen bonds. It is worth reiterating here that most nanomaterials are made up of regular, repeating unit structures, which is a feature of chemical composition that is vastly more common in inorganic compounds than in organic compounds.

Inorganic nanomaterials divide into two broad material types: *carbon* and *metal*. Carbon nanomaterials are, perhaps obviously, either composed solely of carbon atoms, or gain their main structure from carbon atoms and contain only a few other elements in specific parts of the nanomaterial¹⁰. Interestingly, nearly all carbon nanomaterials have a similar bond structure, or geometric arrangement of atoms relative to one another within a given material sample.

Carbon nanomaterials can in general be described as variations on two-dimensional topologies: the thin film graphene is a flat Euclidean plane, a single sheet of hexagonally-bonded carbon; buckminsterfullerene (C_{60}) is spheroidal—a two-dimensional surface embedded in three-dimensional Euclidean space—and carbon nanotubes are cylinders, which are again two-dimensional surfaces embedded in Euclidean 3-space. Even multi-walled carbon nanotubes are simply multiple concentric cylinders. Metal nanomaterials, by contrast, cannot be described in terms of two-dimensional topological variations on a single basic bond structure.

Metal nanomaterials display a wide variety of bond structures, as well as a variety of elemental (in the periodic-table sense) compositions. Some metals form nanomaterials more readily than others. Most contemporary metal nanomaterial research is devoted to the formation of nanoparticles from the "noble" metals—gold, silver, platinum, ruthenium, osmium,

¹⁰For example, it is possible to "trap" small molecules inside a buckminsterfullerene frame.

iridium, rhodium, and palladium—and the transition metals more generally.¹¹

Metal nanomaterials can be made of single metals or of combinations of metals. Of the latter, the metals can combine in three distinct ways. *Janus particles* are composed of distinct halves, one half of one metal and the other of another. *Core-shell* particles have one metal in the central "nucleus" of the particle and another metal surrounding. *Alloy* or *intermetallic* particles are composed of multiple metals approximately evenly distributed throughout the particle. The crystal lattice type of a given metal nanomaterial is usually similar to the crystal lattice seen in macroscopic samples of the same material; for instance, macroscopic solid gold generally forms into cubic-close-packed arrangements, and so do gold nanomaterials. However, there are notable differences in bond structure more broadly speaking between macroscopic metals and metal nanomaterials, which will be discussed in Chapter II.

b. Dimensionality Nanomaterials can be described as either 0D, 1D, or 2D, according to the number of dimensions in the material that *exceed* the nanoscale (approximately 1–100 nm). For instance, small spherical particles like fullerenes have no dimensions above the nanoscale and are thus 0D; nanowires and nanotubes have one dimension above the nano scale and are thus 1D; and thin films like graphene have two dimensions above the nano scale and are thus 2D.

The 1–100 nm range is a useful point of reference for determining a material's nanoscale dimensionality. It is not hard and fast. The more important consideration is the relative size of each dimension of the nanomaterial. If the particle is approximately the same size in each dimension, it is 0D. If it is elongated in one dimension, it is 1D; if it is elongated in two dimensions, it is 2D. So there exist 0D nanoparticles that are 175 nm along each side (e.g. Wiley et al. 2005) and 1D nanoparticles where two dimensions are around 10 nm and the third dimension is around 60–90 nm (e.g. Sau and Murphy 2004). Again, this flexibility of dimensional assignments highlights the importance of understanding that nanomaterials are

¹¹Metals from the early (leftmost) groups on the periodic table, namely the alkali and alkaline earth metals, do not readily form nanomaterials. These elements do not commonly exist as pure substances in nature, because their electronic structure is relatively unstable in elemental form and is, consequently, readily stabilized by the formation of compounds (e.g. NaCl, common table salt). In other words, the disposition of noble metals to form nanomaterials more readily than early-group metals is *not* a scale-dependent feature of the metals, but rather a consequence of electronic behavior at any scale.
first and foremost materials that *have nanoscale properties or engage in nanoscale behavior*, and that identification of the dimension of the material is done in the service of understanding the material's behavior. It is not the size that makes the nanomaterial: it is the size that makes possible the novel properties, and it is the novel properties that make the nanomaterial the kind of system synthetic scientists are interested in making.

Dimensionality often plays a central role in constraining the behavior of a nanomaterial and, in turn, the uses to which a nanomaterial can be put. For instance, 2D materials can be used to coat the surfaces of macroscopic objects, whereas 0D materials are more often spread, via suspension in a solute, throughout the body of an object. And electric currents and propagating plasmons can be transmitted over macroscopic distances along the elongated dimensions of 1D and 2D nanomaterials, but not along 0D materials. Conversely, only 0D materials can exhibit LSPR.

c. Shape Within each dimension, a wide variety of geometries can occur. Beginning with 0D particles, there is a basic division between *isotropic*, or spherical, particles and *anisotropic* particles. Isotropic particles are relatively more stable than anisotropic particles, and consequently they are easier to synthesize. Most syntheses of anisotropic 0D particles are achieved by modifying or building on to isotropic particles. Anisotropic particles can be compared for *relative* isotropy: for instance, a cube is more isotropic than a long, thin wire.

Anisotropic 0D nanoparticles come in a variety of geometries. Many are polyhedra, commonly cubes (e.g. Gräf et al. 2006), octahedra (e.g. Tao et al. 2006), and icosahedra such as C_{60} (Buckminsterfullerene). Another common shape family is the class of plate-like particles with distinct edge symmetries, such as triangular (e.g. Wiley et al. 2005) and hexagonal (e.g. Yang et al. 2007) plates. Finally, 0D anisotropic nanoparticles can also take on more exotic geometries like the "inukshuk" particles reported in (Aizawa et al. 2005).

1D nanomaterials are all elongated in one dimension, as discussed above. But their shapes can vary within this geometric constraint. The particles can be either solid or hollow. The cross-section of 1D nanomaterials can either be circular (for tubes, rods and wires) or polygonal (for prisms and belts). Finally, the termini of the nanomaterial can either be convex, concave, or neither. Differences in each of these features lead to differences in the behavior of the nanomaterials and differences in the ability of the nanomaterial to function within larger systems.

2D nanomaterials are, by contrast, relatively simpler to classify by shape. They are all thin films, and the perimeter shape of the film itself does not have much bearing on the behavior of the material. However, film thickness can vary greatly. This variation can give rise to changes in material properties, due to variations in electron transport phenomena as the path length from the bottom to the top of the film increases. Many optical and electronic properties of materials are affected by variations in the thickness of a film.

d. Surface Chemistry Nanoscale surfaces are unlike the surfaces of macroscopic materials, both in terms of their role in constraining and shaping the properties and behaviors of the material, and in terms of the modeling strategies used to describe, predict and explain what happens at the surface of a nanoscale object. I will discuss strategies for modeling nanoscale surfaces at length in Chapter II and implementations of those models in synthesis projects in Chapter III. For now, I want to focus on a few very basic facts about the surfaces themselves:

I. Surface atoms make up a measurable proportion of the atoms in the material. This is a difference between nanoscale and macroscopic materials. While surface atoms make up a negligible fraction of the atoms in a macroscopic material, even one as small as a human hair, surface atoms account for significant proportions of the atoms in a nanomaterial. See Figure II.4 in Chapter II for the trend in ratios of surface atoms to overall atoms in a material.

This difference between nanoscale and macroscale materials implies that both the nature of atomic bonding and some macroscopically observable properties change between macroscale and nanoscale materials.

II. No nanosurface is an island. Nearly all metal nanoscale surfaces are bound by capping agents. This is the most important feature of nanosurfaces, in terms of understanding the chemical activity at the surface of the material. Capping agents can affect the catalytic behavior of the material, the morphology and dimensionality of the material, the stability and persistence of the material over time, the kinds of materials with which

a given nanomaterial can interact, and any number of other properties of interest. The effects of capping agents are discussed below.

- III. The energy of a material system is minimized when its surface-to-volume ratio is minimized. The minimal surface-to-volume ratio of a material occurs when the material is spherical. It is worth noting that this point underlies many of the difficulties in synthesizing anisotropic nanoparticles. Energy minimization and surface-to-volume ratios are taken up in detail in Chapter II.
- IV. Surface atoms are bound less tightly to a material than interior atoms. This is true of both nanoscale and macroscopic materials. What it means it that it takes relatively more energy to keep a surface atom bound to a material object than it does to keep an interior atom so bound. Having relatively more atoms in the interior of a material, especially a crystalline material, creates a stronger "pull" on the surface atoms and stabilizes the material.

From these two facts one can see that nanoscale surfaces are relatively higher-energy than their macroscale counterparts. They have relatively fewer interior atoms binding them to the system. As a consequence, nanomaterial surfaces require *stabilization*, which can come in a number of forms. For instance, surface atoms can rearrange themselves into lower-energy configurations. However, this often does not provide sufficient stabilization on its own—or, if it does, it changes the overall morphology of the material. Much more frequently, especially in metal nanomaterials, surfaces are stabilized by the action of molecules foreign to the nanomaterial, which affix themselves to the nanomaterial and affect the surface chemistry and material properties of the material. These stabilizing agents are generally known as surfactants or *capping agents*.

The upshot of these features is that *different capping agents lead to different material behavior*, which is why surface chemistry joins composition, dimension, and shape as a fundamental classificatory feature of nanomaterials. Capping agents can influence both the dimension and morphology of a nanomaterial, as well as its catalytic behavior, its ability to interact with other materials and its environment, and a number of other material properties of interest. Further discussion of how different capping agents affect nanosynthesis and the behavior of nanomaterials is developed in Chapter III.

E. SCALE-DEPENDENCE AND NANOSCALE KINDS

Nanoscale kinds present a sort of double-edged Sorites paradox: when a sample of a material gets too small or too big, it ceases to be a nanomaterial, but the lines between molecular, nanoscale, and macroscale materials are fuzzy at best. It is common for nanoscientists to speak of molecularly-structured, 13-atom clusters of gold as nanomaterials, while there are many materials five or ten times larger than those clusters that get counted as molecular materials. For nanoscientists, while boundaries of size are an acceptable rough first pass, it is really the *scale-dependent properties* of nanomaterials that define them as nanomaterials. These properties are the same ones that raise the question of whether it makes sense to treat gold nanoparticles as *of the same kind* with either atomic or macroscale gold. Extending the structure–property paradigm to nanomaterials, then, requires building up a classification scheme that captures the relevant structural features of nanomaterials that correlate with properties of interest.

Most accounts of chemical kinds, including the one I sketched above, are founded on the idea that composition, that is, elemental makeup, is the foremost grounds for classification in chemistry. For many chemical substances and processes, composition dictates structure, and is thus an essential sortal for chemical properties. The centrality of composition is easily understood by considering the relationship between composition and molecular, as well as electronic, structure. The microscopic composition of a substance is its atoms, arranged into molecular (covalent or ionic) structures. While composition alone underdetermines the structure of a given compound, it *constrains* structure by limiting the available electronic configurations of an atom, which in turn constrains the number, kind, and arrangement of chemical bonds into which the atom can enter. For both "fundamental" chemical kinds (i.e. individual elements) and higher-level kinds (e.g. metals, noble gases, alcohols), composition is the primary constraint on structure.

One of the central differences between nanomaterials and chemical materials writ large is that for nanomaterials, composition no longer takes such a center stage. As discussed above (Section I.D), the relevant groups of properties of nanomaterials, and the relevant structures that correlate to those properties, are underdetermined by the specification of composition alone. While specification of non-compositional properties of macroscopic chemicals, such as phase (solid, liquid, aqueous, gaseous), temperature, and quantity is often useful for understanding reaction dynamics—that is, how quickly and through what intermediate phases a reaction will proceed—specification of such physical properties is inessential to the determination of macroscopic chemical kindhood. Macroscopic gold is still gold, whether solid or molten. But at the nanoscale, one needs to know significantly more about the physical properties and chemical environment of a material in order to determine its relevant reactive properties.

Dimensionality, shape, and surface chemistry all play central roles in determining the physical and chemical properties of nanomaterials. For nanomaterials, specifying composition—even when composition is supplemented by information about microstructure drawn from the structure—property paradigm—is insufficient to specify kindhood. The tetrad of composition, dimensionality, size, and surface chemistry provides a more useful, though perhaps a dauntingly fine-grained, set of classificatory principles. I will comment on the fineness of grain shortly, arguing that it is endemic to nanoscience as a synthetic science. But before I do, I want to step back and consider how this four-part classification scheme fares with respect to the methodological problem of kindhood I defined in Section I.B.

I said earlier that the reason philosophers of science concern themselves with kinds is neither strictly metaphysical nor wholly grounded in concerns about reference; rather, I argued that the motivation for questions about kindhood in philosophy of science stems from a need to develop an account of the role of classification in the carrying-out of science. Whether one takes science to be carving nature at its joints or merely developing empirically adequate, usefully projectable, or inferentially reliable categories, classification is a tool used to help develop scientific concepts and determine whether or not a theoretical principle applies to a given situation. So the fact that the classification of nanomaterials differs from the classification of macroscopic materials, as well as from the classification of molecules, implies that the concepts and theories that make sense of either molecular or macroscopic matter.

While this result is not necessarily surprising in its own right, paying attention to how the interests of classification change across scales can provide insight into the question that arose out of Section I.C, namely, how do macroscopic and microscopic classification schemes in chemistry connect to one another? The emphasis on surface chemistry in nanoscale classification is not repeated at either a higher or a lower scale, which suggests that the transition from microscopic to macroscopic phenomena is more conceptually complex than merely aggregating properties of atoms over a large-number sum. Likewise, the appearance of novel phenomena unique to the nanoscale countermands traditional, reductionist accounts of macroscopic properties of substances being explicable entirely in terms of properties of individual atoms. And, rather than urging an about-face from reduction to its philosophical counterpart, emergence, the appearance of novel conceptual and classificatory structures in nanoscience shows how attention to the movement between scales, rather than to a destination at the very big or the very small, can reveal complex conceptual mechanics.

At the nanoscale, the properties and behaviors that stem from the shrinking scale of the material in question become more important than the properties and behaviors that originate from the elemental identity of the material. Another way to frame to this point is to say that for gold nanoparticles, the change in the scale of the material serves as a change in the properties associated with microstructure. By the structure–property paradigm, this change needs to be reflected in the classification scheme used to categorize nanomaterials. And this change in classification will itself percolate up to changes in the concepts and theories used to describe those materials.

The development of new conceptual infrastructure at the nanoscale is an instance of the scale-dependent introduction of new concepts, theories, and models across especially the physical (but likely also the life and human) sciences. Wilson (2012) addresses shifts in the conceptual structure of a scientific theory that result from changing theoretical contexts by discussing the interrelations between point-particle, rigid-body, and flexible-body mechanics. While he does not frame his discussion explicitly in terms of scale, one of his central concerns is the ways in which larger-scale systems suppress information about smaller-scale behaviors, a process that requires mathematical approximation techniques and variable reduction and which he calls "physics avoidance." Batterman (2001; 2012), on the other hand, has emphasized changing scale as a driving force for changes in conceptual structure, which has led to an ongoing debate (cf. e.g. Butterfield 2011; Butterfield and Bouatta 2011;

Reutlinger 2014) over the essentiality of renormalization-group approximation techniques for developing conceptual understanding of certain physical systems. This debate continues to highlight the fact that many if not most scientific concepts are anchored to particular scales. So the scale-dependence of the classification schemes—and, thus, of the conceptual and theoretical infrastructure—of nanoscience is not unique to the relation between chemistry and nanoscience. Rather, it is simply a particularly evocative instance of a pervasive scale-dependence in contemporary scientific theories.

While the scale-dependent motivation for changing classification schemes is not unique to nanomaterials, there is one difference worth mentioning between the classification of nanomaterials and other scale-dependent developments of novel classification schemes. I mentioned earlier that the tetradic division of structural features of nanomaterials would appear to lead to a particularly finely-wrought taxonomy of nanomaterial kinds, and this is no illusory appearance—a common goal of articles on the synthesis of nanomaterials is to introduce the reader to a supposed new kind of material with new structures and new properties.

Recall that the methodological problem of kinds, which motivates debate about kindhood in philosophy of science, is to determine how classification aids the carrying-out of scientific projects. Given this framework, one can see how differences in the goals of scientific projects can and should lead to different sorts of classification schemes. Classificatory projects in some sciences aim to find a small, finite number of phenomena or objects, such as the 61 particles identified by the Standard Model of Particle Physics or the 92 natural chemical elements in the periodic table. Classification schemes of this sort aid *descriptive* scientific projects, such as the development of dynamical theories. On the other hand, many scientific projects—including much of chemistry and nearly all of nanoscience—aim to synthesize new substances and materials, and thus to generate new kinds to be added to an open-ended classificatory scheme. Even the projectability of the classificatory principle underlying the periodic table (that is, integer increases in nuclear charge) produces an open-ended, though not anarchic, classificatory scheme.

The tetradic classification scheme of nanoscience, which sorts materials by composition, dimension, shape, and surface chemistry, generates a dense network of possible kinds with a variety of interrelations. One can partition the network into regions by any of the four sortal criteria, or by combinations thereof. This map of structural features identifies not only individual nanomaterial kinds, but similarity of structure among higher-level classes of kinds, allowing for comparisons, e.g., of properties among metal nanowires of different compositions or of differences in optical response between gold nano-spheres of different sizes. These comparisons are not merely hypothetical; they are descriptions of current research projects in synthetic nanoscience laboratories (e.g. Sun and Xia 2002; Wiley et al. 2005; Ringe et al. 2010 and Zeng et al. 2011).

This classification scheme also indicates to researchers different ways in which an existing particle may be modified to obtain a desired difference in properties of interest to various applications of nanomaterials. The structure–property paradigm indicates that changes in structure correlate to changes in properties, and as research in nanoscience progresses scientists are developing better control over desired properties by refining their ability to manipulate the structure of synthesized materials.

Such projects are rather far-removed from questions aiming to locate the essence of a newly-synthesized material, and those with heavily naturalist intuitions might worry that the very synthetic nature of such materials belies attributions of essence. And although there is a rather practical worry about consistency of naming of similar particles created by different labs, the question of how reference attaches to newly-synthesized particles is less of a philosophical concern—nearly every new particle architecture is baptized by reference to some or all of the four structural features, e.g. "poly-(vinyl pyrrolidone)-capped silver nanocubes"—than it is a headache for researchers whose particles lie in a nebulous somewhere between "sphere" and "cube." But in order to manipulate and control the properties of newly-synthesized particles, and in order to understand the relation between those properties and the structures of nanomaterials, the four-part classification scheme is an essential tool. This is the point of classification in nanoscience, and consequently, deconstructing the tetrad is a more productive project for philosophers of science than is arguing for the (lack of) essence of nanoscale kinds or projectability of nanoscale kind terms.

F. CONCLUSIONS

In Section I.B I began with a negative project of sorts, arguing that the kindhood debate in philosophy of science is neither about metaphysics nor reference. As an alternative I offered a positive thesis, namely that what interests philosophers of science about kindhood should be, and usually is, what I have called the methodological problem. The methodological problem is the question of how classification aids the practice of science. I have offered a response to the methodological problem that both is specific to nanoscience and which nonetheless explains a recent spate of interest among philosophers of science in the role of scale in scientific concepts and theories. With the three cases of gold atoms that began this discussion I showed that microstructure, as it is commonly conceived in the literature, underdetermines kindhood in chemistry and nanoscience because it fails to capture distinctions between relevant properties that arise from different arrangements of atoms in each of the samples. These properties correlate with a more sophisticated understanding of microstructure, which in chemistry includes molecular geometry and in nanoscience includes dimensionality, size, and surface chemistry alongside composition. I called this ontologically deflationary approach the structure–property paradigm, following the chemical vocabulary.

I showed that in nanoscience, this refined notion of microstructure is an asset rather than a hindrance to the synthetic goals of that science, which suggests that the classificatory practices of nanoscientists are successfully aiding the carrying out of that science's projects. However, the account of nanoscientific classification I sketched fails to align with many tacit best-practices assumptions in philosophy of science for how to set up classification systems. In particular, the tetradic, bottom-up classification scheme fails to be a traditional taxonomy, in which kinds are grouped at successively greater levels of generalization, such as biological taxa are. I argued that this is a problem for philosophical preconceptions about classification, rather than for scientists. As a synthetic science, nanoscience has a rather different set of classificatory needs that typical descriptive sciences. While taxonomy may work for those sciences (and the effectiveness of taxonomy in those sciences is still an open question), it is clear that in nanoscience and other synthetic sciences, greater specificity of structure is well worth the cost of a clean-cut taxonomic framework. This point has received relatively little philosophical inspection up to this point, largely because synthetic scientific projects themselves are not often the subjects of philosophical analysis. Understanding the considerations of synthetic science, though, can aid in an exposition of the latter point. Consider that in synthetic scientific projects, the aim is to make something new, and very rarely is the target creation blindly constructed; instead, a set of heuristics or design principles guide and constrain the scientist's decisions about what sorts of objects, substances, or materials to employ. Materials are selected for their properties, and they are manipulated by altering their structure, further rationalizing the structure–property paradigm.

I also showed that the role of scale as a sortal for classification has been largely overlooked. By this I do not mean that scientists and philosophers have failed to notice that groups of big things are different than groups of small things; rather, scale has been overlooked as a constraint on what sorts of properties, structures, and objects are useful ones for partitioning up a collection of things of scientific and everyday interest. Scale constrains what kinds of kinds are useful to a given scientific setting, and it explains why talk of, e.g., differences in numbers of carbon atoms does not make sense when one is trying to determine whether two macroscopic samples of steel are of the same kind.

To clarify, I return to the three cases of gold atoms that began this essay and consider how each case would be analyzed in a synthetic setting. The macroscopic lump of gold has macroscopic properties, ductility and malleability, as well as non-reactivity, that are similar to other noble metals (e.g. palladium and platinum), as well as to many plastics. The gold atoms, on the other hand, have electronic symmetry, namely d10s1, that is shared by copper and silver. The nearby kinds in the macroscopic case are different that the nearby kinds in the microscopic case; scale has dictated a difference in the kinds of kinds that are relevant in each setting.

