

**DUHEM'S BALANCING ACT: QUASI-STATIC
REASONING IN PHYSICAL THEORY**

by

Meghan Danielle Dupree

B.A., Philosophy, University of California, Los Angeles, 2006

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

2014

UNIVERSITY OF PITTSBURGH
DEPARTMENT OF PHILOSOPHY

This dissertation was presented

by

Meghan Danielle Dupree

It was defended on

April 14, 2014

and approved by

Mark Wilson, Department of Philosophy

Robert Batterman, Department of Philosophy

James Woodward, Department of the History and Philosophy of Science

Giovanni Valente, Department of Philosophy.

Dissertation Advisors: Mark Wilson, Department of Philosophy,

Robert Batterman, Department of Philosophy

Copyright © by Meghan Danielle Dupree
2014

DUHEM'S BALANCING ACT: QUASI-STATIC REASONING IN PHYSICAL THEORY

Meghan Danielle Dupree, PhD

University of Pittsburgh, 2014

The celebrated philosopher-physicist Pierre Duhem appears to maintain virtually contradictory views. On the one hand, he claims that science does not aim to explain natural phenomena, where he assumes that an “explanation” strives to reveal the natural world underpinnings hiding “behind the veil” of observable phenomena. Despite these strong disavowals, he also insists that successful scientific theories should converge on “natural classifications” which allegedly provide “hints concerning the true affinities of things.” But won't such relationships also lie “behind the veil”? These warring inclinations have created significant exegetical confusions, leading his interpreters to classify him as an antirealist, a realist and everything else in between. Duhem is clearly trying to get across some important methodological lesson about science. But what is it?

The trick is to align his philosophy more closely with the forms of physics he endorses. On this basis, I argue that Duhem's disavowals of “explanation” actually represent arguments against a dynamic laws picture of science: the doctrine that science must seek laws that track material systems according to the basic patterns of D-N explanation. He argues that many of nature's most important hidden quantities (e.g., entropy and potential energy) were not discovered in a dynamical manner but were instead uncovered by stringing together relationships in the quasi-static manner employed in thermodynamics. Indeed, it is the deep relationships of the latter subject of the latter subject that supply paradigms of the “natural classifications” that Duhem seeks.

Once one follows through the details of his recommendations, employing concrete scientific examples, one realizes that Duhem's reflections on the scientific method greatly enlarge our appreciation of what the many varieties of "good science" can look like. This challenges many dogmatic presumptions about the scientific methodology that still prevail in contemporary philosophy of science.

TABLE OF CONTENTS

PREFACE	ix
1.0 A PUZZLE ABOUT NATURAL CLASSIFICATIONS AND EXPLANATIONS	1
1.1 Towards a Broader Conception of Physical Theory	13
2.0 ON EXPLANATION	18
2.1 Deductive-Nomological Accounts of Explanation	18
2.2 Causal Process Theories	21
2.3 An Explanation is a Perfect Theory	24
2.4 Dynamic Laws vs. Classifications	28
2.5 The Pendulum	29
2.5.1 The dynamic laws picture of the pendulum	31
2.5.2 Duhem's approach to the pendulum	32
3.0 LAGRANGIAN MECHANICS	34
3.1 Motivation for Lagrange's Mechanics: Constraints in the Point Mass Setting	34
3.1.1 Generalized Coordinates and Configuration Space	37
3.2 Lagrange's Statics and the Principal of Virtual Work	39
3.2.1 The Principle of Virtual Work	42
3.3 A Classification of Statics	44
3.4 Lagrange's Dynamics and D'Alembert's Principle	51
3.5 The Role of Lagrange	56
3.6 The New Mechanics: Energetics	57
3.6.1 Virtual Modifications	57

3.6.2	The Structure of Energetics	61
3.7	Conclusion	62
4.0	THE CARNOT CYCLE	64
4.1	A Brief Overview of Thermodynamics	64
4.2	The Discovery of Entropy	66
4.3	Thermodynamics as a Natural Classification	71
4.4	Statistical Mechanics	73
4.4.1	A Brief Sketch of Statistical Mechanics	74
4.4.2	A Reduction or an Illustration?	75
4.5	Conclusion and Summary	79
5.0	REFRAMING DUHEM	81
5.1	The Complexity of Coupling	81
5.2	Disguised Appeals to Variational Principles	85
5.2.1	Dissolving the Tension	95
5.2.1.1	Physical Theories do not Explain	96
5.2.1.2	Convergence on a Natural Classification	98
5.2.2	Summary	106
6.0	CONCLUSIONS	108
6.1	Consequences for Contemporary Theory	109
6.1.1	An Alternative Model of Scientific Theory	109
6.1.2	Theories of Causation	111
6.2	Mixed Level Theories	113
6.3	Physical Insight	115
6.3.1	An Objection Based on Equivalent Formulations	118
6.4	Final Thoughts	120
	BIBLIOGRAPHY	122