As for the gold nanoparticles, it should now be clear that they were themselves underspecified in the original example—their surface chemistry, shape, and size is indeterminate from the introductory description. But knowing that these are the features of the nanoparticles that need to specified to generate relevant alternatives still illustrates the point, which is that classification in the sciences depends on scale and scientific intent, and the properties that aid synthetic scientists in selecting among a range of materials vary with the scale of the system being studied.

II. SURFACE TENSION

At the macroscopic scale, surfaces are often ignored as boundary conditions, that is, parts of a material that do not significantly impact the physical or chemical behavior of that material. At the nanoscale, however, the behavior of surfaces dominates the behavior of materials. This scale-dependent change in the role of surfaces leads to a change in what it means to *be* a surface; this is a form of conceptual instability that Wilson (2008) has termed "multi-valuedness." Responding to this multi-valuedness requires a new account of inter-theory and inter-model relations. This account, which I call the model interaction account, is an alternative to traditional views of reductive and emergent relations between theoretical frameworks, and it hinges on the exchange of information across theoretical frameworks through attention to boundary behaviors. I argue that philosophical theories of explanation, prediction, modeling, and design need to attend more closely to the dependence of these activities on boundary conditions. This dependence is glaringly apparent in the field of nanoscience, where the boundary behaviors of materials dominate the overall behavior of materials.

A. INTRODUCTION

Philosophers of science often characterize explanation in the physical sciences by separating the components of an explanation into two categories: natural laws, and initial and boundary conditions. Laws are descriptive or governing relations among the parameters of the system, and initial and boundary conditions are variable parameters, usually specified by empirical methods or by fiat. This distinction has been crucial in the development of deductivenomological and unificationist theories of explanation, as well as in many conversations about idealization, models, and the role of mathematics in the physical sciences.

Most accounts of explanation focus on the role of natural laws in generating (or failing to generate) explanations and predictions: in addition to the deductive-nomological and unification accounts of explanation, there exist boutique literatures on *ceteris paribus* laws, laws and idealization, and whether laws govern or describe their target systems (cf. e.g. Dretske 1977; Cartwright 1983; Earman et al. 2003; Lange 2009). Initial and boundary conditions are rarely discussed as features of an explanation or prediction that do explanatory or predictive work, and when they are discussed, they are rarely differentiated from one another. But unlike initial conditions, boundary conditions typically contain modeling assumptions crucial to the development of explanations and predictions, as well as to the accomplishment of activities less frequently discussed by philosophers of science, such as the design of materials. In this essay, I argue that philosophical theories of explanation, prediction, modeling, and design need to attend more closely to the dependence of these activities on boundary conditions, and I show that this dependence is glaringly apparent in the field of nanoscience, where the boundary behaviors of materials dominate the overall behavior of materials.

I begin in Section II.B by distinguishing boundary conditions from laws in an explanation of how a violin produces sound. I show that the mathematical model of the instrument's harmonics ignores the features of the instrument actually responsible for producing audible pitch. An additional model of the structure and behavior of the instrument is needed to understand why the violin makes noise at all, and an explanation of the violin's sound is only achieved when the information from both models is synthesized in a non-reductive, non-emergent relationship that I term *model interaction*.

I argue next in Sections II.C and II.D that materials modeling requires model interaction. The boundary conditions in the model of the violin's harmonics are a pair of zero-dimensional points that represent the endpoints of the vibrating string. But in most materials modeling, boundaries are represented as two-dimensional surfaces, i.e., the surfaces of materials that serve as interfaces between a material and its environment. In physical models of materials, the structure and behavior of surfaces are suppressed as boundary conditions, like the endpoints of the string in the model of the violin's harmonics. In chemical models the structure and behavior of surfaces is detailed, but the models ignore the behavior of bulk materials,

analogous to the way in which the structural model of the violin ignores the production of harmonic overtones.

Aligning these two modes of description with one another is a non-trivial and nonreductive project, and I show in Section II.E that it is one of the chief modeling challenges of nanoscience. When physical and chemical models of materials are exported to describe nanoscale materials, the structure and behavior of surfaces change to the extent that what once was a 'mere' boundary now occupies the majority of the material. This change indicates an area of conceptual instability that Wilson (2008, ch. 6) has called "multi-valuedness." In the case of surfaces, the tension between chemical and physical models of surfaces is responsible for one kind of multi-valuedness of the concept, and the move from the bulk scale to nanoscale is responsible for another. This multi-valuedness makes modeling surfaces challenging at both the macroscale and the nanoscale, but at the nanoscale the challenge is more urgent. Using model interaction to develop explanations (and predictions, descriptions, etc.) of boundary behaviors can resolve this challenge.

B. THE SOUND OF AN A STRING: BOUNDARY CONDITIONS, MULTI-VALUED CONCEPTS, & MODEL INTERACTION

When a violinist draws her bow along a properly-tuned open A string, the string vibrates at a fundamental frequency of 440.0 Hz, creating a sound wave with a wavelength of approximately 80 cm. The sumptuous, complex texture of the note—what distinguishes its plaintive tone from the digital drone of a pure-440.0 Hz sound wave—is partially the result of vibrations at this fundamental frequency combining with *harmonic overtones*, vibrations of shorter wavelength that constructively interfere with the fundamental wavelength. The harmonics of a fundamental wavelength are waves whose wavelength is $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{4}$, etc. as long as the fundamental (See Figure II.1). If the violinist lightly touches her finger to the string at the $\frac{1}{2}$, $\frac{1}{3}$, etc. distances on the string, without pinching the string against the neck of the instrument, the harmonic overtone is emphasized and becomes the dominant pitch.¹ In the figure, these points are labeled "harmonic nodes."

This harmonics model of the sound of open A on a violin relies on relating the lengths of sound waves to lengths of line segments. This model requires that the end points of the vibrating string (and the harmonic nodes dividing the string to generate new frequencies) be treated as zero-dimensional boundaries. But in reality the vibration of the A string is not contained by zero-dimensional endpoints; if it were, the sound of the violin would not be audible. This is because the sound listeners hear when the violinist plays open A is not, primarily, the sound of the string vibrating, but rather the amplifications of that vibration in the instrument's sound box. These amplifications lead to further interference, complicating and enriching the sound of the note. The mechanics of sound waves in the violin's sound box cannot be explained by the harmonics model, but rather require an explanation from a *structural model* of the instrument's anatomy. This structural model of open A requires jettisoning the treatment of the string's endpoints as zero-dimensional boundaries, but of course, in so doing it cannot explain the presence of harmonic overtones.

The structural model of a violin (See Figure II.2) facilitates an explanation of how vibrational energy from the A string transfers to other parts of the instrument. As the A string vibrates across the instrument's bridge, it transfers this vibration to the spatially extended region of the string that is in contact with the bridge. One end of the bridge (See Figure II.2) in turn vibrates against the front of the sound box. The other end of the bridge is connected via the sound post to the back of the sound box, and some of the vibrational energy from the bridge is transferred to the back of the sound box via the sound post. The sound box amplifies all these vibrations and, especially at the openings of the sound holes, generates further interference among the various phases and frequencies, as well as between the A's harmonic overtones. All this vibrational complexity results in an A that is comprised of much, much more than just 440.0 Hz waves and a tone that cannot be described simply by its wavelength.

In this explanation of the sound of a violin's open A, the natural laws at work are laws

¹This technique of 'natural overtones' produces a differently-textured pitch than normal playing, and it is denoted in sheet music by a small circle above the fundamental pitch.



Figure II.1: The harmonic series on a vibrating string. Grey circles marking fractional lengths of the string represent harmonic nodes. Harmonic overtones listed are for violin's G string. [Author illustration]



Figure II.2: Anatomy of a violin's soundbox, used to model transmission of sound waves from a violin string, through bridge, body, etc. [Author illustration, adapted from Creative Commons licensed sources]

about interference of waves. The harmonics model presents these laws more centrally than the structural model, building a network of relationships among wavelengths in order to explain why constructive interference occurs where it occurs. But even in the harmonics model, the boundary conditions—that is, the idealization of harmonic nodes and the string ends as dimensionless, fixed, and infinitesimal—play the crucial role, not natural laws. While the laws of wave interference are necessary for the explanation of open A, that necessity does not automatically confer upon the laws the status of 'hardest-working' explanantia: the boundaries, or endpoints and harmonic nodes, of the string are what generate new frequencies in the harmonics model; they do most of the work in generating the harmonic overtones, or modes, that are used in explanations and predictions of the string's sound. It is no coincidence that many models of violin harmonics do not mention the wave equation, nor even simplified phenomenological 'laws' relating wavelength, frequency, and the speed of sound. Moreover, without endpoints on a line segment, mathematical relationships between harmonic overtones could not be defined. In order for the harmonics model to convey useful information about the sound of an A string, it is necessary to suppress as a boundary condition the structural information about the transmission of sound through the instrument.

Wilson (2008, pp.184–192) (2012) has called suppressions of this sort "physics avoidance." He argues that areas of overlap between theoretical frameworks are often found where one model "avoids" complicated physical behavior at an interface by treating a part of the system as an infinitesimal parameter and another model represents that part of the system as a richly-textured set of structures and behaviors. In such cases, it becomes evident that scientific concepts such as *weight* and *hardness* are represented by multiple, often conflicting theoretical frameworks, and Wilson shows that this multi-valuedness is not only unproblematic for science, but it is often necessary, and it is productive in its ability to inspire further investigations into areas of overlap between theoretical frameworks. Wilson's concern lies primarily in justifying the multi-valuedness of scientific concepts by exploring areas of overlap between theoretical frameworks on a case-by-case basis and showing that in each case, areas of conceptual and theoretical breakdown parallel changes in the mathematical frameworks used to describe the system. His message is that physics avoidance is a necessary and successful strategy, and that it should be acknowledged as an alternative to traditional accounts of inter-theoretic relations that rely on notions of reduction and emergence.

This message can be seen clearly in the explanation of open A. The harmonics and structural models are not in a reductive nor an emergent relationship with one another, and where they overlap is precisely at the point (literally) where the harmonics model suppresses information about the system in the form of a boundary condition. In order to develop an explanation of the movement of sound waves through the violin string and into the air around the instrument, information from both models is needed. As an alternative to reductive and emergent relationships between models, then, I propose to call inter-model relations of this sort cases of *model interaction*. Wilson's physics avoidance is one type of model interaction; Batterman's asymptotic explanations (2001), Winsberg's handshaking models (2006; 2010), Mitchell's integrative pluralism (2009), and Wimsatt's rainforest landscapes (2007) are others.

Model interaction keys in on the idea that explanatory (or descriptive, predictive, or synthetic) work is often accomplished neither by reducing one model to another, nor by constructing emergent relationships between phenomena, but by attending to areas of overlap between theoretical or modeling frameworks. These overlaps occur where one model or theory suppresses as a boundary condition precisely the part of a system that another model or theory aims to describe. Not coincidentally, boundary conditions frequently bear the brunt of explanatory burdens in a way many philosophical accounts of explanation have so far failed to appreciate.

The violin example highlights two related tools for the evaluation of modeling activity in the sciences: conceptual multi-valuedness and model interaction. In the remainder of this chapter, I show how these tools can be used to identify and resolve a central challenge in the modeling of nanoscale materials, namely the problem of constructing effective models of nanoscale surface structure and behavior.

Physical models of macroscopic material surfaces usually treat surfaces as two-dimensional boundaries, analogous to the 0-dimensional endpoints in the harmonics model of vibrating strings, whereas chemical models of macroscopic material surfaces treat surfaces as richlystructured three-dimensional regions, analogous to the moving sound waves in the structural model of the violin. Chemical models cannot suppress surface structure into a boundary condition, because chemical reactions occur at the surfaces of solids. Interior atoms are less likely than surface atoms to participate in reactions for two reasons. First, interior atoms are bonded to, or coordinated with, more other atoms than surface atoms (See Figure II.3); another way to phrase this point is by saying that surface atoms, but not interior atoms, have *available reaction sites*. Second, interior atoms are physically screened off or blocked from interaction with new chemical species by the presence of nearby atoms, much like it is impossible to reach a friend in the middle of crowded room without pushing past the people on the edges of the room. Whole fields of chemistry are devoted to the principle that increasing surface area increases chemical reactivity.² On the other hand, many physical behaviors, such as bending, transmitting electricity, and changing temperature, are not surface phenomena. Section II.C discusses the multi-valuedness of the concept *surface* in more depth, and Section II.D argues that surface models participate in model interaction.

In Section II.E, I move from macroscopic materials modeling to nanoscale materials modeling. Nanomaterials modeling presents a particularly apt case of model interaction because the role of surfaces in constraining and explaining the behavior of nanomaterials is different than the role of surfaces in constraining and explaining the behavior of macroscopic materials. This difference is due to the fact that as materials get smaller, the proportion of surface atoms to interior atoms changes. While surfaces make up a negligible proportion of macroscopic materials, surface atoms make up anywhere from 5% to 100% of nanomaterials (See Figure II.4). Consequently, at the nanoscale physical models of surfaces, which treat surfaces as two-dimensional boundary conditions, lose explanatory and predictive power. While physical models of surfaces interact with chemical models of surfaces—which do not treat surfaces as two-dimensional boundary conditions—even at the macroscale, the need for and benefits of that model interaction become more apparent at the nanoscale.

Additionally, one of the central activities of nanoscience is the design and synthesis of materials. These activities make use of models and theories, but unlike accounts of theories

²This is one of the founding principles of heterogeneous catalysis. Catalysis is the use of an additional chemical agent, the catalyst, to increase reactivity or rate of reaction. Heterogeneous catalysis uses a catalyst that is in a different phase of matter than the reactants, either a solid in the presence of liquids and/or gases, or a liquid in the presence of gases. The heterogeneous catalyst provides extra reaction sites at which the reactants can interact, and the number of sites available is directly proportional to the amount of catalytic surface with which the reactants are in contact. Heterogeneous catalysis is one of the largest fields of industrial chemistry.



Figure II.3: A ball-and-stick cartoon of a crystalline solid. Inset A shows an interior atom with no available reaction sites. Inset B shows a surface atom with an available reaction site at the top of the atom. Note also the difference in bond lengths between surface atoms and interior atoms. Bonds on the surface of the material (the uppermost layer) are shorter by n than bonds between interior atoms. Greyed-out circles and lines represent atoms and bonds farther back in the crystal. The foremost, blacked-in layer represents merely a cross-section of the material, rather than a material surface. [Author illustration]

and models in explanation, prediction, and description, accounts of theories and models in synthesis are rare in the philosophical literature. When they do appear, they are often only sparsely sketched. A fuller account of theory and model use in activities like synthesis and design appears in Chapter IV.

Finally, nanoscience is still a young science. The question of how theories and models interact in the modeling of nanomaterials is a live one not only for the philosophical community, but also for the scientific community. This paper is aimed at the former audience, but applications of its findings have already proven useful to communities of nanoscientists.

C. THE MULTI-VALUEDNESS OF SURFACE

Surfaces are divisions between bodies or materials. No solid or fluid object exists without being bounded by a surface. Nonetheless, many physical models of materials *iqnore* the true physical behavior of surfaces by treating surfaces as boundary conditions or otherwise parameterizing away surface behavior. Commonplace mechanical models of simple machines, such as blocks sliding down planes or wheels rolling on flat surfaces, ignore both the physical and the chemical behavior of surfaces. As Wilson has argued in his (2010), fluctuations in the microscopic trajectories of the point masses that make up blocks and planes cannot be individually accounted for in a modeling framework that makes macroscale predictions about the trajectory of the block as a whole. Strategies for treating surfaces as boundary conditions in macroscale materials modeling abound: consider continuum models of deformation in solids, computational fluid dynamics, as well as the high-school physics models of frictionless planes and pendula swinging from massless strings. Despite treating surfaces as boundary conditions, these models are successful predictors and explainers of macroscopic material behavior in many situations (Phillips 2001), and some philosophers (cf. e.g. Batterman 2001; Wimsatt 2007; Wilson 2010, 2012) have argued that this success occurs because of, not in spite of, the models' suppression of surface behaviors.

While these authors are correct to conclude that the suppression of surface considerations is often an effective modeling strategy, there are contexts they have overlooked in which it becomes useful to study the structure and dynamics of the surfaces themselves—in the design of adhesives and lubricants, for instance, or the development of touch screens or sterile medical instruments. These contexts are often more "applied" or materials-based than the physical modeling contexts mentioned in the previous paragraph, and they often use conceptual resources from both physics and chemistry to model the structure and behavior of surfaces.

These surface-centered contexts appear most often in contemporary materials science, chemistry, and engineering, where small differences in surface structure or behavior can lead to larger-scale changes in a metal's electrical or thermal conductivity, or differences in mechanical, electronic, or chemical properties.³ Rather than treating surfaces as boundary conditions or single-variable parameters, in these materials models surfaces are decomposed into structural elements. Those structural elements explain and predict the behaviors of materials, from frictional forces to conductivity and chemical reactivity.

The differences between treatments of surfaces in chemical, physical, and applied-materials contexts exhibit the multi-valuedness of *surface*. Fans of conceptual parsimony might still point out that each of these models can define a surface as a boundary region in a material that takes up an negligible proportion of the material. But even this thin unifying characterization does not hold up at the nanoscale. There, the percentage of atoms on the surface of a material becomes statistically significant: rather than making up a negligible fraction of the material's mass and structure, surfaces make up significant proportions of the material, anywhere from 5–100% (See Fig. II.4). This means that physical surface behaviors, which are modeled by interfacial dynamics or ignored by physical models of materials, become the dominant physical behaviors of the material. So the role of the surface in influencing the physical behavior of the material changes.

Increased surface area also increases chemical reactivity. Additionally, novel physical and chemical phenomena such as LSPR (See Section II.E) and quantum confinement arise at

 $^{^{3}}$ It is worth mentioning that contemporary materials science has developed techniques for modeling *interfacial dynamics*, which focus on the structure and behavior of spatially extended regions that contain parts of both a material and its environment, or that contain multiple materials. The interfacial-dynamical strategy for modeling surfaces conceives of surfaces not as sites of reactions, nor as ignorable boundary conditions, but as extended and complexly structured regions of spacetime. The interfacial-dynamics approach to modeling surface behavior is thus an additional alternative approach to framing the concept *surface*.



Figure II.4: Relative amounts of material that make up the material's surface for spherical palladium clusters. Far-right abscissa marks at 10^5 nm are around $\frac{1}{4}$ the width of a human hair. [Reprinted with permission from Nützenadel et al. 2000]

the surfaces of nanomaterials. At the nanoscale, surfaces simply cannot be written off as boundaries. As the structural role of surfaces changes, the concept *surface* itself changes. It is a change in the scale of the system being studied that induces this conceptual change. Recently, Phillips (2001), Winsberg (2010), and Batterman (2012) have written about the role of scale in models and explanations of physical systems. These authors all invoke the term "multiscale modeling" and argue that attention to scale leads to explanatory and predictive insights. But although Winsberg uses nanoscience as an example of multiscale modeling, none of these authors acknowledge the unique position of nanoscience as an *entire* discipline framed around the study of a length scale, nor the role of this length scale as a motivating force for combining modeling strategies from multiple disciplinary traditions. The nanoscale is important precisely because it is where the curve in Figure II.4 peaks. When the size of the material approximates the size of its component atoms, what was once an effectively infinitesimal boundary is now a robust area of the material, and new modeling strategies are needed. The particular modeling challenge in nanoscience, then, is that of extending a familiar but multi-valued concept to a new domain of application. After some preliminaries on the mechanics of macroscopic materials modeling in Section II.D, I show in Section II.E how the model interaction account helps modelers respond to this challenge.

D. MODELING MACROSCOPIC SURFACES

So far I have indicated a distinction between physical and chemical models of materials that is analogous to the distinction between the harmonics model and the structural model of the sound of a violin's A string. But the distinction between what is 'properly' physical or chemical is hazy at best and arbitrary at worst. And the two kinds of models differ in their methodology and assumptions significantly enough that it is possible to distinguish them without referring to the disciplines out of which they arose: harmonics-type models are *continuum* models of materials, and structure-type models are *atomic*, *discrete*, or more generally *structural* models of materials. Continuum models treat surfaces as two-dimensional boundary conditions. Atomic models can either represent bulk structure or surface structure, but rarely do they represent both.

1. Continuum Models

Continuum models of materials fall into a variety of categories, based on factors such as the kind of response they are intended to model, the kind of material to which they apply, and time-, length-, and energy-scales they employ in order to draw relationships between a material's structure and its behaviors or properties. For instance, in linear elastic modeling, Young's modulus (also known as the elastic modulus) is a measure of how much a material will elastically deform when subjected to a certain amount of stress that is evenly distributed over a given area on the material's surface. It addresses only elastic, as opposed to plastic, responses in the material. Elastic deformations, like stretching a rubber band, return the material to its original, pre-deformed shape once the deforming force is removed, whereas plastic deformations, like crushing a soda can, leave the material in a new shape even after the deforming force is removed. Young's modulus is time-independent, and it applies over any macroscopic length scale. It applies to any type of material that can generate an elastic response to stress, regardless of whether the material is metal or wood, plastic or crystalline.

This wide scope of applicability results from the fact that Young's modulus is a continuum parameter. It treats materials as if they were made up of uniformly distributed and continuously deformable parts. Young's modulus is one of many continuum parameters found within the catalog of physical models of materials. Continuum models do not consider the atomic or microscopic structure of materials in their evaluations of the relationships between structure and behavior. Rather, they focus on macroscopic structural features of materials such as shape and size. For instance, one expression of Young's modulus is:

$$E = \frac{FL_0}{A_0 \Delta L}$$

where F is the magnitude of force applied to the material, L_0 is the length of the object before a force is applied, A_0 is the cross-sectional area of the object before a force is applied, and ΔL is the change in the object's length after application of force. E is Young's modulus itself. The value of E is considered to be an *intrinsic property* of a material, one that is consistent across all samples of the material and which depends minimally on the environment in which the material is embedded. Young's modulus in particular is an invaluable resource for both engineers and materials scientists; it comprises a whole section of the *CRC Handbook of Chemistry and Physics*, a definitive resource for academic and practical information about properties of materials.

Continuum models, like those that employ Young's modulus as a parameter, ignore the microstructure of the materials they model. Nonetheless, these models are often among the most useful for solving both scientific and engineering problems, such as understanding how sound might travel through a material or constructing buildings that will resist collapse. In a comprehensive review of material properties, Ashby (1989) provides an overview of a variety of physical models of materials and relates them to various projects. Young's modulus features prominently in this review, which demonstrates that one of the most effective means of classifying materials for these sorts of purposes is by relating properties that ignore the material's microstructure, like Young's modulus, to other intrinsic properties of the material, like density.⁴

A central feature of the review is the development of charts, such as the one in Figure II.5, that graphically depict the kind of classification schemes Ashby has in mind. They are effective or phenomena-capturing, rather than 'fundamental'; the models do not reduce descriptions of behaviors to descriptions of atomic motion or quantum states. The success of these models in aiding engineers and scientists in the construction of complicated objects and in the understanding of material behavior flies in the face of reductive attitudes toward theories (cf. e.g. Nagel 1979), which argue that understanding a scientific system is best achieved by understanding the smallest components of the system.

Phillips (2001) and Wilson (2012) have pointed out a number of the problems with using such bottom-up models, for instance employing point-mass mechanics to try to construct materials or understand material behavior. Wilson, along with Batterman (2012), has been

⁴Density is a particularly interesting material property, because it can be calculated without referring to the material's microstructure by simply quotienting mass and volume. However, it is often rationalized in microstructural terms, by referring to the atomic mass of the material's components and the spacing between atoms in the material.



Figure II.5: An Ashby diagram relating density to Young's modulus. [Image courtesy Granta Design (M. Ashby and D. Cebon), licensed under Creative Commons]

interested in the success of continuum models in spite of the fact that they ignore microstructure. But what is interesting for present purposes is the success of these models in spite of the fact that they ignore the role played by the material's surface. Young's modulus is an intrinsic property of a material, so it does not change depending on the interaction between the surface of the material and the environment in which the material is embedded. Young's modulus predicts the same material response if the beam is submerged in the ocean as if it is suspended in a vacuum.

These environmental considerations can be ignored in macroscopic materials because the surfaces make up a negligible proportion of the entire material.⁵ As a result, continuum models, which treat surfaces as boundary conditions, provide accurate and reliable predictions of the behavior of such materials. But as previously noted, at the nanoscale, the mass (and volume) of the material's surface reaches the same order of magnitude as the mass

⁵For instance, the percent of atoms on the surface of a typical 6mm-diameter steel ball bearing, which is a small macroscopic material, is only about 1.88×10^{-5} .

(and volume) of the rest of the material. As a result, models that treat surfaces as boundary conditions risk neglecting a significant proportion of the behavior of materials at the nanoscale.⁶

In many cases of modeling, continuum models are used successfully without supplementation from atomic models: one can explain why a steel beam bends without having to mention that steel is made up of atoms. But continuum models do not fare so well in explaining the behavior of systems made up of multiple kinds of materials. This defect of continuum models can have catastrophic consequences: When engineers designed the first commercial jetliner, the de Havilland Comet, in the early 1950s, they built the aircraft with square windows. The sharp corners of the glass windows caused massive metal fatigue and led to a number of the jets ripping open in the sky, killing dozens of passengers. In order to understand why the corners of the windows led to metal fatigue in the body of the jetliner, it is necessary to develop an interactive-model explanation using both continuum models and structural models.

2. Structural Models

Structural models of materials represent the material's microstructure, that is, the constituent atoms or groups of atoms that make up a material, just as the structural model of the violin represented the parts of the instrument that transmit and amplify sound waves. They use geometric relations between parts of the material to predict and explain material structure and behavior. Nearly all molecular and crystalline models of materials refer to such geometric arrangements, as outlined in Chapter I's discussion of the structure–property paradigm. Structural models can target either molecular or crystalline bonding behaviors of materials. Structural models needn't explicitly represent atoms; models of grain boundaries in crystalline materials are still structural models, because they divide a material into non-identical parts.

⁶The limits of predictive accuracy of individual continuum parameters such as Young's modulus are not yet well-understood or well-defined. For instance, Young's modulus may be applicable to metal nanomaterials as short as a few dozen nanometers in length. In my work with the Millstone group, we have devised experiments that test the limits of predictive accuracy of Young's modulus in silver. The experiment is still underway, and no results are presently available.

Metals like the steel in the Comet aircraft are modeled using both crystalline and grainboundary models. Because metals are polycrystalline, the atoms in a metal are grouped into small patches called grains. Grains are irregularly sized, shaped, and distributed in most metals (see Figure II.7). Within each grain, atoms are arranged in monocrystalline arrays that is, in regular, repeating, structures (such as the one in Figure II.3) that can be modeled mathematically by translation operations on a hypothetically infinite lattice. Mathematical tools such as group theory, Bravais lattices⁷, and Miller indices are all used to model the structure of individual crystals. In an elemental metal, different grains all have the same lattice structure but different orientations; in alloys, lattice structure as well as orientation can differ from one grain to the next, and regions of one metal are often composed of multiple grains.

Atoms in metals are highly mobile, which is why metals are malleable and ductile. But it requires less energy for atoms in a grain to move as a group than to deform the grain by changing grain boundaries. So when a metal deforms in response to stress, grains rearrange. The dynamics of grain movement in a metal are similar to the movement of grains of uncooked rice in a bowl: when a small surface area is stressed, fewer grains move, but the ones that move, move more. When a larger surface area is stressed with the same force, more grains move, but each grain moves less. This is why you can easily plunge your finger into the bowl, but not your open palm.

For the square-windowed Comet, the sharp corners of the windows stressed the metal that surrounded them. That stress was applied over a small surface area, leading to a buildup of high-energy grain movement away from a small region of the material, which led to cracks in the metal panels. Historically, continuum models concerned with this sort of high-metalstress situation modeled the plastic deformation of metals via single-scale and, obviously, continuum strategies. This way of modeling catastrophic metal fatigue necessarily involved idealizing the areas of fatigue as boundaries—singularities and/or infinities—a strategy that met mixed success and, in the case of the Comet, some clear failure. Models that are

⁷A Bravais lattice is an infinite array of discrete points in n-dimensional space, generated by a set of translation operations on a single point. Three-dimensional Bravais lattices are used to represent the microstructure of most crystalline solids; these are summarized in Table II.6. Bravais lattices can be used to predict geometric, optical, and electronic properties of a variety of crystals. Information about bulk material properties like hardness, solubility, and optical clarity can be inferred from crystalline geometries.

Lattice Sys- tem	$a_1, a_2, a_3 =?$	$\varphi, \psi, \theta = ?$	Centerings	Primitive Unit Cell
Triclinic	$\mathbf{a_1} \neq \mathbf{a_2} \neq \mathbf{a_3}$	$\varphi,\psi,\theta\neq90^\circ$	Primitive	
Monoclinic	$\mathbf{a_1} = \mathbf{a_2}, \neq \mathbf{a_3}$	$\begin{array}{rcl} \varphi,\psi &=& 90^{\circ},\\ \theta\neq 90^{\circ} \end{array}$	Primitive, Base- Centered	
Orthorhombic	$\mathbf{a_1} eq \mathbf{a_2} eq \mathbf{a_3}$	$\varphi,\psi,\theta=90^\circ$	Primitive, Base- Centered, Body- Centered, Face- Centered	
Tetragonal	$\mathbf{a_1} = \mathbf{a_2} \neq \mathbf{a_3}$	$\varphi, \psi, \theta = 90^\circ$	Primitive, Body- Centered	
Rhombohedral (Trigonal)	$a_1=a_2=a_3$	$\varphi = \psi = \theta \neq$ 90°	Primitive	J.J.
Hexagonal	$\mathbf{a_1} = \mathbf{a_2} = \mathbf{a_3}$	$arphi = 120^{\circ}, \psi, \theta = 90^{\circ}$	Primitive	
Cubic	$\mathbf{a_1} = \mathbf{a_2} = \mathbf{a_3}$	$\varphi,\psi,\theta=90^\circ$	Primitive, Body- Centered, Face- Centered	

Figure II.6: 3D Bravais Lattices [Author illustrations]

able to predict the cracks as resulting from changes in the surface and internal structure of the material, unlike single-scale continuum models, have fared better and have received prolonged interest especially in humanity's airborne decades. Single-scale continuum models can describe expansion and contraction of the metal in response to changes in surrounding temperature and pressure, and they can describe the deformation of metals in response to external stresses, but they cannot predict the buildup of localized stresses along grain boundaries that can lead to the appearance of cracks. Here is where multiscale models and mixed continuum/structural models fare better.

Mathematical models of crystal structure cannot predict cracks either, and for a similar reason: crystal models like Bravais lattices assume infinite arrays of lattice points, and so they do not contain the conceptual tools necessary to represent surfaces. The response of materials modelers has been to adapt lattice models of materials to develop structural models of surfaces. Surface-layer atoms are not in the same relationships to other atoms in the material as interior atoms, so they cannot be mathematically modeled as points on an infinite, repeating lattice. But they can be visually modeled using the graphical notion of a lattice (See Figure II.3). Surface atoms bond to fewer other atoms in the material. This leads to changes not only in the lattice representation of the surface layer, but also in the energetic and chemical behavior, as well as the geometric structure, of the surface. If it were the case that surface atoms were simply bonded to fewer atoms, but no other changes in physical or chemical properties occurred at the surface, then the problem of how to model surfaces would be largely uninteresting. A mathematical caveat would have to be written in to deal with the failure of the infinite-lattice representation of material structure at the material's surface, but no other part of the theory would need to change.

However, surface atoms often differ in bond length and atom placement from interior atoms, as a result of the relative instability of surfaces. So the relative locations of points in a lattice-like model of the material changes at the surface. Surface structure is, consequently, defined by contrasting the actual locations of surface atoms with the expected locations of atoms in an infinite lattice. Predictions about surface geometry and surface energy are obtained by comparing expected values from a lattice model to actual values and inferring what possible reconfigurations could perturb the expected values to bring them closer to the



Figure II.7: Left: Scanning electron microscope image of grain boundaries in a metal. Right: cartoon of grain boundaries indicating differences in lattice orientation between grains. [L: Image by Edward Pleshakov, reprinted under GNU Free Documentation License R: Author illustration]

actual values. The resulting structural model can be further tweaked to represent defects in a surface and areas of increased reactivity. This is a case of extending conceptual resources from one model into a new domain of application, and it illustrates the non-reductive method of model interaction that is seen in materials modeling of surfaces.

Bravais lattices cannot represent surfaces, and yet the cartoon of differently-oriented pieces of a lattice in Figure II.7 is a sensible and informative model of the structure of a polycrystalline metal. This cartoon is a *hybrid model*, one kind of product of model interaction (I discuss others in Chapter III). The construction and application of hybrid models is one way of using model interaction to respond to the challenge of extending conceptual resources from one model into the domain of applicability of another model. Hybrid models are neither reductions of higher-level models nor reducible to lower-level models. Against reductive intuitions, separating the pieces of the lattice-grain-boundary model into higher-level (grain-boundary) and lower-level (lattice) components makes it less useful in explanations of the mechanics of crack formation, not more.

The lattice-grain-boundary hybrid model exemplifies the method of model interaction, in which representational, predictive, or explanatory elements from multiple models are combined in response to modeling challenges posed by novel domains of application. Because this modeling challenge arises when something that was a boundary condition in one context becomes a domain of interest in another context, it makes sense both that model interaction is the appropriate strategy for modeling boundary conditions and that boundary conditions, rather than natural laws, perform the majority of the explanatory, predictive, and descriptive work in these cases.

E. NANOSCALE SURFACES

Surfaces are the part of nanoscale materials that differ most noticeably and most systematically from macroscopic materials. Changes in the surface structure of a nanomaterial or in the proportion of the material on the surface can lead to changes in known material properties or behaviors, or to the appearance of genuinely novel, scale-dependent behaviors. Changes in known properties include the semiconductivity of the nanomaterial graphene.⁸ Novel, scale-dependent behaviors include the optical phenomenon *localized surface plasmon resonance*. LSPR is the collective, resonant oscillation of electrons at a material surface. Resonant oscillation means that the surface electrons are moving collectively at a frequency similar to that of the photons or electrons that stimulate the oscillation. In order to exhibit resonant oscillation, the electron field (the plasmon) must evolve a standing wave. This only occurs when the surface is significantly smaller than the wavelength of the incident photon or electron. (See Figure II.9.) This optical phenomenon is responsible for the appearance of colors in stained glass.

Nanoscientists are interested in modeling LSPR in order not only to understand it, but

⁸Graphene, which is a 2D carbon nanomaterial made up of (ideally) single-layer or (realistically) very thin multi-layer sheets of hexagonally arranged carbon atoms, is analogous to the macroscopic material graphite (pencil lead). Graphite exhibits some range of electromagnetic behaviors, as do all materials; it so happens that graphite's response to electromagnetic stimulation is not that interesting, as it is neither a conductor nor diamagnetic.

At the nanoscale the electrical conductivity of hexagonally-arranged sheets of carbon atoms improves, to the extent that it is possible to pass a current along a graphene surface. Graphene is considered a semiconductor, meaning that in graphene, the energies of conduction-band and valence band electrons are very similar to one another. This electronic property distinguishes graphene from graphite, which is significantly less conductive due to intermolecular interactions between the sheets of carbon that make up bulk graphite.



Figure II.8: When the resonant frequency of an LSPR effect is in the visible spectrum, brilliant colors appear. LSPR is the primary phenomenon responsible for the colors in historic stained glass, like these panels from Gaudi's La Sagrada Familia Cathedral in Barcelona, ES. The colors in stained glass are achieved by the addition of finely divided metals or metal oxides to molten glass; the deep reds in this image, for instance, are most likely due to LSPR in colloidal gold particles. [Image by Tony Hisgett, reprinted under Creative Commons Attribution 2.0]



Figure II.9: Localized Surface Plasmon Resonance: The (grey) electronic field around a group of atoms is perturbed by external electronic or photonic stimulation (not pictured), and the perturbation is restored by nuclear forces. This collective perturbation and restoration generates a collective movement away from the stimulated nuclei of the electron probability distribution around the nuclei (black), which is a plasmon. When the stimulation frequency matches the frequency of the plasmon, the plasmon will resonate at that frequency. In macroscopic materials, plasmons will travel (propagate) down the length of the material; in nanoscale materials, when the resonant wavelength is longer than the material's length, the plasmon is "trapped", or localized, in a standing wave of a given frequency. [Author illustration]

also to control the intensity of its output and the input frequency at which it responds. The Mie solution to Maxwell's equations (See Figure III.4) models output intensity in terms of input frequency and structural features of the material. The Mie solution is itself a sort of hybrid model: it is neither entirely continuum-based (it refers to electron distribution) nor discretely structural (the radius of nanomaterial is allowed to vary continuously); it makes certain idealizing assumptions about the nanomaterials it models (e.g. treating them as perfectly spherical); and it does not model surfaces explicitly. Rather, it makes an oblique reference to an interfacial-dynamical conception of *surface* by including terms that represent interaction between the dielectric (electric susceptibility) of the material and the dielectric of its environment.

In order to construct nanomaterials with specific LSPR responses, it is insufficient to play a game of guess and check with these dielectric terms. In fact, once scientists are engaged in a synthesis design, they rarely refer to the Mie solution at all. The Mie solution model rationalizes expected trends between nanomaterial structures and behaviors, but the real work of nanoscale synthesis is in designing conditions that will generate a particular structure and maintain it over time in a stable environment.

To accomplish this task, scientists need additional conceptual resources to model surface structure. Unlike the surfaces of macroscopic materials, the surfaces of nanomaterials need to be stabilized. This stabilization requires scientists to model the energetics of surfaces using a combination of techniques including Bravais lattice models, structural models of surfaces, and dynamical models of colloid nucleation and growth at the surfaces of nanoparticles.

None of these models originated in nanoscience, nor did the Mie solution. And each model fails to predict or explain under various realistic conditions, because of the different ways in which surfaces are treated in the different models. Dynamical models of colloid nucleation, for instance, treat nanoparticles as continuous bodies and so lose predictive accuracy when the particles are under a certain diameter. Structural models based on Bravais lattices fail to describe the surface energy of spherical particles, because the lattices can only predict crystals in the shapes of lattices themselves.

In order to model nanoscale surfaces, these models are combined and adapted from their original domains of application at larger (Bravais lattice and colloid nucleation models) and smaller (molecular models) scales. In each of the models' original domains, developing a model of surface structure and behavior was not the primary modeling consideration. But rather than starting from scratch and building new models to represent and manipulate surface structure and behavior in nanomaterials, scientists use parts and pieces from models that they already know to be useful. In this way, model interaction circumvents traditional worries about incommensurability in the development of novel theoretical and modeling frameworks. The effectiveness of this strategy provides evidence that despite its multivaluedness, the concept *surface* is still gainfully employed at a variety of widely separated length scales.
F. CONCLUSIONS

I have shown that *surface* is multi-valued at both the macroscale and the nanoscale. This multi-valuedness is interesting in its own right as a demonstration of the patchy, piecemeal conceptual foundations of a scientific concept that plays important roles in all the physical sciences. Further, attending to the modeling strategies developed to accommodate the multi-valuedness of *surface* has exemplified the modeling strategy I have termed "model interaction" and presented it as an alternative to traditional views of reductive and emergent relationships between theories and models.

As a way of responding to new modeling challenges, model interaction calls to mind Neurath's famous (1932) adaptation of the Ship of Theseus, wherein sailors repair and rebuild a ship out of its own parts in the midst of an ocean voyage. While Neurath's concern was primarily in confirmation rather than modeling, a parallel moral stands: Rather than expecting to build new models out of whole cloth to fit novel linguistic, conceptual, and modeling situations—such as scientific research into a new scale of matter—modelers adaptively recycle models, retrofitting and hybridizing them to accommodate previously unexpected consequences of changes in domain.

There is no algorithm for extending models beyond their original domains of application, but I have argued that turning one's attention away from natural laws and toward boundary conditions often indicates areas ripe for model interaction. In the case of surfaces, the tension between treating surfaces as boundary conditions and treating them as structural elements of the material gives rise to both the multi-valuedness of the concept *surfaces* and the modelinteraction strategy for responding to this multi-valuedness. This discussion of surfaces is meant literally, but it also stands as a useful allegory for the challenges the scientists and philosophers can expect to face as nanoscience continues to mature.

III. MAKING THINGS SMALL

This chapter considers further scale-dependent conceptual changes in nanoscience. The discussion moves from an account of the structure and behavior of nanomaterials to the synthesis processes used to fabricate these materials and pays particular attention to the close relationship between these two types of accounts. I argue that contemporary philosophers of science have thus far failed to appreciate this aspect of scientific reasoning. I develop an account of iterative interpolation in the following chapter under a more general consideration of the differences between synthetic and descriptive scientific practices.

A. INTRODUCTION

In Chapter II I argued that the scale-driven relationship between the proportion of the material found on the material surface and the appearance of novel surface behaviors like LSPR are reasons to think that *surface* is a scale-dependent concept. I showed that the scale-dependence of the behavior of surfaces, and the resulting multi-valuedness of the concept *surface*, also constrain the kinds of reasoning strategies that can be used successfully to carry out the projects of nanoscience. Chapter I showed that the reasoning strategies used in common theoretical activities like classification can depend on the nature of the projects one is interested in accomplishing, and one of the central projects of nanoscience is synthesizing nanoscale materials. To that end, this chapter and the next explore the reasoning strategies employed in synthesis in more depth, arguing in particular that philosophy of science has failed so far to incorporate into its purview useful accounts of these strategies.

While philosophical considerations on the nature of explanation, theory structure and

interpretation, inter-theoretic relations, and confirmation all touch on interrelated themes about the reasoning strategies that scientists employ, they largely ignore the reasoning strategies that support the effective design and implementation of synthesis experiments. This chapter presents an overview of some of the special synthetic constraints imposed on nanoscale synthesis by the very scale of the materials being synthesized. Starting with this narrow focus provides a vivid example of the reasoning strategies that scientists need to solve synthetic problems in light of material constraints, and it continues the previous chapter's theme of treating the concepts around which models and theories are built as scale-dependent. I construct my account of reasoning strategies in the synthetic sciences by beginning with nanosynthesis in part because of a certain scale-dependence of philosophical argument: without a clear picture of the challenges and solutions particular to a given synthetic endeavor, general remarks on the nature of reasoning in the synthetic sciences may seem rather thin and unconvincing. This chapter's continued study of nanosynthesis is intended to thicken and substantiate my broader remarks on reasoning strategies in the synthetic sciences, which are found in Chapter IV.

The chapter proceeds as follows: Section III.B familiarizes the reader with the central challenge of nanosynthesis, that of limiting the growth of nanomaterials above the nanoscale. This challenge is unique to nanosynthesis, and it illustrates not only how scientists use theories and models in synthetic sciences, but also why attention to the scale of the system being modeled is important for an understanding of the models' uses. I survey models of growth, used to address this challenge, in section III.C, and I discuss how these models are combined in nanoscientific settings in section III.D. Here I explicate my vocabulary of "model interaction," and I show how it proposes a non-reductive, non-emergent set of relations among models of a problem system. Section III.E applies model interaction to specific synthetic problems in nanoscience, and section III.F extracts a related use of models and theories, which I call "tuning," from further examples of synthetic challenges in nanoscience. I close in section III.G by comparing my account of model interaction and tuning with another recent account from Eric Winsberg of non-reductive, non-emergent relations among models and theories.

B. LIMITING GROWTH: THE CHALLENGE OF NANOSYNTHESIS

Surfaces impact overall material behavior differently at the nanoscale than at the macroscale. In the previous chapter, I showed how this difference affects conceptualization and theory development. But it also has a very practical consequence: unlike in bulk-scale synthesis, in order to synthesize nanoscale materials one must, in general, come up with a plan to stabilize the material's surface. This has led to the development of specially-adapted, scale-specific synthetic practices, the reasoning strategies behind which are often piecemeal and patchy but, perhaps surprisingly, often quite effective. Theories and models have been poached from both macroscopic (larger) and molecular (smaller) scales in order to support those practices. These theories include accounts of the kinetics of individual molecules, the acidity or basicity of a system, the probability of migration of atoms on the surface of a solid, and more. The particular ways in which these models are sewn together to solve individual synthesis problems varies depending on the kind (composition, dimension, shape, surface chemistry) of nanomaterial being made and the synthetic approach of the scientist.