LIST OF FIGURES

1	An example of a simple pendulum.	30
2	An illustration of Duhem's perspective of the pendulum. In each state of the pendulum, the dotted line represents the potential energy and the solid line represents the kinetic energy. The total energy remains the same.	33
3	Pressure-volume diagram for a Carnot cycle.	68
4	An example of a pith ball electroscope.	99
5	Example of a classification from chemical theory.	105

PREFACE

This document would not be possible without the unceasing support and limitless genius of the philosophical community at the University of Pittsburgh.

In particular, I want to thank Mark Wilson for his patience, wisdom, creativity and foresight throughout this project. Mark has an uncanny ability to see value in even the foggiest ideas, mine included, and has taught me that there is always room to dig deeper. For this, I am forever indebted to him. I am similarly indebted to Robert Batterman, for his endless availability and willingness to hash out ideas, his unusual capacity for philosophical clarity, and his untiring devotion to his graduate students. Thanks to James Woodward and Giovanni Valente for their novel perspectives, both of which helped to broaden the scope of this project and ensure its technical accuracy.

I am also thankful for the conversations with and comments from, in no particular order, Sean Kelsey, Japa Pallikkathayil, Sandra Mitchell, James Rocha, Kristen Inglis, Katherine Brading, Andrew Chignell, Elizabeth Miller, Karen Rae Zwier, Katharina Nieswandt, Ulf Hlobil, Joshua Hancox, Bihui Li, and the University of Pittsburgh Dissertation Seminar 2012-2014. A very special thanks to Jim Bogen for reminding me, at least daily, to finish my thesis.

On a more personal note, I am grateful for BJ Woodworth and the congregation of the Open Door Presbyterian Church of Pittsburgh. Your prayers, meals, and emotional, spiritual, and financial support reflect your undying commitment to be the physical manifestation of Christ's love for his children. I am incredibly blessed to have been a part of this community. Special thanks to Ellen Huang, Joy Hickey, Melanie Montgomery, Kathryn Lindeman, Autumn Boyer, Amy Edwards, Kim Wachenfeld, Rachel Eanes, Anna Duchene, and Erika Aylwin for your friendship and love. Extra special thanks to Krista Aylwin for taking care

of Sojo the Bear when I was working late, Lee Lawson for all of her prayers, and Dr. Lee for his enduring mentorship.

A very special thanks to Christian Page, without whom I would have plausibly starved or, at the very least, had innumerable outstanding parking tickets. I am incredibly blessed by your support, love, willingness to do laundry, and incisive wit which always helps put life into perspective. I love you to the moon and back.

To my mom, thanks for believing in me far more than rationality permits. I honestly can't imagine a better mother than you have been. And to my father, thanks for your love and prayers.

Finally, I want to dedicate this dissertation to the best of all possible sidekicks, King Jennifer Modals Jhun. Jenn, you have made graduate school not only bearable but enjoyable. I cannot count the times or ways in which you have helped, supported, loved and comforted me. Your wisdom, devotion, and creative sense of humor extend far beyond what anyone could hope for in a friend. I can't wait for you to get a motorcycle.

1.0 A PUZZLE ABOUT NATURAL CLASSIFICATIONS AND EXPLANATIONS

Pierre Duhem, a physicist and chemist working at the turn of the twentieth century, wrote extensively on the nature of scientific theory. His scientific work focuses on thermodynamics and its relationship to other fields such as chemistry, fluid dynamics, and electricity and magnetism [Martin1991, 2]. Interpretations of Pierre Duhem’s philosophy of science are as numerous as his scientific works. He is an inspiration to positivists¹, instrumentalists², motivational realists³ and realists⁴ alike, while others claim his philosophy of science flows not from observation of scientific practice but instead from religious dogma⁵.

It is unsurprising that Duhem’s writings give way to a plethora of readings — his philosophical texts are full of poetic prose which fails to specify the scientific results that ground his model of science. He argues that science does not explain, but also claims that scientific theories converge on a natural classification.