The biggest challenge of synthesizing nanoscale materials is keeping the dimensions of the materials at the nanoscale. Recall from the previous chapter that a significant proportion of the atoms in any given nanomaterial can be found on the material's surface. Recall also that atoms bonded at a surface are higher-energy and less stable than atoms in the interior of a material. So nanomaterials are inherently higher-energy, and thus inherently less stable, systems than macroscopic materials.

Once a material's dimensions exceed the nanoscale, the proportion of atoms on its surface becomes vanishingly small and the material is, consequently, more stable. But once a material grows above the nanoscale, it ceases to exhibit most of the properties and behaviors that make it interesting as a nanomaterial. So one of the biggest challenges of nanosynthesis is *how to limit the growth of materials*, or equivalently, how to stabilize nanomaterials. While the vocabulary of stability and minimization of system energy may be more familiar to readers with a physics background, I will primarily use the chemists' vocabulary of "limiting growth," because this vocabulary captures both chemists' native language and the diachronic aspect of synthetic processes in a way that "stabilization" cannot.¹

Because nanoscale systems are inherently unstable, the synthesis of nanomaterials requires additional synthetic techniques, and consequently additional reasoning strategies, beyond those of macroscopic synthesis. Theories and models developed around the concept of *surface stabilization* are needed to support successful nanosynthesis, whereas at the bulk scale surfaces are stabilized by the relative stability and larger proportion of interior atoms. In order to successfully limit growth of nanosynthesis systems, chemists have developed two kinds of tools. The first is a set of models that help chemists to *understand the mechanisms* of growth in a nanosynthesis system; in other words, they are *models of growth*. The second is a set of experimental methods, called *growth limitation techniques*, that help chemists to *manipulate the dynamics* of the synthesis. Models of growth include *growth-limited growth* and *diffusion-limited growth*, as well as models of nucleation and growth dynamics such as *LaMer's model* and *Ostwald ripening*. Growth limitation techniques include *physical growth limitation* and *chemical growth limitation*. I discuss models of growth and growth limitation techniques in some detail here in order to develop the ensuing discussion of how these tools combine to aid in the successful synthesis of nanoscale materials.

C. MODELS OF GROWTH

Nanosynthesis uses three main models of growth:

• Growth-Limited Growth: Models of growth provide information about the *rate*determining factors in a reaction, that is, factors that inhibit reactions from proceeding forward unchecked. In reactions driven by growth-limited growth (GLG), the overall amount of reactant is the rate-determining factor. The principle at work here is a basic conservation law: it is impossible to produce more of a product than is put into the reaction. If a chemist adds enough chloroauric acid (HAuCl₄) to yield 0.01mg of gold, she will not be able to yield 0.02mg of gold nanoparticles. GLG suggests that lower con-

¹Readers interested in the relationship between physical stability and various chemical concepts of stability may be interested in Roald Hoffmann's (1987) essay on the subject.

centrations of particles are more likely to produce smaller particles and/or grow them more slowly.

- Diffusion-Limited Growth: Diffusion-limited growth differs (DLG) from GLG in that DLG is limited by the accessibility of the particle's surface. Instead of the amount of gold in the reaction system determining how quickly a particle will grow, the amount of surface available for adsorption determines how quickly the particle will grow. Diffusion-limited growth (DLG) provides models for the dynamics of synthesis at a different scale than GLG. Whereas GLG is concerned with the entire reaction system, either in bulk-quantity terms or in terms of probability distributions over the motion of atoms, DLG is concerned only with the region near the nanoparticle's surface—that is, the *interfacial* region between the particle and its environment.
- LaMer's Model: LaMer's model differs from GLG and DLG in that it provides information about a continuous process of nucleation and growth, rather than just identifying factors that constrain growth. The model depicts continuous changes in the concentration of reactants over the course of a reaction (see Figure III.1). The concentration of reactants is divided into two regimes: below and above supersaturation. The diagram, and the mathematical and graphical models that support it, represent the idea that the shorter amount of time a reaction spends in supersaturation, the more uniformly distributed in size, or *monodisperse*, the resulting particles will be.

LaMer's model is often invoked as a sort of heuristic during nanosynthesis, and the basic relationship between duration of supersaturation and monodispersity is as deeply entrenched a guiding principle in nanosynthesis as the ideal gas law relations are in thermodynamics. Rarely, however, is the model tested against fundamental physical or chemical laws, or revised in light of new *ceteris paribus* conditions, or even invoked to represent the real dynamics of a real reaction as it proceeds. Rather, it is used more often to justify changes in experimental procedure or to explain the mono- or polydispersity of a synthesis. In the latter case, it is often used in conjunction with reference to Ostwald ripening, the last model of growth considered here.

• Ostwald Ripening: When a synthesis undergoes Ostwald ripening, its polydispersity (non-uniformity of size) increases over time according to deterministic physical principles



Figure III.1: LaMer's Model of Growth. [Creative Commons]

laid out in the model. Initial small differences in the sizes of individual particles become more pronounced as the reaction ages, owing to the greater stability of larger particles. Smaller particles give up the thermodynamic ghost, as it were, dissolving and adsorbing to larger particles. This reduces overall surface energy of the system both by eliminating the high-energy surface of the smaller particle and by decreasing the surface-to-volume ratio of the larger particle. Wilhelm Ostwald, one of the founders of physical chemistry, first observed the phenomenon, and it has since been explained via diffusion kinetics in (Lifshitz and Slyozov 1961).

In contrast with the growth limitation techniques I will discuss later, none of these models of growth correspond to actions that a chemist performs during an experiment. Rather, models of growth are ways of describing (or understanding, or explaining) the kinds of limits imposed on growth by dynamical laws and the initial and boundary conditions of the system. Importantly, these models do not in themselves suggest courses of action for modifying a synthesis to improve its yield. Instead they identify the manipulable relationships among parts of a system. With those relationships identified, a chemist may then use the information to design or modify a reaction procedure. Designing a reaction requires additional conceptual resources and new models, such as those that describe growth limitation techniques. These relations between models of growth and growth limitation techniques are neither reductive nor emergent; rather, the two kinds of models are *interacting*.

D. MODEL INTERACTION

Consider a traditional interpretation of the landscape of growth models I have described so far. DLG and GLG seem, at first glance, to cleave the set of possible solution-based syntheses cleanly in twain: either a reaction rate is limited by the overall quantity of reactant, or it is limited by the surface accessibility of reactants. The options seem exhaustive and exclusive, and the models appear to coexist non-interactively with equal claims to (non-) fundamentality.

The truth of the matter is somewhat more complicated than this. It point of fact, it is often the case that both DLG and GLG can both describe the dynamics of a single reaction. This occurs, for instance, when a small amount of reactant is left in a solution and the growing particles are also surrounded by a surfactant that limits surface accessibility. Such situations are not unfortunate accidents of the lab or poorly-implemented synthetic designs; they are the conditions needed to grow certain sorts of nanoparticle shapes (e.g. cubes). In such cases (and really, in all cases), GLG provides information about the behavior of the solution system and DLG provides information about the behavior of the surface structure of the growing particles—so, the two models describe the system at different scales. GLG describes the macroscopic behavior of the solution and its relation to the microscopic behavior of individual atoms, where DLG describes the mesoscale behavior of surfaces. When the solution concentration affects the eventual surface structure and overall shape of the synthesized particles, DLG and GLG have *interacted* to describe the growth dynamics of a given synthesis. The interaction proceeds without the need—or the ability—to reduce one model to the other, which is possible because of the different scales at which the two models apply.

The situation gets more complicated when one imagines the attempt to draw a map of the relationships between DLG and GLG, on the one hand, and LaMer's model or Ostwald ripening, on the other. No one model is a special or limiting case, a *ceteris paribus* exception, to the others. There is no hierarchical set of theoretical principles in which to embed the models. Instead, relations between the models rely on notions of scale-separation and contextual application rather than fundamentality, verisimilitude, or emergence.

More generally, by *model interaction* I mean that information from one model is being embedded in another model or hybridized with another, but without the baggage that comes with assuming one to be the more fundamental, less idealized, or otherwise closer to the real deep truth at the heart of things. Instead, for example, LaMer's model contains the piece of information, call it **J**, that monodispersity is more likely when the supersaturation period of the synthesis is shorter. That is not all of the information in the model, and the model contains a variety of additional relationships between monodispersity, reaction concentration, and time, as well as resources for explaining these relationships. But \mathbf{J} can be extracted from LaMer's model and embedded in a synthesis protocol in the form of a guiding principle K that says chemists should limit the supersaturation periods of their reactions. Additional models of, e.g., molar concentration of the precursor solutions and stoichiometric relations between the reactants can then be used to determine how to minimize a particular synthesis' supersaturation period. Depending on the complexity of the synthesis, additional conceptual resources may be needed to bring about K, but the point is that such models, including the molar concentration models, are neither conceptually prior nor posterior to LaMer's model. This is one sort of model interaction, call it *sequential* model interaction, where models can provide guiding information at differing stages of synthesis. The dynamics of model interaction can get significantly more complicated and are discussed at greater length in Chapters III.G and V.

Model interaction rationalizes, or at the very least makes room for, a robustly pluralistic set of models used to explain, describe, and synthesize a given set of target systems. Where both reductive and emergent approaches to inter-theory/inter-model relations suggest that in cases where multiple models describe a given target, it is preferable to assign one model (relative or absolute) fundamentality, model interaction explains why such projects often fall flat in actual scientific practice.

Fundamentality projects fall particularly flat in synthetic contexts, because the goals of synthesis are rarely to subsume behaviors under a natural law or unify across a variety of theoretical domains. Solving problems in nanosynthesis often requires that scientists treat the inputs from the various models of growth on even footing—that is, as equally useful tools for improving synthetic protocols that may or may not be applicable in a given case rather than as a top-down illumination of the principles of growth dynamics. The ultimate goal is the creation of individual nanomaterials, rather than the unification of a variety of nanomaterial behaviors under one dynamical model or the discovery of natural laws or causal mechanisms. So models of growth play a different role in the practice of nanosynthesis, and enter into different kinds of relationships with one another in synthetic settings, than dynamical models play in descriptive scientific practices.

E. GROWTH LIMITATION TECHNIQUES

Growth limitation techniques are synthetic strategies for constraining the development of nanomaterials to prevent them from growing above the nanoscale, and often to guide their growth in other ways as well. Multiple techniques can be (and usually are) used in a given synthesis. These techniques are as numerous and varied as the researchers who implement them—many nanochemists have proprietary or favorite growth limitation techniques—but they can be broadly classified as either chemical or physical.

Physical growth-limitation techniques constrain the spatial locations in the reaction environment where nanoparticles will form. In other words, they provide a material, physical constraint on the size and shape of nanoparticle growth, such as a solid template that is filled in with growing nanomaterials. In such cases, the template acts as a mold, channel, or barrier for the synthesis, preventing growth in certain directions. The template can be either reusable or sacrificial—that is, it may be used up and disintegrated during the course of the synthesis. An example of physical growth limitation is the use of graphene (singlelayer graphite) to template the growth of MoS_2 thin films, which ensures the synthesis of the desired shape, as well as more regular crystallinity and higher yield than untemplated syntheses (Shi et al. 2012).

In chemical growth limitation, an additional chemical species is introduced into the reaction to constrain growth or direct growth in a particular direction. A clear case of chemical growth limitation is the use of *capping agents* (discussed in Chapter I.D.1.d) as well as limiting growth by preventing further atoms of the nanoparticle material from attaching to the capping agent's adsorption sites. Capping agents are usually responsible for blocking the surface accessibility of a growing particle in DLG, and the overwhelming majority of nanomaterials require capping agents to prevent them from growing above the nanoscale. For instance, the gold nanomaterials that the Millstone Lab regularly studies can be capped by a variety of ligand shell types—sodium citrate, poly-(ethylene glycol) thiol, and others—and the relationships among these capping agents, and between the capping agents and the nanoparticles, has become a central research subject in the lab (cf. e.g. Smith et al. 2015).

One difference between physical and chemical growth limitation techniques is that the latter are much more useful for 0D nanomaterials than the former. While physical templates are very good at constraining growth in one or two dimensions, it is exceedingly difficult to physically constrain growth in three dimensions. By contrast, chemical growth limitation through the use of capping agents is comparatively easy to apply in three mutually orthogonal directions, constraining growth in three dimensions. This is because the constraining agent in chemical growth limitation is both of a similar scale to the nanomaterial, and because the chemical capping agent is quite spatially flexible compared to a physical template.

A comprehensive survey of individual growth limitation techniques could comprise a dissertation or two in itself, so rather than painting a landscape here, I will provide a few portraits of growth limitation techniques in order to develop an impression of the complexities and multiple scales involved in modeling nanoscale synthetic behavior, as well as the interaction between the two.

1. Vapor–Liquid–Solid Synthesis: Limiting Growth and Directing Growth

Consider the Vapor–Liquid–Solid (VLS) synthesis method, a well-known and popular experimental technique for growing nanowires. VLS synthesis occurs by first affixing a liquid droplet to a solid surface, then releasing gas-phase reactants into the reaction environment with the droplet-covered solid. When gaseous reactants come into contact with the liquid droplets, the reactants are absorbed into the droplet. Atoms of reactant filter through the liquid droplet to the liquid–solid interface, where they deposit in the solid phase. The deposited atoms arrange themselves to follow the crystal structure of the solid surface, and they grow vertically from the interface between the liquid and solid.

In VLS, the droplet acts as both a physical and a chemical constraint on growth, not by making growth harder in certain places but by making growth *easier* in certain places—it is relatively easier for reactants to precipitate out of the gaseous phase at the liquid–solid interface than by depositing onto the solid surface in the absence of the liquid droplet. This is because the liquid droplet lowers the surface energy of materials formed under its umbrella by effectively eliminating the energy costs of creating a new surface. So the droplet is a capping agent, in that it reduces the surface energy of the nanoparticle with which it interacts. But it also acts as a template, in that it provides a physical barrier between spatial locations in the reaction environment where the particle will form and spatial locations in the environment where the particle will not form.

Whether the droplet is modeled as a physical or chemical growth limitation technique depends on what design problem the modeler is trying to solve. If the problem is one of modifying the diameter of the resultant wire, then the droplet can be treated as a physical template, because its diameter controls the wire's diameter. If, on the other hand, the problem is one of changing the composition of the wire, then the droplet is better treated chemically, because the chemical interactions between droplet, substrate, and gas must be reconceived in light of the new composition. Again, the relationship between various reasoning strategies supporting synthesis is defined relative to the particular synthesis-design problem being considered.

In addition to acting as an example of both physical and chemical growth limitation and

growth direction, VLS provides an interesting look at how crystal-growth principles influence the size and shape of nanomaterials. When atoms of reactant precipitate through the drop and deposit onto the solid surface, the pattern of deposition aligns with the structure of the solid surface onto which the atoms deposit. That is, newly-precipitated atoms at the solid– liquid interface adsorb to the solid surface in the shape of the solid surface's crystal structure. This deposition pattern is a form of *epitaxial growth*, a term borrowed from macroscopic crystallography to describe the preservation and extension of a crystal structure from one material to another material in close proximity.

2. Anisotropic Metal Nanoparticles

The case of VLS growth introduces an additional role for growth limitation techniques, namely to *direct* the growth of a nanomaterial to produce a desired shape—a nanotube, a nanocube, a thin film, or a wire. Capping agents, templates, and DLG can all be applied to problems of directing growth as well as to problems of limiting growth. For instance, *selectively binding* capping agents can be used to promote anisotropic (asymmetric) growth. Selectively binding capping agents promote growth along certain crystal facets and not others, making them useful for growing both 1D shapes like nanowires and 0D shapes like cubes and triangular prisms.

Nanocubes are anisotropic crystalline nanoparticles that feature a number of interesting optical and electronic properties. They have exciting potential applications in drug delivery systems, cancer therapy, and detection of biological irregularities. So far, nanocubes with edge lengths between 80 and 175 nm have been synthesized from silver (Ag) and gold (Au), and hollow gold nanoboxes have also been observed (Wiley et al. 2005; Chen et al. 2005a,b). Researchers have concluded that the unusual anisotropic growth that permits the formation of these geometric nanoparticles is a result of preferential binding of capping agents on particular crystal faces, but they have not determined why capping agents bind preferentially to the cube-favoring crystal faces.

One particular capping agent, poly-(vinyl pyrrolidone) (PVP, $(C_6H_9NO)_n$), is the only one that has so far been observed to be conducive to the growth of silver nanocubes (Sun



Figure III.2: Schematic of vapor-liquid-solid growth. Left: liquid catalyst droplet (dark grey) on solid surface (black). Middle: Gaseous metal particles (light grey) filter through catalyst. to form metal nanowire (light grey). Right: Wire growth. [Author illustration]

and Xia 2002). The chemical literature does not presently contain a mechanistic explanation of why silver and gold, as opposed to other transition metals, undergo cubic synthesis; nor of why PVP and not other capping agents leads to cube formation.

PVP-mediated synthesis of geometric silver nanoparticles is achieved by adding silver nitrate (AgNO₃) and low concentrations of PVP to ethylene glycol ($C_2H_6O_2$) at 160 °C. In this synthesis, ethylene glycol acts to reduce Ag³⁺ ions to elemental silver. Nucleation occurs and particles stabilize at a critical radius once concentration of reduced silver is sufficiently high. At this point, PVP acts as a capping agent for the silver nanoparticles, binding preferentially to particular faces of the crystal and inducing diffusion-limited growth of the particles. The highly controlled temperature of the reaction allows for rapid nucleation, and thus rapid achievement of a critical concentration of elemental silver, which results in highly monodisperse particles—LaMer's model at work. Fine-tuning the addition of PVP allows for morphological control during the Ostwald ripening and DLG phases of the synthesis (Wiley et al. 2005). The principled and in some sense 'symmetric' anisotropic growth of silver nanoparticles into nanocubes is not seen in most other nanosynthesis processes, where anisotropic growth is more likely to produce long, thin particles such as nanowires, nanobelts, and nanorods.

One central problem in synthesizing silver nanocubes is to gain better control over the final morphology of nanocubes. Understanding the relationship between the concentration of PVP in the synthesis and the final morphology of the particles can help to solve this problem. Improving that understanding requires the interaction of a variety of reasoning strategies in a collaborative, interpolative manner.

As a first pass, it is worth considering what factors influence whether or not the silver seeds twin, as twinning will result in nanowires or quasi-spherical particles and the absence of twinning is required for nanocube synthesis. Xia et al. (2005) report that twinning can be limited by initiating synthesis with a relatively higher solution of $AgNO_3$ and maintaining a low ratio of $AgNO_3$ to PVP. They give two reasons for these conditions: 1) The rapid nucleation that results from higher precursor concentration means the reaction spends less time in the nucleation-and-growth stage of synthesis. This means that there is not time for twin defects to arise before diffusion-limited growth sets in, and 2) Twinning only occurs when the surface energy of {100} faces is higher than the surface energy of {111} faces, and PVP preferentially binds to {100} facets, which indicates that these faces will have lower-than-normal surface energy.

This latter feature of the synthesis provides the strongest clue for explaining the formation of nanocubes. Because PVP preferentially binds to $\{100\}$ faces, available growth sites on $\{111\}$ faces are the main ones that can accept new atoms during the growth phase. Filling in of $\{111\}$ sites on all faces of the truncated octahedral seed crystal out to the $\{100\}$ surfaces will increase the surface area of $\{100\}$ surfaces, at which point PVP caps will be more likely to bind. Growth will slow to a stop when the bulk of the surface is $\{100\}$ and PVP has capped all surfaces of the crystal.

So in some sense preferential binding controls the formation of nanocubes, but the real question, namely what controls preferential binding, still remains. Since anisotropic growth of this sort occurs with PVP and not with other, chemically similar capping agents such as poly-(vinyl alcohol) or poly-(ethylene glycol thiol), the identity of the capping agent must play a significant role in the occurrence of this geometric form of anisotropic growth. I have not found a discussion in the chemical literature of what features of the capping agents might

lead to this particular sort of preferential binding, but it seems most likely that the polymer's steric features, electronic features, or some combination thereof must be responsible. It is worth speculating on each feature, but I focus at present on steric considerations.

Xia et al. (2005) discuss the difference in preferential binding between PVP and another common capping agent, sodium citrate, in their paper, specifically noting that sodium citrate does not bind preferentially to {111} faces. It is worth mentioning that in inorganic settings, sodium citrate is frequently used as a chelating agent for transition-metal complexes (Hornyak et al. 2008). It is thus not inconceivable that the binding action of sodium citrate on silver nanoparticles would demonstrate properties akin to chelation, specifically that the ligand could bond to two empty coordination sites on a surface atom of a silver nanoparticle. If this is indeed the binding action of sodium citrate on silver nanoparticles, then it is probably the case that the caps are rather bulkier than PVP caps, which bind singly, albeit with a bulky aromatic ring in the polymer. The added bulk of sodium citrate caps would have two related effects that may explain preferential {100} binding: 1) It would sterically hinder the formation of {100} surfaces, which require less closely-packed surface atoms than {111} surfaces, and 2) It would lower surface energy compared to PVP caps with single bond attachments, since surface atoms would have additional bonds to reduce the strain brought on by having empty coordination sites.

* * * *

The point of going into this much detail is to illustrate the blooming, buzzing confusion of models that are needed to represent and rationalize the various synthetic constraints on metal nanocube synthesis. Clearly, the (admittedly oversimplified) traditional, hierarchical view of inter-model relations I criticized in Chapter II cannot handle these sorts of cases. Model interaction can, once it is coupled with a conceptual tool that is endemic of synthetic reasoning, namely, *tuning*. Tuning is the use of multiple inter-constrained models to determine the inputs requisite to produce a target output. Unlike more characteristically descriptive modes of model use, such as prediction, tuning requires a keen sensitivity to the relationships among input variables in a given modeling scenario. It is the reasoning strategy employed in adjusting the temperature of a reaction in response to a kinetic model's depiction of a change in reaction rate. An example in the following section, and extended discussions of tuning in Chapters IV and V, illustrate.