... physical theory is not an explanation, but a simplified and orderly representation grouping laws according to a classification which grows more and more complete, more and more natural [Duhem1954, 55].

Problematically, the natural classifications that Duhem describes appear to meet the criteria for the explanations he disavows. On the one hand, Duhem insists science does not take us behind the veil, but in the next breath he exclaims it gives us pieces of the “true nature of things” — information which certainly is not part of the world of appearances. In

¹[Quine2010]

²[Goddu1990], [Ariew2008], [Cartwright1983]

³[Darling2003]

⁴[Needham2008], [Needham1991]

⁵[Brenner et al.2011], [Deltete2008]

this chapter I will explore this tension by unpacking Duhem’s claims about explanation and natural classifications.

To begin, Duhem claims that an explanation reveals the essential structure of nature.

To explain (explicate, *explicare*) is to strip reality of the appearances covering it like a veil, in order to see the bare reality itself [Duhem1954, 1].

Although Duhem attempts to define his notion of an explanation, he does so by invoking the vague phrase ‘bare reality itself’. To clarify Duhem’s definition of explanation, it helps to consider his assumptions about the nature of perception and the limits of the human intellect.

The human intellect does not have direct knowledge or immediate vision of the essence of external things. What we know directly of these things are the phenomena that arise from them and the sequence of these phenomena [Pierre Duhem1996, 31].

The observation of physical phenomena does not put us into relation with the reality hidden under the sensible appearances but enables us to apprehend the sensible appearances themselves in a particular and concrete form [Duhem1954, 7].

On Duhem’s picture, the world we observe is heavily mediated by our senses. There does exist some fundamental “reality” which causes our perceptions to have the qualities that they do, but we do not observe this reality. Instead, our intellect only receives the “appearances” presented by the senses. Based on this distinction between reality and appearances, to go “behind the veil,” as Duhem maintains a theory must do in order to be explanatory, is to describe the world as it is apart from our sensual apprehension of it — a true explanation must unveil the fundamental causes of our perceptions.

Duhem presents the acoustic theory of sound as one example of a scientific theory which poses as an explanation. Acoustic theory gives an account of the phenomenal aspects of sound according to changes in the vibratory patterns of waves. Concerning acoustic theory, Duhem writes

... these abstract notions — sound intensity, pitch, timbre, etc. — depict to our reason no more than the general characteristics of our sound perceptions; these notions get us to know sound as it is in relation to us, not as it is in sounding bodies. This reality whose external veil alone appears in our sensations is made known to us by theories of acoustics. The latter are to teach us that where our perceptions grasp only that appearance we call sound, there is in reality a very small and very rapid periodic motion; that intensity and pitch

are only external aspects of the amplitude and frequency of this motion; and that timbre is the apparent manifestation of the real structure of this motion . . . Acoustic theories are therefore explanations [Duhem1954, 8].

We experience auditory sense data in the form of different sounds. We can categorize this data according to the properties which distinguish one perception of sound from another. For example, we observe that some sounds are louder than others. Loudness is a measure of a sound’s intensity — how strongly a sound appears to a particular observer. But while predicates such as ‘loud’ effectively characterize our observations of sounds, they do not take us “beyond the veil” because such properties cannot be attributed to “the sound itself.” In contrast, acoustic theory models and describes the behavior of the waves that *cause* the sound. It reduces the concept of intensity to the amplitude of a wave, allowing us to refer to the true properties of physical bodies rather than to the phenomena we experience. Acoustic theory *explains* intensity by describing the essential nature of the physical bodies which produce auditory phenomena.

To summarize, I offer the following characterization of explanation a’la Duhem:

Duhemian Explanation (DE): A scientific theory provides an explanation if it presents us with a non-sensual (“behind the veil”) description of the entities which cause the phenomena we observe.

Duhem stresses the importance of an appeal to fundamental entities and the particularly causal nature of explanations. To those familiar with some of the contemporary literature in philosophy of science, these sound like rather stringent demands on what ought to count as an explanatory theory. However, the explanations to which Duhem refers are what we would currently consider to be *metaphysical* explanations. This is evident in his description of what he calls a “perfect theory” — a unified and explanatory theory of everything.

...[the] ideal and perfect theory . . . would be the complete and adequate metaphysical explanation of material things. This theory, in fact, would . . . be the very expression of the metaphysical relations that the essences that cause the laws have among themselves. [Duhem, Ariew, and Barker1996, 67-68].

Perhaps it is more appropriate to call the explanations that Duhem describes *metaphysical systems*. A metaphysical system is a complete ontology of the fundamental nature of the

world and a list of fundamental laws; were we able to access such a metaphysics, we would surely have an explanation of the laws and regularities that govern the appearances.

Were science to provide us with explanations, then an ideal scientific theory would be identical to the perfect theory Duhem describes. However, Duhem denies that science provides us with such explanations:

A physical theory is not an explanation. It is a system of mathematical propositions, deduced from a small number of principles, which aim to represent as simply, as completely, and as exactly as possible a set of experimental laws. [Duhem1954, 19]

On the basis of such disavowals of explanation, many scholars read Duhem as *antirealist* or *instrumentalist*. Instrumentalism about science is a form of pragmatism - a scientific theory true just in case it works (i.e. allows us to make predictions about observable phenomena). Science is an instrument that we use to manipulate and interact with the world, but it does not provide us with any new knowledge about the nature of the world.

Such a view is held in contrast to realism. Realists assert that the claims of scientific theories are straightforwardly true; scientific theories are sets of propositions and truthfully describe the world in which we live. Although antirealism is considered the denial of realism, the two views are actually the poles of a rather large continuum. Some philosophers are realists about *parts* of scientific theories — holding that some parts of physical theory are true while others are merely pragmatic. Precisely which parts of theories are taken to be true vary quite radically over the diverse versions of selective realism.

Traditionally, Duhem has been construed as a firm antirealist and one of instrumentalism's great forerunners. This analysis of his work is linked heavily to his disavowals of explanation.

For example, Nancy Cartwright argues that Duhem's position is almost identical to van Fraassen's constructive empiricism, and she relies on allegedly Duhemian worries for support of her antirealist evaluation of theoretical laws. The most famous component of van Fraassen's constructive empiricism is his claim that accepting a scientific theory does not include a commitment to the existence of unobservable entities in the theory. In his words,

Science aims to give us theories which are empirically adequate; and acceptance of a theory involves only a belief that it is empirically adequate (van Fraassen 1980, 12).

A theory is empirically adequate if it “fits together with” the observable data — that is, if all of the predictions a theory makes about observable entities prove true.

Cartwright attributes this view to Duhem, and takes him to oppose both theoretical laws, and theoretical entities that are postulated but unobservable.

Duhem has no quarrel with phenomenological laws, which can be confirmed by inductive methods. What he opposes are theoretical laws, whose only ground is their ability to explain [Cartwright1983, 88].

Cartwright further suggests that Duhem’s arguments rely on his rejection of inference to the best explanation. In contemporary terms, Duhem rejects the key premise of the *no-miracles* argument. The no-miracles argument is an argument first formulated by Hilary Putnam in “What is Mathematical Truth?” [Putnam1975]⁶. Putnam claims that if our theoretical laws aren’t at least approximately true descriptions of the world, the success of science is miraculous. Supposing that the success of science is not miraculous, it follows that theoretical laws must be at least approximately true. According to Cartwright, Duhem denies that the best explanation of the success of science is the accuracy of its laws.

Duhem rejects theoretical laws because he does not countenance inference to the best explanation. Neither van Fraassen nor Duhem are opposed to ampliative inference in general. They make a specific and concrete attack on a particular kind of inference which they see as invalid — inference to the best explanation — and thereby on the scientific realism to which it gives rise [Cartwright1983, 88].

According to Cartwright, because Duhem does not accept IBE, he also denies the realism associated with the no-miracles argument.

Like Cartwright, David Stump professes that Duhem holds some version of antirealism. He, too, bases this assessment of Duhem on the aforementioned anti-explanation passages.

Duhem argues that science should be seen as instrumentally adequate, not as explanatory ... Duhem is antirealist about both entities and theories ... [Stump1989, 338-340].

Both Stump and Cartwright takes Duhem’s disavowals of explanation as claims about the entities in our scientific theories. When Duhem rejects that science explains, he is rejecting the view that our scientific theories are true in a non-deflationary sense.

⁶For a further discussion of the no-miracles argument see [Matheson1998], [Psillos2013], and [FrostArnold2010].

Although it's easy to see how Duhem's claims about explanation lend themselves to an antirealist reading of *Aim & Structure*, the same text contains claims about natural classifications which resist any clear antirealist interpretation. For example, Duhem asserts

... the more complete [physical theory] becomes, the more we apprehend that the logical order in which theory orders experimental laws is the reflection of an ontological order, the more we suspect that the relations it establishes among the data of observation correspond to real relations among things, and the more we feel that theory tends to be a natural classification [Duhem1954, 26-27].