F. TUNING IN NANOSYNTHESIS

Anisotropic metal nanoparticles, or AMNPs, are 0D metal nanomaterials with non-spherical geometries such as cubes, prisms, discs and wires. AMNPs are commonly grown in colloidal solutions, and to the naked eye these solutions look like innocuous, if brightly colored, fluids, not unlike Kool-aid (see Figure III.3). Rather than resulting from traditional pigments or dyes, these colors are the result of optical phenomena that only occur at the nanoscale, most commonly LSPR or quantum confinement. These phenomena can be manipulated and controlled by controlling the development of the nanoparticles, and the phenomena can be used in technologies as diverse as cancer therapy, drug delivery, surface-enhanced spectroscopy, and solar cells (Sun and Xia 2002; Yavuz et al. 2009; Liao et al. 2006).

An experimental therapeutic technology known as *thermal ablative therapy* consists of injecting a solution of colloidal nanoparticles into the site of a malignant tumor then activating an LSPR response in the nanoparticle solution (Loo et al. 2005). Activating LSPR in these materials generates very high temperatures in very local (around 10nm) regions around the surface of the material, which means it is possible to cause cell death via heat exposure in much more localized target areas than with chemotherapy or traditional radiation therapy.²

In order to use LSPR-active nanomaterials in these or any other settings, the materials must first be synthesized. Synthesizing LSPR-active nanomaterials requires using a variety of theories and models to tune the development of the materials during synthesis. In this section, I describe how theories and models are used for the synthesis of LSPR-active

²At the time of writing, thermal ablation therapy has not been approved for use in human patients. Like many medical applications of nanomaterials, the long-term effects of the therapy have not been assessed to researchers' satisfaction. This sort of testing of nanomaterials is, unfortunately, outside the scope of the present discussion. For philosophical discussion of the biological and ethical implications of nanotechnology, see, e.g., Sandler and Kay 2006; Sandler 2009.



Figure III.3: Coloration of gold nanoparticles in colloidal solution, variation by size. [Courtesy Millstone Lab, University of Pittsburgh]

nanomaterials for cancer therapy. I begin in Subsection III.F.1 by describing the physical phenomenon of LSPR. Section IV.C describes the ways that a material's LSPR response can be controlled by tuning the synthetic conditions under which the material is made, and Section IV.D describes the refinement of these synthetic conditions over the course of an experiment. Throughout these discussions, I refer to the interactive use of multiple models and theories to obtain synthetic conditions that will produce the desired outcome.

1. Modeling LSPR

A *plasmon* is a kind of quasiparticle: it is a unit of plasmic oscillation, which is to say it is a packet of energy arising from a group of electrons in a material as the electrons oscillate between higher- and lower-energy states in a conducting medium (i.e. a plasma). Plasmons are generated by the excitation of electrons. The subsequent relaxation of those electrons gives off energy, and the excitation–relaxation cycle produces a plasmon. Electrons can be excited either by other electrons or, more commonly, by photons. *Plasmon resonance* occurs when the energy of incident photons matches the energy of the electronic transition from the higher- to lower-energy state. Where a non-resonant plasmon leads a momentary existence, resonant plasmons endure so long as the material is being stimulated with incident photons—it is the difference between plucking a fiddle string and sawing it with a bow.

Surface plasmons occur when the plasmon is generated at the surface of a material, and surface plasmons can be either localized or propagating. Propagating surface plasmons, like the wake of boat in the water, spread out from their origin over the surface of a material. Localized surface plasmons, on the other hand, are bounded like a vibrating string, and they only occur when the wavelength of the plasmon is significantly wider than the width of the material in which the plasmon is induced. The wavelength range of visible and near-infrared light is about 400–1200 nm. So, visible- and infrared-light-induced localized surface plasmon resonance can only exist in materials with nanoscale dimensions. LSPR is responsible for the colors of the vials in Figure III.3. As mentioned in Chapter II, the colors in stained glass are also the result of sunlight-induced LSPR in metallic nanoparticles embedded in the glass.

It is worth noting that this description of what LSPR is does not indicate how to make an LSPR-active nanomaterial, and indeed this description is somewhat irrelevant to the project of making LSPR-active nanomaterials for thermal ablation therapy. The converse also holds true: being able to successfully make an LSPR-active nanomaterial does not require a thick understanding of what plasmons or what plasmon resonance is. I have found this much to be the case in my discussions with chemists, but the truth of it can also be seen in the fact that few of the models used to tune the synthesis of LSPR-active nanomaterials rely on knowing what LSPR is.

Of models used to tune the synthesis of LSPR-active nanomaterials, the Mie solution to Maxwell's equations contains the most information about LSPR itself. The Mie solution is an approximative solution to Maxwell's equations for a spherical particle:

$$E(\lambda) = \frac{24\pi N_A a^3 \epsilon_m^{3/2}}{\lambda \ln(10)} \frac{\epsilon_i}{(\epsilon_r + \chi \epsilon_m)^2 + \epsilon_i^2}$$

Here $E(\lambda)$ is electromagnetic extinction, or intensity of LSPR response. λ is the wavelength of incident light. N_A is the electron distribution of nanoparticles, which is determined by the composition of the material. a is the radius of the material particle; ϵ_m is the (compositiondependent) dielectric of the environment surrounding the particle; and ϵ_r and ϵ_i are the

Variable	Parameter	Tunable Feature
λ	incident wavelength	energy of light inducing LSPR
N _A	electron distribution	composition of nanomaterial
a	radius	size of nanomaterial
ϵ_m	dielectric of environment	composition of environment
ϵ_r	real-component dielectric of material	composition of nanomaterial
ϵ_i	imaginary-component dielectric of material	composition of nanomaterial
χ	aspect ratio factor	shape of nanomaterial

Figure III.4: Breakdown of variables in Mie solution into tunable features of a nanomaterial.

real and imaginary components of the particle's (composition-dependent) dielectric, and χ is a factor determined by the aspect ratio, or relative length to width, of the particle.³ To synthesize a system that exhibits LSPR with a large $E(\lambda)$ in response to an incident wavelength λ is to generate a tuned LSPR response.

Table III.4 summarizes the five synthetically-controllable variables that affect the particle's LSPR intensity, as measured by electromagnetic extinction. Extinction is the sum of electromagnetic absorption events and electromagnetic scattering events in a region over a period of time. Incident photons produce absorption events and oscillations of a plasmon produce scattering events. High, narrowly-peaked maxima on extinction graphs indicate places of strong absorption and scattering, which is interpreted as plasmon resonance. (See Fig. ??.) Tuning LSPR in a metal nanomaterial, then, is a matter of figuring out what input conditions will lead to the synthesis of a material that responds to a given incident wavelength of light with a large, narrowly-peaked extinction spectrum.

In any system of interrelated variables, such as the Mie solution or the equations of motion governing the block sliding down a plane, assigning a determinate value to one

³For spheres, $\chi = 2$. Different shapes such as cubes or wires take on different values of χ , but these shapes also require other modifications of the solution, which are better obtained through numerical-methods approximations to Maxwell's equations than through analytic solutions like the Mie solution (Haes and Van Duyne 2004).

or more variables constrains the possible values of the others. Tuning works by assigning a determinate value to a given variable, the target output, and figuring out what inputs will lead to that output. This process of selecting inputs for the output is not, however, usually a matter of straightforwardly plugging and chugging. Once an incident wavelength is selected, and given that the goal is to maximize extinction, there are still a variety of intervariable constraints between composition, shape, size, and environment. Moreover, while the Mie solution captures some of these relations, most of the constraints between composition, shape, size, and environment are not represented in the Mie solution, nor could they be by any reasonable thickening of the mathematical stew, because they vary significantly from one kind of nanomaterial to the next. Relations between size and shape for gold are different than relations between size and shape for silver, and nanomaterials in organic environments tend to take on different shapes than similarly-produced materials in aqueous environments.

Call these sorts of relations *synthetic constraints*. Synthetic constraints are often amassed as sets of rules of thumb or heuristics, rather than full-blown models, and they are pieces of information about the way that chemical substances behave when they interact with one another. The theory of acids and bases is a type of synthetic constraint, as are tables of miscibility, that is, lists of whether or not two solutions will combine with one another. So, it is clear that multiple models are used in solving the present tuning problem. The Mie solution does not contain information about synthetic constraints, so it is inadequate on its own for solving the problem of tuning LSPR for thermal ablation therapy.

Synthetic constraints are used in solving tuning problems, and only accounts of modeling that are sensitive to the varieties of inter-theory/inter-model relations needed to describe these constraints can adequately capture the dynamics of model use in synthetic contexts like the ones considered in this chapter. While the traditional account of inter-theory/inter-model relations I have caricatured here clearly cannot provide the right sort of context-sensitivity, there are some contemporary accounts that can. These include Mitchell's (2009) *integrative pluralism*, Chang's (2004) *epistemic iteration*, and Woodward's (2003) *manipulationism*, each of which articulates feedback dynamics between different models or theories as information among inter-related models is what I will in Chapter IV call *iterative interpolation*.

G. MODEL INTERACTION AND MULTI-SCALE MODELING

There is another sort of recent philosophical account that can accommodate model interaction: contemporary accounts of models and theories that focus on scale-sensitivity. Many of these accounts (cf. e.g. Wilson 2012; Batterman 2012; Morrison 2014) rely on indexing certain modeled phenomena to particular length, time, or energy scales and then explaining particulars of relations across scales. One such account, namely Eric Winsberg's "handshaking" view of inter-model relations, draws some of its details from simulations of nanoscale material behavior.

Winsberg's account is part of his larger aim to develop an epistemology of computer simulation, and so he is interested in inter-model relations more as a means than an end. But along the way he encounters the problem of how to make sense of simulations whose component models come from a variety of theoretical backgrounds, as is often the case in nanoscience. In the nanoscale simulations Winsberg discusses, the problem at hand is how to simulate crack propagation in nanomaterials. The simulations he discusses use three distinct computational models to describe three distinct sorts of dynamics that a nanoscale system experiences, each at characteristic length scales. Continuum mechanics describes large-scale behaviors of the system, molecular dynamics describes somewhat smaller domains that behave like classical rigid bodies, and quantum mechanics describes the smallest interactions among individual atoms. In the multiscale simulation Winsberg describes, the models are combined "in parallel" such that at each time-step of the simulation, each of the models contributes some mathematical quantity to the overall description of the system. The models are combined by developing computational techniques to average over boundary regions where a higher-scale model and a lower-scale model disagree on a predicted quantity. This method of combining models, a form of model interaction, is the phenomenon that Winsberg, following Broughton (1999), calls "handshaking."

The handshaking process Winsberg describes in the fifth chapter of *Science in the Age* of *Computer Simulation* (Winsberg 2010, pp. 72–92), drawn from Broughton et al.'s (1999) and Abraham et al.'s (1998) work in computational physics, is the development and implementation of algorithms to combine state descriptions from smaller- and larger-scale models of a solid-state physical system. Winsberg argues that these algorithms illustrate the same non-hierarchical relationship between smaller-scale and larger-scale models that I identified in Chapter II.B's violin example:

One issue that has received perennial attention from philosophers of science is that of the relationship between different levels of description. Traditionally, the focus of this inquiry has been debate about whether or not, and to what extent or in what respect, laws or theories at higher levels of description are reducible to those at a lower level. Underlying all of this debate, I believe, has been a common intuition: the basis for understanding interlevel interaction—to the extent that it is possible—is just applied mereology. In other words, to the extent that the literature in philosophy of science about levels of description has focused on whether and how one level is reducible to another it has implicitly assumed that the only interesting possible relationships are logical ones—that is, inter-theoretic relationships that flow logically from the mereological relationships between the entities posited in the two levels. But if methods that are anything like those described above become accepted as successful in nanoscale modeling, that intuition is likely to come under pressure. The reason is that parallel multiscale modeling methods are forced to develop relationships between the different levels that are perhaps suggested, but certainly not logically determined, by their mereology. Rather, developing the appropriate relationships, in Abraham's words, "requires physical insight." What this suggests is that there can be a substantial physics of interlevel interaction—a physics that is guided but not determined by either the theories at each level or the mereology of their respective entities. Indeed, whether or not the relationships employed by Abraham and his group will turn out to be the correct ones is an empirical/physical question and not a logical/mereological one. (Winsberg 2010, p. 84–85)

Winsberg is, ultimately, much more interested in representational relationships of the sort typical to the models literature than he is in inter-theory relations, and so he does not prod the details of the simulation once he has made this point that in simulations like this one, inter-level interactions are empirically determined. This is, I think, a mistake, as it overlooks the question of *why* the simulation manages to successfully model propagating-crack systems, even in the absence of a determination of whether the inter-level interactions are, in point of empirical fact, "the correct ones." I take it that part of Winsberg's epistemology-ofsimulation project is to rationalize the use of simulation models to make scientific progress in just these sorts of cases, namely when the complex, multiscale dynamics of the modeled system are as-yet unknown.

In the case of these crack-propagation simulations, it turns out that details of the mathematics of the simulations—and in particular, the sorts of assumptions built into the averaging techniques used to handshake across scales—answers the question of why these models are non-hierarchically combinable in the parallel-multiscale mode that Winsberg describes. In the example from Broughton et al (1999) that Winsberg develops, there are three length scales of interest and, therefore, two handshaking algorithms. The macroscopic length scale describes regions of a solid-state system whose dynamics are close to equilibrium, and it is modeled by finite-element (FE) methods, which are derived from the continuum-mechanical elastic theory of solids. In particular, the model divides a continuous volume into cells and describes the kinetic energy (displacement) and potential energy (strain) on individual cells (hence, finite elements). These cells are joined together into a mesh, and cell borders are called mesh points. One of the paper's innovations is their particular method for determining overall displacement and strain as a function of values of those quantities at the mesh points. So their model discretizes the continuum description. This will be important shortly.

The mesoscopic length scale describes regions of the system that are slightly perturbed from equilibrium but which are nonetheless not dynamically central to the simulation—in crack-propagation models, these are not the areas where bonds are breaking and forming, but instead the trailing wake of dynamical disturbance left by a propagating crack. These regions are modeled by molecular dynamics (MD), which is derived semi-empirically from statistical mechanics. Molecular dynamics is what it sounds like, that is, a model of the movement of individual atoms or molecules. It is a classical (non-quantum) model. Since the model is of a solid, it is populated as a lattice of atoms whose movement (vibration, rotation, and the breaking and forming of bonds) is described by interatomic electronic potentials. Think of a network of balls connected by springs.

To develop a handshaking algorithm for these two models, Broughton⁴ begins by drawing an imaginary surface between a region described by the FE model and one described by the MD model. The interatomic potentials governing the movement of atoms in the MD model range across that surface into the FE model. Since the FE mesh describes continuously distributed matter, the mesh points can be lined up anywhere along the surface—including right along the lattice of MD interatomic potentials. As Broughton puts it, what the algorithm needs is a "one-to-one mapping of a mesh point to an atom site." (Broughton et al.

⁴For narrative ease, I refer to the first author of the paper in this explication. Unless otherwise noted, references and page numbers attach to (Broughton et al. 1999).

1999, p. 2396) Farther away from the handshake interface, the mesh can be spaced out for computational efficiency. The energy (Hamiltonian) at a point on the interface is calculated by taking the arithmetic mean of predictions from each model.

I want to pause here to highlight exactly why this handshake algorithm works, because it is important to realize that the effectiveness of the energy-averaging strategy the physicists end up with is contingent on both features about the world and features about the computational models. Averaging the Hamiltonians is effective exactly because the MD lattice has been lined up with the FE mesh at the contact region; otherwise there would not be commensurable values along the imaginary surface at the handshake region. That alignment is possible because the FE mesh describes a region of the system where continuum mechanics is the appropriate description of physical behavior, and the continuous deformability of the mesh is what allows modelers to line mesh points up with the MD lattice. The macroscopic scale, and associated near-equilibrium, of the FE-modeled region of the system, is what licenses the use of continuum (FE) methods to describe it.

So the physics of the system, namely the scale-dependent deviation from equilibrium, constrains the application of different models to different portions of the system. This feature of the world, namely the scale-dependence of the system's behaviors of interest, is reflected in the selection and application of models to the system—and the resulting picture of intermodel relations happens to be neither reductive nor emergent, but instead to exhibit model interaction. The models applied to the difference scales of interest are in turn constrained by a combination of scale-dependent physics, mathematical descriptions thereof, computational effort associated with implementing each model⁵, and, somewhat less so, by other pragmatic constraints, such as the particular goal of the simulation and the theoretical background and imagination of the modeler.

To emphasize the role of particularities of the physical details in understanding the development of cases of model interaction, I proceed first by working through the other algorithm in the Winsberg/Broughton handshaking example, namely the movement from MD models to quantum-mechanical models of bonds breaking and forming at the site of a propagating crack in a solid. In contradistinction to the MD/FE handshake region, Broughton deems the

⁵MD is much more computationally costly than FE. For more on this see Winsberg or Broughton.

handshake region between MD and the tight-binding (TB) model of quantum chemical and physical behavior to be the "active zone" (p. 2398) of the system and, therefore, of different physical interest than the MD/FE region.

One of the distinctions with which Broughton begins his discussion of the TB/MD handshake algorithm is that the TB/MD models both proceed under assumptions of atomicity, whereas in the FE/MD region the FE model's mesh points are not meant to represent spatial points—indeed just the opposite, mesh points are meant as abstractions away from spatial details of material behavior.⁶ So the problem of connecting the spatial regions of the system that are being represented by the different models is going to be solved differently than in the MD/FE case, where the FE mesh plane was deformed to line up with the MD lattice plane.

Both TB and MD can model atoms in a lattice and their electronic interactions. So the algorithm constructed to combine them is composed of operations performed on a specially-defined set of simulated "atoms." In Broughton's example of crack propagation in a silicon material, these "atoms" are called "silogens," because they have some silicon-like properties and some hydrogen-like properties (the latter introduced for the sake of localizing electronic behavior so it can be modeled by MD). The difficult (interesting) move in this combination of models comes in assigning properties to the silogens. Some properties, such as the spacing in a crystal lattice, are defined by the dynamics described by the MD model. Other properties, such as electronic symmetries, come from TB. The result is a messy, chimeric picture of silogens in the interfacial region, where multiple "atoms" exist at a single lattice site, and these "atoms" behave neither entirely like silicon nor entirely like hydrogen.

The handshake algorithm assigns contrived, un-physical electronic behaviors to silogens (i.e. reducing the number of electronic symmetries associated with atoms of a particular element) and to their near neighbors, in order to generate predictions about the local energetic behavior of the interfacial region. By reducing electronic symmetries, the models develop a localized description of the energetic behavior of the handshake region, whereas in the non-handshake region of the quantum-mechanical TB model, energetic interactions are

⁶It remains curious to me that despite the modeling literature's interest in representation, the difference in representational intent between continuum and molecular models is overlooked by many a discussant.

delocalized.⁷

Here again modelers are constrained by the physics (e.g. by the genuinely quantum behavior of the TB region, which cannot be classically predicted) and by the computational costs of their models, but they are able to choose some features of the handshake region (such as the reduced symmetries) to best suit their needs. However, the kinds of *conceptual influences* on their considerations of physical constraint and computational costs differ greatly than those in development of the FE/MD handshake algorithm.

Broughton puts the point thus:

In contradistinction to the FE/MD handshake algorithm, where a plane between rows of atoms was defined, the MD/TB handshaking takes place conceptually across a plane consisting of atoms. This different approach is necessitated because it is difficult to apportion (localize) energy in a computationally efficient way to specific bonds in an electronic structure calculation. The total energy is a property of the entire system. Attempts to define a 50/50 Hamiltonian, such as was used for the FE/MD interface, run into issues ... (Broughton et al. 1999, p. 2398)

The basic idea is this: one can treat computations of the modeled system's energy as *conceptually* the same between the MD and FE regions, although the *entities* that produce changes in that energy differ. However, the conceptions of energy required to make sense of the MD and TB models are really quite different from one another, but the *entities* are, while not identical, both discretized and inter-relatable. So to combine the MD and TB models, different strategies are needed than those used to combine FE and TB—and those strategies require different conceptual resources.

Here is another way of thinking about the relationship between the two handshake strategies. In the MD/FE handshake, the modeling strategy that gets the algorithm off the ground is kinematic—lining up the FE mesh points with MD lattice spacing. On the other hand, in the TB/MD handshake, the thing that needs to happen for the algorithm to go forward is dynamic—changing way that the behavior of atoms is modeled. In neither of the handshakes is one model conceived of as containing "truer," "more correct," or "better" physics, and in both handshakes, each participating model is conceptually and computationally compromised in order to develop a model of the overlap. The handshakes are similar in that physical

⁷This is another instance of Wilsonian "physics avoidance," and strategies of this sort play a leading role in Batterman's (2001; 2012) "asymptotic" account of inter-theoretic relations.

details of the modeled behaviors at each characteristic length scale in the system constrain, but do not dictate, the ways in which handshaking algorithms can be developed. But what conceptual resources are accessed to develop a handshake differs greatly from case to case.

The MD/TB handshake postulates new entities, where FE/MD does not. MD/TB then runs dynamical simulations on these un-physical entities and their near neighbors in the MD and TB regions, where FE/MD instead develops a contrived kinematic picture and averages dynamical results from both simulation models across this very particular region of simulation-space. Winsberg calls both these strategies "handshaking" and does not distinguish between them, but the ways in which the models interact, and the sorts of constraints put upon the component models, vary from case to case.

Winsberg is right to identify that the relationships between these three component models are better understood as multiscale or interactive, rather than reductive or emergent. There is no fundamental model that either represents most accurately or explains most lucidly the behavior of the target simulated system. However, by lumping the interactions between Broughton's models into one category, Winsberg overlooks the intricate system of physical, computational, and pragmatic constraints upon the individual handshake algorithms. By attending to precisely these differences, one can glean that the algorithms contain different strategies for combining models across scales. These strategies in turn generate different representational content, as well as different ways of rationalizing both the permissibility and effectiveness of modeling the crack propagation behavior in this multiscale mode.

The reason for these differences is exactly the sort of scale-dependence I described at the beginning of this section and throughout the discussions of nanoscale material behavior in this dissertation. Morrison (2014) has lately demonstrated that in addition to scaledependence, systems must demonstrate *scale separation* in order to get multiscale modeling strategies such as Broughton's off the ground. In my extended discussion of Broughton's algorithms for combining the component models of his simulation, I showed that there are in fact a plurality of such strategies available for combining models across length scales, and that Winsberg's "handshaking" category is too broad to capture the particulars of these strategies. I call these strategies "model interaction" instead, and I have gathered a few varieties of model interaction—from the positing of fictional, unphysical entities and the manipulation of a model's kinematic background to the iterative movement between conflicting assumptions about the continuity of a physical system—in the examples discussed throughout this chapter. These join the previous chapters' examples of hybrid models and sequential model interaction in the growing menagerie of kinds of model interaction.

One final point worth noting is that in developing both the FE/MD and the MD/TB algorithms, Broughton needed to attend to the boundary-region behavior of the component models. Indeed, these boundary regions are, more or less by definition, where the models interacted. Just like in the previous chapter's violin example, the challenge of developing a multiscale model lay in accounting for the behavior at the boundary between regions of the system modeled by a higher-scale model and regions modeled by a lower-scale model. In this way, one might recover some sense from "lumpy" accounts of model interaction like Winsberg's: while there are many differences in the specific physical, computational, and pragmatic constraints associated with combining models across scales, these constraints all inter-mingle in an attempt to solve the same sort of problem, namely how to account for behavior in the boundary regions of modeled systems.

IV. HOW THEORIES (AND MODELS) SHAKE HANDS

Philosophers of science have written extensively about the mechanics of theory and model use in activities like prediction, explanation, and description. These activities are chosen because they are, either implicitly or explicitly, taken to be some of the central aims of scientific enterprise. But much of scientific activity is about *making*, *producing*, or *synthesizing*, rather than describing or explaining. Here I present an account of the mechanics of theory and model use when the goal of a scientific activity is synthesis, rather than description. The primary use of theories and models in nanosynthesis is to *tune* the input parameters in a synthesis reaction in order to produce a target product. I contrast this use of theories and models, which I call *the tuning account*, with explanatory and predictive uses of theories and models.

A. INTRODUCTION

Philosophers of science have written extensively about the mechanics of theory and model use in activities like prediction, explanation, and description. These activities are chosen because they are, either implicitly or explicitly, taken to be some of the central aims of science. But much of scientific activity is about *making*, *producing*, or *synthesizing*, rather than describing or explaining. This chapter presents an account of the mechanics of theory and model use when the goal of a scientific activity is synthesis, rather than explanation or description. This account leads to a new view of inter-theoretic relations, which dissolves a variety of contemporary worries about reduction, emergence and taxonomies of theories and models. I use this inter-theoretic relations result to explain why so little of contemporary philosophy of science has addressed the synthetic sciences, and I suggest that philosophers of science may resolve additional contemporary debates by turning their attention to these sciences.

To begin, consider the following collection of statements from philosophy about aims of science and the role of theories in supporting those aims:

- I. Science aims to give us, in its theories, a literally true story of what the world is like. (van Fraasen 1980, p.8)
- II. Science aims to give us theories which are empirically adequate. (van Fraasen 1980, p.12)
- III. Science aims to provide the best possible explanatory account of natural phenomena. (Ellis 1985, p.51)
- IV. What is the role of the sciences in a democratic society? Some people, let us call them the "scientific faithful," say this: "The sciences represent the apogee of human achievement. Since the seventeenth century, they have disclosed important truths about the natural world, and those truths have replaced old prejudices and superstitions. They have enlightened us, creating conditions under which people can lead more satisfying lives, becoming more fully rational and more fully human. The proper role of the sciences today is to continue this process, by engaging in free inquiry and by resisting attempts to hobble investigations for the sake of any moral, political, or religious agenda." (Kitcher 2001, p.3)
- V. The three cardinal aims of science are prediction, control, and explanation; but the greatest of these is explanation. Also the most inscrutable: prediction aims at truth, and control at happiness, and insofar as we have some independent grasp of these notions, we can evaluate science's strategies of prediction and control from the outside. Explanation, by contrast, aims at scientic understanding, a good intrinsic to science and therefore something that it seems we can only look to science itself to explicate. (Strevens 2006)

Statements (1)-(3), collected from Ellis's criticism of van Fraassen, suggest an image of science as an attempt to coat the world in description—no matter whether the description is of causal relations between seen and unseen phenomena or noumena, or simply a description of correlations between those phenomena that are empirically observable. Science collects data by observing, and then it interprets those data, and it structures its theories around

them. Statement (4) reinforces that image. Although it is far from Kitcher's own view of the aims and methods of science, its emphasis on science as truth-revealing continues the theme that science is primarily observational. Kitcher's own view is still a view of "inquiry," (Kitcher 2001, p.44) and despite the fact that his view "allows a place for human values and human interests in the constitution of the goals of the sciences," (Kitcher 2001, p.44) his primary concern is still how to use science "for the purpose of describing nature." (Kitcher 2001, p.45) Even though he, like many constructivists before and after him, believes that humans make or modify the world around them as they develop scientific theories, he writes about how humans make theories, rather than how theories help humans make things.

On the other hand, the reference to control in statement (5) acknowledges that the goals of science range wider than description alone. While Strevens' mention of control in the quote above is oblique and somewhat dismissive, Woodward (cf. e.g. Woodward 2003) and other so-called 'interventionists' have made control a central theme in recent accounts of causal explanation. They argue that identifying control- or manipulation-based relationships between phenomena is both the goal and the method of causal explanation in the sciences. On this view, theories organize information about what would happen to the output phenomenon if input conditions were altered or, to paraphrase Woodward, information about how a system of scientific interest changes when its levers are wiggled.

Still, the interventionist theory is ultimately a theory of *explanation*, an account of how theories help scientists and laypeople to make sense of the world. Knowing how a system responds to having its levers wiggled is not the same as actually wiggling its levers, although the two are clearly related. And when a scientific theory is being used to make something, the theory is useful insofar as knowing how levers wiggle can help to generate the particular desired outcome, whether or not further descriptive or causal-explanatory goals are met.

For instance, consider the project of synthesizing nano-cubes such as those in Figure IV.1. The goal of this project is to bring into existence some new stuff that used not to exist, namely the cubes themselves. These cubes are synthetic objects: they would not occur in nature without the coaxing of a human hand (with the help of accurate, precise, tailored instruments). In order for synthetic objects like these cubes to exist, some people had to decide that they wanted to make something, and then they had to figure out which



Figure IV.1: Scanning electron microscope image of $AuCu_3$ nanocubes. Average edge length: 44.9 nm. Cubes are capped with 1-dodecanethiol. [Reprinted with permission from Liu and Walker 2010]

Science	Title	
Biology	"Tuning the Dials of Synthetic Biology" (Arpino et al. 2013)	
Neuroscience	"Correlations on Ion Channel Expression Emerge from Homeostatic Tun-	
	ing Rules" (O'Leary et al. 2013)	
Physics	"Electrical Tuning of Valley Magnetic Moment through Symmetry Control	
	in Bilayer MoS_2 " (Wu et al. 2013)	
Materials Science	"Tuning Molecular Adhesion via Material Anisotropy" (Zhang et al. 2013)	
Chemistry	"Tuning the Surface Chemistry of Pd by Atomic C and H: A Microscopic	
	Picture" (Aleksandrov et al. 2013)	
Climate Science	"Climate Models Sensitive to Tuning of Cloud Parameters" (Balcerak	
	2013)	
Nanoscience	"Tuning the Electrical and Optical Properties of Graphene by Ozone Treat-	
	ment for Patterning Monolithic Transparent Electrodes" (Yuan et al. 2013)	

Figure IV.2: A cross-section of articles published in top journals in 2013 that refer to tuning. Among articles published in 2013, Google Scholar searches show over 10,000 results for "tuning" in each of these fields besides climate science. These numbers are similar to those for "explanation" in each of these fields, and "tuning" is more common than "explanation" in both nanoscience and materials science.

levers to wiggle to make it. Using theories and models to figure out which levers to wiggle in order to make something is how theories and models are used in the synthetic sciences. In Section IV.B, I argue that this use is distinct from the ways in which scientists use theories and models to predict, explain, or describe. Section IV.C offers an account of the mechanics of tuning through an extended example of tuning LSPR in nanosynthesis, and Section IV.D considers the diachronic aspect of tuning, a phenomenon I call *iterative interpolation*. Section IV.E orients this view of inter-theoretic relations within contemporary philosophy of science and uses the tuning account to sketch the foundations of a general account of scientific reasoning in applied science.

B. TUNING: THE SYNTHETIC USE OF THEORIES AND MODELS

Tuning is a term used in synthetic sciences to refer to the use of theories and models to produce certain desired synthetic consequences, like the creation of the gold nanocubes in

Figure IV.1. As evidenced by Table IV.2, tuning is a widespread scientific activity. It is also one that philosophers of science have largely overlooked. Tuning indicates the use of theories and models for the purposes of bringing about a particular, targeted outcome, rather than merely describing or explaining an outcome.

To get a clearer picture of what is meant by tuning, consider the block sliding down the plane in Fig. IV.3). This block is modeled by the graphical depiction in the figure and by the equations describing its motion, which are in turn derived from the laws of Newtonian classical mechanics. Solving problems about this block requires the use of these models and this theory, and different problems will make different use of the models and theories. Here are three kinds of problems, and three ways of using models and theory:

- I. The problem of predicting how fast the block will be moving at the bottom of the ramp, if it begins from rest: Solving this problem requires identifying the equations of motion that describe all block-ramp systems of this type, e.g. $F_{g_x} = -mg\sin\theta$, and the initial and boundary conditions specified by the diagram, e.g. the block's mass is 1kg and the length of the ramp is 2m. From these laws and parameters, it is possible to derive a value for the speed at the bottom of the ramp. The simplest solution, then, is, e.g.: The block will be traveling 5.26 m/s at the bottom of the ramp. While this is sufficient for a solution to the prediction problem, often problems of this sort are more satisfying when the steps of the derivation are explicated. To that end, a fuller solution to this problem would be: From Newtonian mechanics, $V_f^2 - V_i^2 = 2ad$, where V_f and V_i are final and initial velocities, a is acceleration and d is displacement (in this case, 2m). a is $mg\sin\theta$, or 9.8 sin $\pi/4$. Since V_i is 0, V_f is $\sqrt{2 \times 9.8 \sin \pi/4 \times 2} = 5.26 m/s$.
- II. The problem of explaining why the block slid down the ramp at all: The Newtonian theory of gravity, because it is a universal law, is recognized to apply to this system. The influence of gravity on the block can be spelled out in terms of mark transmission or counterfactual dependence. The boundary conditions represented in the graphical depiction are by and large irrelevant. Other modeling parameters¹, such as coefficients of

¹The "modeling parameters" vocabulary here is a loosening of the notion of boundary conditions on a modeled system, and it is meant to encompass both the boundary conditions on partial differential equation models, as well as other parametric features of models that fall into neither fully nomic nor fully incidental categories. Many thanks to Mark Wilson for suggesting this more expansive vocabulary, which will be

friction that would overcome the force of gravity on the block, are neglected in both the graphical depiction and the causal explanation of the block's sliding, except as *ceteris* paribus conditions. The solution is: The block slides down the ramp because it is experiencing gravitational force. If one is partial to unificationist lines of reasoning, one might instead say, The block slides down the ramp because all things experience gravitational force, and this block is one of them.

III. The problem of getting the block to arrive at the bottom of the ramp at a speed of 10 m/s: This problem requires using theories and models to change the values of various input variables, or modeling parameters, in order to produce the desired outcome. Clearly, given the solution in (1) above, leaving the system alone will not ever produce the desired result. However, by setting V_f to 10 in the system of equations in (1) above, it becomes apparent that various changes to V_i , a and d, or some combination thereof will lead to the desired outcome. Some solutions include: Lengthening the ramp to 7.2 m, increasing the weight of the block to 3.6 kg, or beginning with the block moving at 8.5 m/s.

In each problem case, theories and models are used to obtain information about relations between different components of the inclined-plane system. This should come as no surprise, but perhaps as some comfort. One of the most elegant consequences of interventionism has been to reveal that whatever else theories and models might be, at heart they are simply ways of organizing information about relations between parts of systems.

However, what is done with that information once it is in hand differs from case to case and is determined by the problem setting and the activity that is meant to be accomplished by the use. Philosophers of scientific explanation working in the pragmatic tradition (cf. e.g. van Fraasen 1980; Achinstein 1985; Batterman 2001) have used the framework of whyquestions or explanation-requests as determinants of the context for application of a theory to a problem. Tuning problems can be defined via similar interrogative contexts; however, the questions asked in tuning settings are not "Why did [x] happen?" or "Why do all [x]s behave similarly?" but rather "How can [x] be made?" or "What needs to happen to make [x]?" Tuning problems are, consequently, a kind of *design problem*.² In the same way that

developed in earnest in future work.

²There is an extensive literature on design problems in contemporary psychology, information science and


Figure IV.3: An inclined-plane problem. [Author illustration]

deductive-nomological analysis, unification, and intervention are ways of solving prediction and explanation problems, tuning is a way of solving design problems.

Returning to the inclined plane, problem (1) is a deductive-nomological use of theories and models to predict. Problem (2) is a causal (or deductive, or unifying) use of theories and models to explain. Problem (3) is a tuning use of theories and models to design. It differs from problem (1) because the fixed parameter is an output, rather than an input. It differs from problem (2) in that causal/explanatory relations are only interesting in problem (3) insofar as they are a means of obtaining information about how to produce the desired target, rather than as an end in themselves. And it differs from both problem (1) and problem (2) because it is the modeling parameters rather than the equations of motion that draw the lion's share of the attention in the problem's solution.

These differences highlight three central features of tuning problems: they are constrained by output rather than input, causal (and nomological) relations are treated instrumentally, and solutions emphasize the role of boundary conditions and other modeling parameters over that of laws. The former feature defines tuning problems as a kind of design problem, and the other two suggest particular attitudes about inter-theoretic relations that will be required for the problem's solution. The latter feature is interesting in light of recent debates about the role of initial and boundary conditions in modeling problems in classical physics (cf. e.g. Batterman 2012; Wilson 2012; Butterfield 2011; Butterfield and Bouatta 2011). In these debates, the thermodynamic limit is a boundary condition that plays a central role in the modeling of phase transitions, and Batterman has argued, *contra* Butterfield and Norton, that the use of a particular model (namely, infinite limits in the renormalization group) is only justified by careful attention to how boundary conditions are modeled, rather than by attention to the dynamical laws that describe the system.

In tuning problems, emphasizing modeling parameters also helps to justify the use of a given set of models. However, it does so not by attending to the mechanics of initial

computer science. This literature, influenced particularly by the work of Simon and Newell (1972), studies human cognition in the solving of design problems. It is closely related to the present discussion of the use of theories and models to solve tuning problems, but it emphasizes cognition rather than theories and models. A detailed discussion of the overlap between these subjects falls outside the scope of this thesis. Likewise, the literature on inverse problems in applied mathematics provides another fruitful area of overlap with the present discussion of design problems and will be pursued in future work.

and boundary conditions in developing solutions to the problem, but rather by using these conditions—as opposed to dynamical laws—to identify what models and theories will apply to the problem setting. Because modeling parameters, and not laws, are the sorts of things that can be manipulated to achieve a target output in a tuning problem, the information that needs to be obtained from models and theories in order to solve tuning problems is information about what sets of modeling parameters will lead to the target output. If that information is accompanied by a story about causal or nomological relations in the system, so much the better, but that kind of information is only useful insofar as it improves the user's ability to determine what modeling parameters will result in the desired state of affairs. Consequently, concerns about reduction, emergence, and the ontological accuracy and relative 'fundamentality' of models and theories are largely irrelevant to solving tuning problems.

Which is not to say that all inter-theoretic relations are irrelevant. In the next section I use an example from nanosynthesis to argue that multiple models and theories are usually required to solve tuning problems. In such cases, models and theories are used *iteratively*, *interpolatively*, and *interactively* to determine what inputs will lead to the desired output. This way of understanding the mechanics of theory and model use in tuning problems presents an alternative picture to traditional accounts of reduction and emergence in discussions of inter-theoretic relations.

C. MODELING SYNTHETIC CONSTRAINTS

In order to tune LSPR in metal nanoparticles, one must make metal nanoparticles. Recalling the discussion in Chapter I, these nanoparticles will have some size, shape, composition, and surface chemistry. The Mie solution in Chapter III demonstrated that there is a set of relationships between these tunable properties and the strength and input wavelength of an LSPR response. However, the Mie solution left out a number of relationships between size, shape, composition, surface chemistry, and environment, which constrain the possible syntheses a researcher may carry out. Like the capping agents discussed in previous chapters, there are patterns of affinities that constrain, e.g., shape in terms of composition, or size in terms of environment. These patterns are *synthetic constraints*.

A full spectrum of the synthetic constraints imposed on a given synthetic system is not, in practice, the product of a single model or theory. Instead, a variety of models or theories contain information about different constraints—that is, different relationships between size, shape, composition, surface chemistry, and environment—and these models or theories must *interact* to produce information about the tunable properties of a nanomaterial.

Some synthetic constraints, like those described in common miscibility charts, are not tied to the nanoscale features of the system. Other synthetic constraints, such as the need for capping agents for metal nanomaterials, are constraints that only appear in nanosynthesis, as opposed to macroscopic synthesis. Information about these synthetic constraints is gathered from models and theories that have been adapted to the nanoscale from higher-scale, macroscopic materials modeling or from lower-scale, molecular modeling, as illustrated in the following example of models of crystal structure applied to nanomaterials. What is important to note is that rather than developing an entirely bottom-up or top-down theory of nanomaterial structure and behavior, information from these models comes together to form a patchy, incomplete, but useful-for-synthesis map of synthetic constraints that can be used to make nanomaterials with a tuned LSPR response.

Crystal structure is the arrangement of atoms in a crystalline solid (or liquid), that is, a solid or liquid made up of regularly repeating unit cells of atomic geometry, as opposed to an amorphous or structured but non-repeating arrangement of atoms. Common crystalline solids include salts and gemstones, and most metals are considered polycrystalline solids, meaning they are made up of multiple crystalline patches, called crystal planes or grains, that are bonded together along asymmetric grain boundaries. (See Figure II.7.)

Metal nanomaterials are either monocrystalline or polycrystalline solids. In macroscopic metals, crystal structure does not significantly influence the shape of a material, because the individual atoms that make up the material are negligible relative to the size of the material itself. This is obvious, as metal objects appear in a dazzling variety of shapes and sizes in our everyday lives. However, individual atoms make up significant proportions of metal nanomaterials, and so crystal structure constrains the possible shapes available to many metal nanomaterials.³

There are a variety of models and theories of crystal structure, and these models and theories come from a variety of fields. In macroscopic materials modeling, crystal structure is taken into account when materials scientists model fracture planes in order to gather information about how a crystalline material will respond to various stresses like squeezing and cracking. These models ignore the structure of unit cells within the material and focus instead on the notion of boundaries between crystal planes in order to predict or explain how a material will respond to stress. In crystal lattice theory, on the other hand, crystal structure is modeled by group theory and mathematical lattices, which provide information about symmetries in the crystal's atomic unit cell. Each chemical element has a unique crystal structure, which is determined by symmetries that arise from the electronic structure of the atoms of that element. These symmetries determine whether or not two elements can form a crystalline compound, and if so, what lattice structure that compound will take. Crystal lattice theory relies on the notion of an infinitely extended symmetric lattice with no interruption in the crystal plane in order to predict and explain what structural arrangements of atoms are likely for crystals of different compositions.

So fracture plane modeling and crystal lattice theory make conflicting assumptions about the structure of a unit cell and about the existence of asymmetrical boundaries in a crystalline material. Consequently, one model is not reducible to the other, nor are the two models able to enter into a logically consistent taxonomy. However, it is still possible to model crystal structure in a polycrystalline material by determining its unit crystal structure from crystal lattice theory and orienting crystal in different directions around fracture planes (See Figure II.7, left). At the nanoscale, interacting models of this sort can be used to relate information about the overall shape of a material to information about its composition. For instance, applying interacting models of crystal structure to the nanoscale can identify whether a particular element or compound is more likely to form cubes or triangular prisms (Tao et al. 2008).

³American toy companies sell spherical neodymium magnets as toys under the commercial names "Zen Magnets" and "Buckyballs." Playing with these magnets is one of the best ways of understanding the increasing constraint of crystal structure on shape as a material decreases in size. A large bunch of the magnets can be deformed nearly continuously into many shapes, but smaller numbers of magnets only form stable structures in particular crystalline arrangements.

At the nanoscale, then, composition and shape are related to one another through crystal structure. The Mie solution does not contain information about those relationships. Instead, multiscale, interacting, mutually inconsistent models relate information about the arrangement of atoms in a nanomaterial to information about the overall shape and size of that material. This information can be used to determine whether it is possible, for a given set of input compositions, sizes, and shapes that will produce the target LSPR response, whether that combination of inputs is synthetically possible.

This picture of multiple models that represent multiple kinds of constraints on a system of variables can be used to generalize the picture that I have in mind when I speak of *model interaction*. As a further example, consider a system of algebraic equations, e.g.

 $w - 5x^{2} + 3y + 7z = 18$ $7w^{3} + 16x - y + 12z = 30$ xw + 2x + 5y - 8z = 27

Different equations represent different relationships between inputs (i.e. variables), or, equivalently, different models of the system of variables. In order to solve the system of equations for w, x, y, and z, the models must interact. Different methods of combining information from each of the equations and solving for w, x, y, and z—whether to use addition or substitution, or simply eyeball it, or write a computer algorithm to test randomly generated input values—are different kinds of model interaction. In many cases, such as the one above, there is more than one set of inputs that solves a tuning problem via model interaction.

Synthetic constraints such as interdependencies between composition and shape make up one set of interacting models used to solve the problem of tuning AMNPs for a particular LSPR response. In order to solve the problem in full, a variety of other kinds of constraints between size, shape, composition, and surface chemistry need to be accommodated. Importantly, no one kind of constraint is the most 'fundamental' or first kind of input that is assigned a definite value; that is, one is not always solving for shape and size in terms of composition or for surface chemistry in terms of shape. The context of particular tuning problems will determine which constraints take priority. Or, to put the point in terms of the algebraic analogy, the context of a particular tuning problem is what defines dependent and independent variables among a system of interdependent inputs.

This context is often external to the mechanics of the synthesis itself. In the problem of tuning LSPR for thermal ablation therapy in cancer patients, for instance, it is important to take biological compatibility into account. Biological compatibility constrains this particular tuning problem in two distinct ways. First, in order to activate LSPR in nanomaterials inside the body, the incident light that activates LSPR must be able to reach the nanomaterials, which means it must be able to pass through the skin. Incident wavelengths in the near-infrared range, around 1000 nm, can penetrate the skin (where visible light cannot), so researchers aim to synthesize nanomaterials with strong LSPR responses around 1000 nm.

Second, nanomaterials that are synthesized in an environment that is highly toxic to the body are, by definition, not biologically compatible, nor are nanomaterials made of poisonous metals such as elemental cadmium. Gold is a particularly biocompatible metal in small doses, and so most nanomaterials synthesized for thermal ablation are either composed entirely of gold or coated in gold (Loo et al. 2005). This means that composition takes priority as a determiner of other synthetic constraints in the problem of tuning LSPR for thermal ablation therapy. By way of contrast, other sorts of tuning problems have focused on how to control shape in a given composition (Tao et al. 2008), or how to grow nanomaterials in a particular chemical environment (Brust et al. 1994).

D. ITERATIVE INTERPOLATION AND INSTRUMENTAL ATTITUDES

One final note will wrap up this discussion of the mechanics of tuning. Rarely does a firstpass synthesis result in the production of the target output. Instead, it is usually the case that a first synthesis makes something that resembles the target in recognizable ways—an attempt to make gold nano-prisms results in gold nano-spheres, for instance. Synthesis is not a one-shot endeavor. Instead, like tuning a guitar, solving tuning problems requires repeatedly adjusting inputs over a series of attempts. I find it useful to call the process of revising synthetic conditions over a number of trials *iterative interpolation*.

During iterative interpolation, a researcher changes the values of various input parameters over a series of trials. She makes choices about which parameters to change and how to change them, and these choices are influenced by two main sorts of factors: previous research and experience, on the one hand, and models and theories, on the other. These factors interact cooperatively to suggest revisions to initial synthetic conditions. Over the course of a series of trials, the goal is to bring the actual output closer to the target output. The "interpolation" of iterative interpolation is the ϵ - δ -like approach toward the synthetic target from an initial trial.

For instance, it is possible to design a solution to the LSPR-for-thermal-ablation tuning problem in which one synthesizes gold nano-cubes with edges about 100 nm long in a biocompatible environment. If the first attempt to carry out this synthesis results in spheres instead of cubes, the experimenter will revise the synthetic procedure. One way of revising the procedure is to change the capping agent that is being used to stabilize the particles. As discussed in Chapter III.E.2 some capping agents, such as sodium citrate, tend to produce spherical particles, whereas others, such as poly-(vinyl pyrrolidone) produce particles in a variety of non-spherical geometries. These non-sphere-producing capping agents are known as *selectively binding* capping agents.

Synthesis is a process of particulars, and researchers cannot make materials simply by thinking, "I should use a selectively-binding capping agent." They must choose a specific chemical (and solvent, and concentration, etc.). And while they may not spend much time consciously justifying their choices to themselves, the choice is rarely arbitrary. There are at least three ways a researcher could justify her choice of a particular selectively binding capping agent: Some researchers have worked extensively with selectively-binding capping agents and may simply select one from a mental catalogue. Others may read of successful syntheses with similar parameters in a journal and attempt to replicate the procedure. Still others may employ some sort of theoretical justification.

When a researcher employs theoretical justification during iterative interpolation, it is another case of model interaction. But even in the absence of such reasoning, the process of iterative interpolation underscores the idea that information about relations between inputs and outputs—i.e., the information gained by using models and theories—is interesting for its instrumental value in guiding perturbations of synthetic conditions, rather than for its own sake. This instrumental attitude toward models and theories further distinguishes tuning contexts from some descriptive contexts. In most descriptive contexts, one looks into the black box out of feline curiosity, in order to see what is inside: to uncover some system of relations that around which an explanatory or predictive narrative can be built. In synthetic contexts, however, the black box is opened in order to *use* what is inside to produce the desired output.

E. CONCLUSIONS

This chapter has introduced *tuning* as a use of theories and models that is distinct from description, explanation or prediction. In Section IV.B, I framed tuning as a process of varying a set of input conditions to obtain a fixed, target outcome, and I contrasted this process with prediction, description, and explanation, which generally seek to uncover fixed input conditions which, once revealed, provide additional understanding of an outcome or a set of possible outcomes. Section IV.C described how theories and models are used to solve tuning problems in nanosynthesis, such as making nanoparticles whose LSPR response may be used for thermal ablative therapy in cancer patients. I argued that different models and theories, such as the Mie solution to Maxwell's equations and crystal lattice theory, refer to the same or similar tunable features of nanomaterials, such as size, shape, composition, and surface chemistry. These tunable features make up the modeling parameters that are varied, manipulated, or perturbed in order to generate the target output LSPR response. Different models contain information about distinct kinds of relationships or constraints between these parameters, which prevents one model from reducing to another. Sometimes, these models are even mutually inconsistent. This inconsistency, while potentially troubling for some theories of explanation or prediction, is not offensive to tuning, because combining information from the different models can nonetheless generate a set of mutually compatible inputs that lead to the target output.

The cases in this chapter have shown that in solving tuning problems, models and theories *interact*, rather than reducing to one another or emerging from one another. The need for non-reductive accounts of inter-theory and inter-model relations has been discussed elsewhere in the philosophical literature. For instance, Mitchell (2009) has used her theory of "integrative pluralism" to argue against the desirability of a "complete" representation of a scientific system. Recalling the ramp example from Section IV.B, a model or theory that describes the dynamics of the block sliding down the ramp will not contain information about the material structure of the block or the ramp, nor about the complex electrostatic interactions between the surfaces of the two materials.

Integrative pluralism, Wilson's physics avoidance, and Winsberg's handshaking are all examples of non-reductive model interaction in descriptive, input-oriented contexts. One feature of all these accounts is their emphasis on the need for local philosophies of science that contextualize conditions for successful representation, explanation, etc. to the goals and methods of individual sciences. Rather than a one-size-fits-all picture of how science works, these accounts and others (cf. e.g. Batterman 2001; Chang 2004; Wimsatt 2007; Love 2012) paint pictures of scientific reasoning in individual sciences such as classical mechanics, thermodynamics, evolutionary-developmental biology, and population biology. Sometimes their conclusions generalize to other sciences, other times the conclusions remain local. The account of tuning presented here further develops this bottom-up picture of scientific reasoning by describing how theories and models contribute to reasoning in synthetic contexts.

While this dissertation has been locally anchored in nanosynthesis, the tuning account can be used to understand why there are relatively few extant discussions of scientific reasoning in the applied sciences, and how that oversight might be rectified. I began this chapter by demonstrating that tuning is an activity that occurs frequently across the sciences, and one which philosophers of science have thus far overlooked. Tuning is the use of models and theories to design a system of inputs that will lead to a target output. In nanosynthesis, the output is some desired structure or behavior of a nanoscale material. The example in the previous section considered the behavior of strong LSPR response to a given incident wavelength, but there are indefinitely many other target structures and behaviors to which a nanosynthesis could be tuned. In any case, a target outcome is defined by a problem context, such as the context of designing nanomaterials for thermal ablation therapy. Problem contexts can be either internal to scientific practice or external to it. Designing materials for thermal ablation therapy is an example of a problem context external to nanosynthesis (in this case, from medicine) that is being solved via nanosynthesis. An internally defined problem setting, on the other hand, is a question about how nanomaterials themselves work—whether platinum and palladium can form an alloy in a nanoscale material, or whether single-layer graphene is more conductive than thin multi-layer graphene. Figuring out what synthetic conditions are optimal for nanoscale alloy formation or production of highly conductive graphene is still solving a tuning problem.

I mention this distinction in kinds of problem context because the first, externally-defined problem setting is clearly a case of applied science. It is less clear whether the second, internally-defined problem setting is a case of applied science. And this haziness indicates that it is worth reflecting on the relationship between tuning problems and applied science. I introduced tuning as an alternative to explanation, prediction, and description, and I argued that it occurs in scientific settings where something is made or produced or synthesized, rather than merely described or explained. This difference between describing and making is one way of distinguishing theoretical from applied science, and it is a fairly natural one: it easily accounts for the tendency to lump together applied science and technology, and it rather naturally captures the tendency of applied science to yield some sort of product, whether it be medicine, building materials, weaponry or solutions to the energy crisis.

It is not a perfect distinction, for there are applied sciences, such as meteorology, whose central "product" is a kind of prediction. Moreover, there are instances of making, especially in the synthetic sciences, where it is not clear whether the product will have a use or function once it is made. This describes the situation for the alloys mentioned above, and, I suspect, for most products produced in the context of internally-defined problem settings. In such cases, however, while there may not be a *particular* use in mind at the time of design, it is still more likely than not that if the product is successfully made, its structures and behaviors will suggest a variety of potential uses. And if this hypothetical sounds unbelievable, consider the battery, the steam engine, the cathode ray tube, or bronze.

Solving tuning problems in internally-defined problem contexts may or may not be a

case of applied science, and I suspect that it is ultimately a matter of taste. But the difference between productive and descriptive scientific activity, which tuning has brought into sharper relief, is a useful first pass at capturing the difference between applied and theoretical sciences. This difference also helps to explain why philosophers of science, whose attention has been on the descriptive activities of prediction and explanation, have so often overlooked epistemological activity in the applied sciences. It also suggests that Woodward's interventionist theory of explanation, which focuses on the tuning-friendly concepts of control and manipulability, may prove particularly effective in applied contexts. So tuning may just serve as a useful inroad toward an epistemologically-oriented philosophy of applied science.

V. CONCLUSIONS

A. OVERVIEW

This dissertation has surveyed the structures and uses of theories, models, and heuristics in synthetic nanoscience in order to show (1) that some concepts and some phenomena are *scale-dependent*, (2) that scale-dependence influences both the structure of models and theories and the uses to which models and theories may be put, and (3) that traditional accounts of theory structure have no means of accommodating these roles of scale-dependence. I have developed an alternative account of theory and model structure that can and does accommodate scale-dependence. This account (a) centralizes the role of boundary conditions in some scientific explanations, (b) uses scale-dependence to articulate relationships between models other than those of reduction or emergence—relationships that I call varieties of *model interaction*—and (c) develops a contextual approach to model and theory use in which models may interact differently depending upon the purposes—e.g. prediction, explanation, synthesis—to which they are put. In particular, I have highlighted synthesis as a distinct use of theories and models, which influences the selection, refinement, and application of models appropriate to a given problem-solving context.

Synthesis, the production of a new stuff out of other stuffs, is the signature activity of chemical practice. Unlike physics, which aims principally at the identification, prediction and explanation of stable systems, or biology, which aims to understand the how and why of life, chemistry is very often aimed at making something. This difference in motivation between chemistry and the sciences more commonly subjected to philosophical analysis has been pointed out before, e.g. by the chemist and Nobel laureate Roald Hoffmann: I have not made a rigorous examination of the education of philosophers of science. But my anecdotal feeling is that, for those who entered the profession directly, an exposure to mathematical logic is more likely than to geology or chemistry. And, for many of the philosophers of science who came to their field after an initial scientific career, their scientific expertise was likely to be in the first instance physics, after that biology, and rarely chemistry. I will argue that this matters, for chemistry is different. (Hoffmann 2007, p. 322)

However, no one has taken the time to point out how the influence of this motivation can shape the ways in which chemists construct and deploy conceptual models of the systems that interest them—in short, no one has taken the time to describe how chemists solve problems. This dissertation has been an attempt to begin rectifying that oversight.

Chemical theory consists of a wide variety of semi-empirical models that represent particular classes of natural systems. Sometimes two models contradict one another; other times they complement one another. In nearly every case, these models do not pretend toward applicability over the class of all chemical systems nor even (with the notable exception of the periodic table of the elements) toward dominance as a tool for conceptualizing a set of relations between two or more parts of the world that are subject to chemical study. Instead, these models are sewn together to produce a patchwork view of many concepts central to the practice of chemistry. What chemists mean when they refer to concepts like *chemical bond* and *reaction mechanism* can, and often does, change depending on which models they choose to anchor their references and inferences. And what they mean by various chemical concepts influences their experimental and theoretical practices.

This dissertation has concerned how chemists and nanoscientists choose which models will anchor their use of a concept in a given research setting, and how that choice shapes the structure and activity of chemical research. In particular, I have been interested in research settings in which a scientist is aiming to synthesize a substance or set of substances. At least in these settings, I showed, there is rarely if ever one overarching model that can provide a chemist all the information that she will need both to successfully perform the synthesis and to describe what she has created. Rather, different models dominate different stages of her reasoning process, and she needs to rely on information gleaned from one model as an input in another model—like the synthetic process itself, her model use comes in distinct stages, each stage building on the last. The kinds of inferences she draws in these cases are rarely strictly inductive or deductive, in part because reasoning across models requires the use of analogy or other forms of abduction to map the properties of a concept predicted by one model onto the properties of that concept predicted by another.

One of the worries that generated this dissertation was how to characterize this form of abductive reasoning. I have termed various aspects of synthetic reasoning processes *model interaction*, *iterative interpolation*, *tuning*, and *synthetic conceptual interpolation*. The earlier chapters of this thesis have served as existence proofs for each of these aspects of synthetic reasoning processes. This chapter articulates in more detail how these various components fit together. Again, I emphasize the role of scale as a constraint on these reasoning processes. Once I have developed my account of synthetic reasoning here in more detail, I close by considering how this account may apply outside the chemical sciences to synthetic problems in physics, biology, engineering, and even outside the natural sciences.

B. CLASSIFICATION, SCALE-DEPENDENCE, AND MODEL INTERACTION

In Chapter I I showed that the problem of classifying nanomaterials is distinct from the problem of classifying chemical substances. I attributed this distinction to differences in the relevant structural features of nanomaterials relative to chemical substances. While the properties of interest in chemical settings can be explained by reference to molecular structure alone¹, the properties of interest in nanomaterials require additional structural information for their explanation, namely information about *dimension*, *shape*, and *surface chemistry*. This difference between classification in nanoscience and classification in chemistry marked the first of many instances of *scale-dependent material behavior* considered in this thesis. Whether a material exhibits scale-dependent behaviors—such as changes in conductivity and catalytic activity, or the appearance of quantum confinement or LSPR—is a *de facto* field test of whether or not it is a nanomaterial. If so, determining what kind of nanomaterial

¹For a fascinating historical discussion of how structure became the central sortal in organic chemistry, see Hepler-Smith's recent article on the 1892 Geneva Nomenclature Congress (Hepler-Smith 2015).

it is requires structural analyses of the sorts I suggested.

One purpose of this discussion was to articulate philosophy of science's interests in the natural kinds debate in philosophy more broadly. I argued that that debate should be understood as composed of three inter-related but distinct questions about kinds, which I called the metaphysical, reference, and methodological problems of kinds, and the discussion of chemical and nanoscale kinds showed how philosophy of science may gain the most from responding to the latter. In particular, it is methodological analysis of the changing importance of various structural features in synthesizing and making nanomaterials that reveals why classifying them solely as chemical substances is inadequate.

The four-part classification scheme proposed in Chapter I serves as an archetype to motivate ensuing the chapters' discussions of *model interaction*. The classification scheme I laid out will, in its ideal state, produce robust individuation—nanomaterials can be uniquely specified by sufficiently narrow specification of composition, dimension, shape, and surface chemistry—but there is no clear hierarchy of fundamentality or sub-taxonomy among the four characteristics. Instead, specifying a material's location along one dimension of analysis constrains but does not determine its location along the other dimensions, e.g., one cannot have a 2D nanocube, or given a design for a gold particle, one might seek a sulphur-containing surfactant, in order to take advantage of the uncommonly strong bond between gold and sulphur.

None of the four characteristics is fundamental relative to the others, and in practice, nanoscientists have developed research programs around each individual characteristic as well as around combinations thereof. Each characteristic requires a slightly different structural understanding of the nanomaterial under consideration, which in turn requires different models of the material. These models, like their targets, collectively inter-constrain one another and occasionally conflict with each other. Which model serves as primary is indexed to a given problem setting, during which a path is drawn through the system of models that specifies which one(s) are used at which points during a particular synthesis or explanation. These paths are non-unique and are limited by the physical behavior of the systems they describe, as well as by computational and representational limitations and the parameters of the problem on the table. This inter-dependence is what I mean by *model interaction*,

which I described at length in Chapters III and IV. There I discussed varieties of model interaction and argued that extant accounts of inter-theory relations often fail to capture and characterize these varieties.

C. IMPLEMENTING MODEL INTERACTION

In Chapter III, I cited Winsberg's (2006; 2010) "handshaking" models as an account of inter-theory relations that aims at the same sort of story I have presented, but which falls somewhat flat on the details of how models interact in the case he describes. In that case, he considers the parallel multiscale simulation of crack propagation in a silicon matrix with nanoscale features. The component models in the simulation range across scales, and he describes the combination of these models as an averaging technique rather than series of reductions of higher-scale models to lower-scale ones. This is correct, so far as it goes, but it misses the details of how averaging, or "handshaking," develops. Each of the three component models, I argued, are accompanied by differing and sometimes conflicting background assumptions. Following the mathematical details of the interfacial "handshake" models is the only reliable way to assure ourselves that in that store of varied assumptions, no lion lays in sheep's clothing, as it were—the mathematical details, such as the continuous deformability of the finite-element lattice sites to align with atomic sites in the molecular-dynamics model, legitimize the process of combining models and ensure that we are not sneaking un-physical or un-chemical assumptions into the interface.

In this dissertation I have concentrated on instances of model interaction that arise in multiscale modeling situations in physical, chemical, and nanoscientific settings. In these instances, from the sheared sides of the De Havilland Comet and the acoustics of a violin to the synthesis and LSPR of gold nanocubes, as well as in the crack-propagation simulation, I have found a common problem, a common feature, and a common strategy toward solution. The problem is how to model phenomena that occur at a variety of length scales and which influence smaller- and larger-scale material behaviors. In the Comet, for instance, the buildup of metal fatigue is a small-scale phenomenon that led to large-scale shears in the metal siding of the jetliner, and in the nanocube, small-scale surface structure and large(r)-scale particle shape both contribute to the output plasmon resonance. In none of these cases is the largerscale or smaller-scale model the "right" one, and information from both (or all) models is needed to solve the sorts of problems that the models are being used to solve.

The common feature is the tendency of the model of the larger-scale system to parameterize out the details of the smaller-scale material behavior; that is, to treat as featureless exactly the feature that the smaller-scale system aims to model. The harmonic model of the violin holds fixed the endpoint of the vibrating string, where the model of sound transmission through the violin system models the vibration of the string against the bridge, into the sound post, and so on. This is the feature of modeling scenarios that Wilson (2008) has called "physics avoidance." This feature signifies a particular kind of modeling situation, in the same way that easy separation of features of a system into component entities and activities operating at different stages might signify a modeling situation appropriate to the New Mechanists' (Machamer et al. 2000) account. This parameterization effect is a result of the multiscale nature of the systems that fall within this problem domain, and in practice it brings with it conflicting descriptions of a particular aspect of the system. One model will treat this aspect as continuous, and the other will discretize it—this is what happens between the continuum and atomistic models in the Comet and the crack propagation models, and the pull between treating a surface as continuous and discretized is the nominal "tension" of Chapter II.

The solution strategy, in all these cases, is to give a mathematical or conceptual story that rationalizes (a) the use of multiple models, and (b) the movement between smallerscale and larger-scale models. The first part of this story is usually told in terms of the separation of scales at which the phenomena of interest occur. For instance, consider the crack-propagation system. Why is it legitimate to use both finite-element and moleculardynamics models on the same system, and how do Broughton et al. get away with using both, despite the conflicting (continuous vs. discretized) assumptions the component models make about the material makeup of the system? Well, for most of the system, i.e. the part that is *not* cracking, the material being modeled is governed by relatively macroscopic behaviors that can be "seen" only by continuum models like the finite-elements model. If one were to use molecular dynamics to describe this part of the system, not only would the simulation be computationally intractable, but it would in a very real sense get the dynamics wrong.²

The second part of this story requires attention to the mathematical details of individual models. In the crack propagation case, two distinct techniques are used. To combine the finite-element (FE) and molecular-dynamics (MD) models, Broughton et al. make the mesh points of the finite-element model line up with the lattice spacing in the molecular-dynamics model near the interface between the regions modeled by each component. To combine the MD and tight-binding (TB) models, they instead develop fictional entities whose energy they evaluate in each model and which they then average. The first technique can be thought of as *embedding* the MD model within the FE framework, where the second relies on *hybridizing* the two models at the interface. Additional ways of combining mathematics across scales are, of course, possible—Batterman, for instance, has written extensively about *asymptotic approach* techniques (cf. e.g. Batterman 2001). But what should not be done—and indeed what cannot be done in my account—and what Winsberg does³, is *ignore* the mathematical details in order to adopt a "lumpier" stance.

D. SYNTHETIC USES OF MODELS AND THEORIES: TUNING, DESIGN, AND ITERATIVE INTERPOLATION

All of the examples considered in any detail in this dissertation are cases of model interaction. Many, but importantly not all, are examples of synthetic modeling contexts, that is, problem-solving situations where one of the modelers' goals is to make something. The crack propagation simulation is not such a model, but the gold nanocube models are. I argued in Chapter IV that when synthesis is a goal of an instance of model use, the use to which

²Note that not all dynamical systems exhibit scale separation, such as in the case of turbulent fluid flow. When there is not a scale-separation story to be told, other means of rationalizing the use of multiple models must be laid out.

³I do not intend to single out Winsberg here as the sole perpetrator of this approach. Others, such as Bokulich in her "model explanation" account (Bokulich 2012), commit the same error. In some ways, Bokulich's approach is an even more extreme version of the problem, as on her account, the legitimacy of applying a model to a given explanation rests entirely in a "justificatory step," whose success and failure conditions are under-specified.

the model is put is as a means for *tuning* the input conditions in order to achieve a target output.

Tuning differs from explanation and prediction in a few ways. First, the kinds of information that are needed from a model or theory in order to generate a successful instance of tuning can be different than the kinds needed to explain or predict. To explain why a particular bridge collapsed in the way that it did at the time that it did requires an account of weak points in the shape of the bridge, or of the wear on the materials in the bridge over time, or perhaps of an environmental influence—climate change, a tornado—that was unaccounted-for by the engineers. To make a bridge that will not collapse, all of these factors and others (e.g cost) must be co-accommodated. Relations between design considerations like material choice and shape, such as that the cables on a suspension bridge are stronger when made of certain sorts of steel than others (and why cables not made of something silly like twine, marbles, or coffee) are at most secondary considerations in the explanation of a bridge's collapse, but they are central in the building.

A natural objection here is that tuning is simply explanation of a different phenomenon e.g., explaining how something is made, rather than explaining its physical or chemical behaviors. But this is not the case. In order to explain how a batch of gold nanocubes has come into existence, that is, how they are made, one must refer to models of the following sorts:

- Crystal structure models, which carry information about the arrangement of gold atoms within the cube,
- Surface structure and behavior models, which carry information about how the cube is constrained to the nanoscale,
- Chemical equations, which carry information about the reactions that take place on the way to the nucleation of the cubes, and
- Nucleation and growth models, which carry information about the kinetics of the cubes over time.

To make gold nanocubes, one needs some but not all of these models, as well as some additional models. In particular, crystal and surface structure models are not terribly relevant to the activity of synthesizing a batch of cubes, since the structural features of the material that are relevant to the selection of reaction conditions are embedded in models of nucleation and growth. Contrariwise, to design the synthesis, but not to explain it, one needs solubility heuristics that indicate what sorts of precursor solutions can mix with one another (e.g. whether a solution is water- or oil-based), as well as models that carry information about what sorts of surfactants are better at constraining growth in different kinds of metal solutions.

In addition to the use of distinct models, the ways in which a single model is used for explanation can differ from the way that it is used for tuning. Take LaMer's model of nucleation and growth (See Section III.C). LaMer's model contains information that relates the distribution of particle sizes in a batch of nucleated substances to the concentration of precursors and the amount of time a reaction spends in various stages of its reaction. Figure III.1 shows distinct stages of a reaction as metal ions are reduced out of a metal salt solution. In the nucleation stage, instances of reduction result in the formation of new particles. This stage is usually brought about by inducing a supersaturation state on the solution. In the growth stage, instances of reduction result in the addition of new atoms to existing nucleated particles. The two phases can overlap, but the basic principle of the model is that the shorter the nucleation period is, the more likely it is that the synthesis will produce particles of all the same size. In an explanatory context, this model might be used to explain why a certain batch of nanoparticles is more or less monodisperse on the basis of the amount of time the synthesis spent in a state of supersaturation. In a tuning context, the model is used to justify the selection of particular reaction conditions in order to minimize the amount of time the reaction spends in its nucleation state. Such a model could also be used to refine a synthesis over multiple trials, a process that in Chapter IV I called *iterative interpolation*.

As a use of theories and models, tuning is distinct from explanation or prediction. Synthetic problems, such as producing a chemical synthesis or building a bridge, can require different models and theories, as well as different uses of the same models and theories, than descriptive activities like explanation and prediction. The focus of this dissertation has been on synthesis in the strictest sense of the word, namely chemical synthesis of substances and nanomaterials. But "synthesis" signifies a broader class of scientific and everyday activities characterized by the bringing together of components to form a novel system, such as the design and construction of bridges, buildings, and scientific instruments. It may also turn out that some of the characteristic patterns of reasoning I have outlined here apply to activities that are synthetic in a more metaphorical sense, that is, creative activities like the making of art, conversation, and essays. While I have left these dimensions unexplored in this document, I suspect from my experience in drafting (and re-drafting, and re-drafting) this dissertation that synthetic reasoning mechanisms like tuning and iterative interpolation may apply to synthetic activities outside the sciences as well.

E. CODA: USE AND STRUCTURE

Many discussions of models and theories in philosophy of science aim to uncover the *logical* structure of these tools for scientific reasoning. Historically, this aim was made explicit in the titles of such accounts—Nagel's *The Structure of Science*, e.g., from which Kuhn's *The* Structure of Scientific Revolutions presumably drew its inspiration, or Popper's Logic of Scientific Discovery. The debate over semantic vs. syntactic approaches to characterizing scientific theories centralizes the role of theory structure in understanding scientific reasoning. More recent debates in philosophy of modeling tend to coalesce around how models represent their target systems, and these debates often rely on notions of mappings between structural features of models and the structure of their targets in order to explain how models manage to work.

The account I've presented here has had a different aim. Rather than beginning by scrutinizing the structures of theories and models, I began with the question, "What do theories and models *do* for scientific reasoning?" Uses of theories and models are many and varied, and they include activities such as explanation, prediction, and organization of the principles of nature, but they also include synthesis, technology, and design. I have aimed in the previous chapters to describe a particular kind of use to which theories and models are put, namely, the synthetic use.

Behind this approach lies the idea that theories and models are ultimately tools for

reasoning, and as with any tool, there is a relationship between structure and use. Further, this relationship shapes the tool in question from both ends—use guides structure; structure suggests use. This relationship excludes some tools from some applications: T-shirts, an effective human tool for blocking the sun from the torso, aren't much help to octopodes for any number of reasons, and one wants a hammer, not a bouquet of roses, to drive a nail into the wall. Among scientific theories, the same relation holds: quantum chromodynamics is not an effective tool for solving problems in population genetics, and the theory of synaptic plasticity can't fix climate change. On the other hand, the relationship between structure and use can be gainfully employed to tailor theories and models to better serve specific situations (use guides structure)—such as when LaMer's model was adapted from colloidal solution synthesis to nanosynthesis—and similar structural features in different contexts can suggest what theory or model may be apt (structure suggests use)—such as in the case of universality among critical phenomena.

One selects tools—indeed, one designs tools—based on what job needs to be done, not based on the structural contours of the object. That's what makes it a tool. Investigations of the structure of theories and models that do not begin by asking what the structure is *for* may nonetheless hit on correct descriptions of theory and model structure, but without a story of how that structure supports the activity of scientific reasoning, such investigations will fail to describe the kinds of tools that theories and models are.

By investigating the ways in which theories and models are used to synthesize substances and materials, I have revealed some characteristics of the kinds of scale-dependent, interactive, and context-sensitive tools that theories and models are.

BIBLIOGRAPHY

- Farid F Abraham, Jeremy Q Broughton, Noam Bernstein, and Effhimios Kaxiras. Spanning the Length Scales in Dynamic Simulation. Computers in Physics, 12(6):538–546, 1998.
- Peter Achinstein. The Nature of Explanation. Oxford University Press, 1985.
- M. Aizawa, A.M. Cooper, M. Malac, and J.M. Buriak. Silver Nano-Inukshuks on Germanium. Nano Letters, 5(5):815–819, 2005.
- Hristiyan A Aleksandrov, Francesc Viñes, Wiebke Ludwig, Swetlana Schauermann, and Konstantin M Neyman. Tuning the Surface Chemistry of Pd by Atomic C and H: A Microscopic Picture. *Chemistry–A European Journal*, 19(4):1335–1345, 2013.
- K. An and T. Hyeon. Synthesis and Biomedical Applications of Hollow Nanostructures. Nano Today, 4(4):359–373, 2009.
- James AJ Arpino, Edward J Hancock, James Anderson, Mauricio Barahona, Guy-Bart V Stan, Antonis Papachristodoulou, and Karen Polizzi. Tuning the Dials of Synthetic Biology. *Microbiology*, 2013.
- M.F. Ashby. Overview No. 80: On the Engineering Properties of Materials. Acta Metallurgica, 37(5):1273–1293, 1989.
- H.A. Atwater. The Promise of Plasmonics. Scientific American, 296(4):56–62, 2007.
- Ernie Balcerak. Climate Models Sensitive to Tuning of Cloud Parameters. Eos, Transactions American Geophysical Union, 94(18):172–172, 2013.
- Robert W. Batterman. The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence. Oxford University Press, Oxford, 2001.
- Robert W Batterman. Falling Cats, Parallel Parking, and Polarized Light. Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 34(4):527–557, 2003.
- Robert W. Batterman. The Tyranny of Scales. In Robert W. Batterman, editor, Oxford Handbook of Philosophy of Physics, pages 255–286. Oxford University Press, June 2012.

- Alisa Bokulich. Distinguishing Explanatory from Nonexplanatory Fictions. *Philosophy of Science*, 79(5):725–737, 2012.
- R. Boyd. Kinds, Complexity and Multiple Realization. *Philosophical Studies*, 95(1):67–98, 1999.
- Jeremy Q Broughton, Farid F Abraham, Noam Bernstein, and Effhimios Kaxiras. Concurrent Coupling of Length Scales: Methodology and Application. *Physical Review B*, 60(4): 2391, 1999.
- M. Brust, M. Walker, D. Bethell, D.J. Schiffrin, and R. Whyman. Synthesis of Thiol-Derivatised Gold Nanoparticles in a Two-Phase Liquid–Liquid System. *Journal of the Chemical Society, Chemical Communications*, (7):801–802, 1994.
- Julia R Bursten. Microstructure without Essentialism: A New Perspective on Chemical Classification. *Philosophy of Science*, 81(4):633–653, 2014.
- J. Butterfield. Less is Different: Emergence and Reduction Reconciled. *Foundations of Physics*, 41(6):1–71, 2011.
- J. Butterfield and N. Bouatta. Emergence and Reduction Combined in Phase Transitions. Arxiv preprint *arXiv:1104.1371*, 2011.
- J. Butterfield and N. Bouatta. Renormalization for Philosophers. Philosophy of Science Archive (forthcoming in Bigaj and Wüthrich, eds., *Metaphylics in Contemporary Physics*), 2014.
- Nancy Cartwright. How the Laws of Physics Lie. Cambridge University Press, 1983.
- Hasok Chang. Inventing Temperature: Measurement and Scientific Progress. Oxford University Press, 2004.
- Hasok Chang. The Hidden History of Phlogiston. HYLE, 16:47–79, 2010.
- Hasok Chang. Acidity: The Persistence of the Everyday in the Scientific. *Philosophy of* Science, 79(5):690–700, 2012a.
- Hasok Chang. Is Water H2O?: Evidence, Realism and Pluralism, volume 293 of Boston Studies in the History and Philosophy of Science. Springer, 2012b.
- J. Chen, F. Saeki, J. Benjamin, H. Cang, M.J. Cobb, Z.Y. Li, L. Au, H. Zhang, M.B. Kimmey, X. Li, et al. Gold Nanocages: Bioconjugation and their Potential Use as Optical Imaging Contrast Agents. *Nano Letters*, 5(3):473–477, 2005a.
- J. Chen, B. Wiley, Z.Y. Li, D. Campbell, F. Saeki, H. Cang, L. Au, J. Lee, X. Li, and Y. Xia. Gold Nanocages: Engineering their Structure for Biomedical Applications. *Advanced Materials*, 17(18):2255–2261, 2005b.

Fred I Dretske. Laws of Nature. Philosophy of Science, pages 248–268, 1977.

John Earman, Clark Glymour, and Sandra Mitchell. Ceteris Paribus Laws. Springer, 2003.

- Brian Ellis. What Science Aims to Do. In Paul M Churchland and Clifford Alan Hooker, editors, *Images of Science: Essays on Realism and Empiricism*, chapter 3, pages 48–74. University of Chicago Press, 1985.
- W. Goodwin. Mechanisms and Chemical Reactions. In Paul Needham Hendry, R.F. and Andrea Woody, editors, *Handbook of the Philosophy of Science*, Vol.6: Philosophy of Chemistry. Elsevier Amsterdam, 2012.
- C.P. Gräf, R. Birringer, and A. Michels. Synthesis and Magnetic Properties of Cobalt Nanocubes. *Physical Review B*, 73(21):212401, 2006.
- Ian Hacking. Working in a New World: The Taxonomic Solution. In Paul Horwich, editor, World Changes: Thomas Kuhn and the Nature of Science, pages 275–310. MIT Press, 1993.
- Amanda J Haes and Richard P Van Duyne. A Unified View of Propagating and Localized Surface Plasmon Resonance Biosensors. Analytical and Bioanalytical Chemistry, 379(7-8): 920–930, 2004.
- R. F. Hendry. Elements, Compounds and other Chemical Kinds. *Philosophy of Science*, 73 (5):864–875, December 2006.
- Evan Hepler-Smith. "Just as the Structural Formula Does": Names, Diagrams, and the Structure of Organic Chemistry at the 1892 Geneva Nomenclature Congress. *Ambix*, 62 (1):1–28, 2015.
- Roald Hoffmann. Marginalia: Unstable. American Scientist, pages 619–621, 1987.
- Roald Hoffmann. What Might Philosophy of Science Look Like if Chemists Built It? Synthese, 155(3):321–336, 2007.
- Jan-Markus Holm. Virtual Violin in the Digital Domain: Physical Modeling and Model-Based Sound Synthesis of Violin and its Interactive Application in Virtual Environment. University of Jyväskylä Press, 2004.
- G.L. Hornyak, J. Dutta, HF Tibbals, and AK Rao. Introduction to Nanoscience. CRC Press Boca Raton, FL, 2008.
- David L Hull. Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science. University of Chicago Press, 2010.
- Philip Kitcher. Science, Truth, and Democracy. Oxford University Press, 2001.

- K.J. Klabunde and R. Richards. *Nanoscale Materials in Chemistry*. Wiley Online Library, 2001.
- Marc Lange. Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature. Oxford University Press, 2009.
- Hongwei Liao, Colleen L Nehl, and Jason H Hafner. Biomedical Applications of Plasmon Resonant Metal Nanoparticles. *Nanomedicine*, 1(2):201–208, 2006.
- Ilya M Lifshitz and Vitaly V Slyozov. The Kinetics of Precipitation from Supersaturated Solid Solutions. *Journal of Physics and Chemistry of Solids*, 19(1):35–50, 1961.
- Y. Liu and A.R.H. Walker. Monodisperse Gold–Copper Bimetallic Nanocubes: Facile One-Step Synthesis with Controllable Size and Composition. Angewandte Chemie, 122(38): 6933–6937, 2010.
- Christopher Loo, Amanda Lowery, Naomi Halas, Jennifer West, and Rebekah Drezek. Immunotargeted Nanoshells for Integrated Cancer Imaging and Therapy. Nano Letters, 5(4): 709–711, 2005.
- Alan C Love. Theory is as Theory Does: Scientific Practice and Theory Structure in Biology. Biological Theory, 7(4):1–13, 2012.
- Peter Machamer, Lindley Darden, and Carl Craver. Thinking About Mechanisms. *Philosophy* of Science, 67(1):1–25, 2000.
- Ruth Garrett Millikan. On Clear and Confused Ideas: An Essay about Substance Concepts. Cambridge University Press, 2000.
- Sandra D. Mitchell. Unsimple Truths: Science, Complexity, and Policy. University Of Chicago Press, 2009. ISBN 0226532623.
- Margaret Morrison. Complex Systems and Renormalization Group Explanations. *Philosophy* of Science, 81(5):1144–1156, December 2014.
- E. Nagel. The Structure of Science: Problems in the Logic of Scientific Explanation. Hackett Publishing Company, 1979.
- Paul Needham. Microessentialism: What is the Argument? Noûs, 45(1):1–21, 2011. ISSN 1468-0068.
- Otto Neurath. Protokollsätze. Erkenntnis, 3(1):204–214, 1932.
- Allen Newell, Herbert Alexander Simon, et al. *Human Problem Solving*. Prentice-Hall Englewood Cliffs, NJ, 1972.
- R. Nussinov and C. Alemán. Nanobiology: From Physics and Engineering to Biology. *Phys*ical Biology, 3, 2006.

- C. Nützenadel, A. Züttel, D. Chartouni, G. Schmid, and L. Schlapbach. Critical size and surface effect of the hydrogen interaction of palladium clusters. *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics*, 8(2):245–250, 2000.
- Timothy O'Leary, Alex H Williams, Jonathan S Caplan, and Eve Marder. Correlations in Ion Channel Expression Emerge from Homeostatic Tuning Rules. Proceedings of the National Academy of Sciences, 110(28):E2645–E2654, 2013.
- Rob Phillips. Crystals, Defects and Microstructures: Modeling Across Scales. Cambridge University Press, 2001.
- Alexander Reutlinger. Why Is There Universal Macro-Behavior? Renormalization Group Explanation As Non-causal Explanation. *Philosophy of Science*, 81(5):1157–1170, 2014.
- E. Ringe, J.M. McMahon, K. Sohn, C. Cobley, Y. Xia, J. Huang, G.C. Schatz, L.D. Marks, and R.P. Van Duyne. Unraveling the Effects of Size, Composition, and Substrate on the Localized Surface Plasmon Resonance Frequencies of Gold and Silver Nanocubes: A Systematic Single-Particle Approach. *The Journal of Physical Chemistry C*, 114(29):12511– 12516, 2010.
- Ronald Sandler. Nanotechnology: The Social and Ethical Issues, volume 16 of Project on Emerging Nanotechnologies. Woodrow Wilson International Center for Scholars, Project on Emerging Technologies, 2009.
- Ronald Sandler and William D Kay. The National Nanotechnology Initiative and the Social Good. The Journal of Law, Medicine & Ethics, 34(4):675–681, 2006.
- T.K. Sau and C.J. Murphy. Seeded High Yield Synthesis of Short Au Nanorods in Aqueous Solution. *Langmuir*, 20(15):6414–6420, 2004.
- G. Schmid, M. Bäumle, M. Geerkens, I. Heim, C. Osemann, and T. Sawitowski. Current and Future Applications of Nanoclusters. *Chemical Society Reviews*, 28(3):179–185, 1999.
- Yumeng Shi, Wu Zhou, Ang-Yu Lu, Wenjing Fang, Yi-Hsien Lee, Allen Long Hsu, Soo Min Kim, Ki Kang Kim, Hui Ying Yang, Lain-Jong Li, et al. Van der Waals Epitaxy of MoS₂ Layers Using Graphene as Growth Templates. *Nano Letters*, 12(6):2784–2791, 2012.
- Ashley M Smith, Lauren E Marbella, Kathryn A Johnston, Michael J Hartmann, Scott E Crawford, Lisa M Kozycz, Dwight S Seferos, and Jill E Millstone. Quantitative Analysis of Thiolated Ligand Exchange on Gold Nanoparticles Monitored by 1H NMR Spectroscopy. Analytical Chemistry, 87(5):2771–2778, 2015.

Michael Strevens. Scientific Explanation. Encyclopedia of Philosophy, MacMillan, 2006.

Y. Sun and Y. Xia. Shape-Controlled Synthesis of Gold and Silver Nanoparticles. Science, 298(5601):2176, 2002.

- Kathyrn Tabb. Biomedical Progress and The Assumption of Diagnostic Discrimination in Psychiatry. Forthcoming in *Philosophy of Science*, 2015.
- A. Tao, P. Sinsermsuksakul, and P. Yang. Polyhedral Silver Nanocrystals with Distinct Scattering Signatures. Angewandte Chemie International Edition, 45(28):4597–4601, 2006.
- A.R. Tao, S. Habas, and P. Yang. Shape Control of Colloidal Metal Nanocrystals. *Small*, 4 (3):310–325, 2008.

Bastiaan van Fraasen. The Scientific Image. Oxford University Press, USA, 1980.

- J.L. West and N.J. Halas. Engineered Nanomaterials for Biophotonics Applications: Improving Sensing, Imaging, and Therapeutics. *Annual Review of Biomedical Engineering*, 5(1):285–292, 2003.
- B. Wiley, Y. Sun, B. Mayers, and Y. Xia. Shape-Controlled Synthesis of Metal Nanostructures: The Case of Silver. *Chemistry-A European Journal*, 11(2):454–463, 2005.
- Mark Wilson. Wandering Significance: An Essay on Conceptual Behaviour. Oxford University Press, USA, 2008.
- Mark Wilson. Mixed Level Explanation. *Philosophy of Science*, 77(5):1–12, December 2010.
- Mark Wilson. What is Classical Mechanics Anyway? In Robert Batterman, editor, Oxford Handbook of Philosophy of Physics, pages 43–106. Oxford University Press, 2012.
- William C. Wimsatt. Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality. Harvard University Press, 2007.
- Eric Winsberg. Handshaking Your Way to the Top: Simulation at the Nanoscale. *Simulation*, pages 139–151, 2006.
- Eric Winsberg. Science in the Age of Computer Simulation. University of Chicago Press, 2010.
- James Woodward. Making Things Happen: A Theory of Causal Explanation. Oxford University Press, 2003.
- Sanfeng Wu, Jason S Ross, Gui-Bin Liu, Grant Aivazian, Aaron Jones, Zaiyao Fei, Wenguang Zhu, Di Xiao, Wang Yao, David Cobden, et al. Electrical Tuning of Valley Magnetic Moment through Symmetry Control in Bilayer MoS₂. Nature Physics, 2013.
- Y. Yang, S. Matsubara, L. Xiong, T. Hayakawa, and M. Nogami. Solvothermal Synthesis of Multiple Shapes of Silver Nanoparticles and their SERS Properties. *Journal of Physical Chemistry C*, 111(26):9095–9104, 2007.
- Mustafa S Yavuz, Yiyun Cheng, Jingyi Chen, Claire M Cobley, Qiang Zhang, Matthew Rycenga, Jingwei Xie, Chulhong Kim, Kwang H Song, Andrea G Schwartz, Lihong V

Wang, and Younan Xia. Gold Nanocages Covered by Smart Polymers for Controlled Release with Near-Infrared Light. *Nature Materials*, 8(12):935–9, Dec 2009.

- Jiangtan Yuan, Lai-Peng Ma, Songfeng Pei, Jinhong Du, Yang Su, Wencai Ren, and Hui-Ming Cheng. Tuning the Electrical and Optical Properties of Graphene by Ozone Treatment for Patterning Monolithic Transparent Electrodes. ACS Nano, 7(5):4233–4241, 2013.
- Shuwen Zeng, Ken-Tye Yong, Indrajit Roy, Xuan-Quyen Dinh, Xia Yu, and Feng Luan. A Review on Functionalized Gold Nanoparticles for Biosensing Applications. *Plasmonics*, 6 (3):491–506, 2011.
- Wenliang Zhang, Yuan Lin, Jin Qian, Weiqiu Chen, and Huajian Gao. Tuning Molecular Adhesion via Material Anisotropy. *Advanced Functional Materials*, 2013.
- Y. Zhu, H. Qian, and R. Jin. Catalysis Opportunities of Atomically Precise Gold Nanoclusters. *Journal of Materials Chemistry*, 21(19):6793–6799, 2011.