Here, Duhem suggests that as theories continue to grow and progress, we begin to suspect such theories are natural classifications — classifications which reflect the ontological order of “bare reality itself.” He further insists that

[Physical theory] assumes, while being completed, the characteristics of a natural classification. The groups it establishes permit hints as to the real affinities of things [Duhem1954, 30]

and

... the order in which theory arranges the results of observation does not find its adequate and complete justification in its practical or aesthetic characteristics; we surmise, in addition, that it is or tends to be a natural classification ... [Duhem1954, 335].

These declarations are unsettling in conjunction with Duhem's disavowals of explanation. Surely the “real affinities of things” lie behind the veil of appearances and are part of “bare reality itself.” But won't such classifications fit the mold of an explanation? How can it be that science fails to explain but manages to naturally classify?

To summarize, I characterize Duhem's notion of a natural classification as follows:

Natural Classification (NC): A natural classification is a classification that provides us with the natural divisions and classifications of “bare reality itself.” Such divisions are part of the nature of bare reality, and not contingent upon the world of appearances.

Antirealists typically assert that the claims of physical theory fail to describe the world. More specifically, antirealism is often formulated as the belief that scientific theories do not contain any information about the world beyond the experimental laws they are constructed to represent.

We observe certain regularities such as “water freezes at 273.16 K.” These regularities are experimental laws. The antirealist asserts that our theories restate and simplify collections of experimental laws, but fail to provide physical insight beyond the data already collected. Duhem denies this version of antirealism by maintaining that theory converges on a natural classification; if our theory provides us with hints of the non-phenomenal nature of things, it does more than just represent experimental laws.

Cartwright and Stump support deflationary readings of Duhem’s claims about natural classifications. Cartwright thinks that the kinds presented by a scientific theory are extremely “rough” and the kinds themselves are not unified by a “real” intension.

Duhem believes that phenomena in nature fall roughly into natural kinds. The realist looks for something that unifies the members of the natural kind, something they all have in common; but Duhem denies that there is anything. There is nothing more than the rough facts of nature that sometimes some things behave like others, and what happens to one is a clue to what the others will do [Cartwright1983, 95].

Stump emphasizes Duhem’s assertion that science can only *approximate* a natural classification, arguing that because scientific practice can never uncover a complete natural classification, theoretical laws are neither true nor false and the use of theories is primarily instrumental.

In contrast to Cartwright and Stump’s deflationary readings, philosophers such as Paul Needham and Andrew Lugg argue that the passages concerning natural classifications encourage a realist reading of Duhem. Needham argues that Duhem is a *moderate realist*, primarily on the grounds of Duhem’s comments about natural classifications. Needham takes natural classifications to be explanatory in a more contemporary sense (although they are held in contrast to Duhem’s notion of an explanation), and he takes Duhem’s insistence that science is a natural classification as an assertion that claims in a scientific theory are straightforwardly true.

I shall emphasize features [Duhem] stressed which are simply not accommodated by the usual instrumentalist, antirealist interpretation of his philosophy . . . I would call his notion of incorporation in the natural classification, illustrated by classical thermodynamics, an explanation; . . . and I see no reason to think he did not hold successful natural classification as the truth . . . My thesis is that he was a moderate realist in the sense of maintaining that whatever statements are regarded as meaningful parts of the body of scientific theory are held to be true [Needham1991, 102-103].

According to Needham, Duhem’s allegedly antirealist assertions are specific claims about the formal mathematics used in developing scientific theories. Only statements that have empirical content can be classified as true or false, and, Needham claims, Duhem is pointing out that some statements in our scientific theories have only mathematical (and not physical) content. The claims that do have physical content, however, should be read as “true” in a realist sense. However, Needham ignores the passages where Duhem explicitly considers explanation in order to focus on his comments concerning natural classifications and the approximate nature of experiment. While Needham’s interpretation is quite interesting, and certainly a step in the right direction, it fails to account for Duhem’s firm rejection of explanation.

Andrew Lugg maintains that Duhem is a *convergent* realist. Scientific theories approximate the truth — as science progresses, the theories it provides are closer and closer to the truth. While Lugg agrees with Stump’s analysis that theories can only ever approximate the truth, he denies it implies antirealism. On the contrary, while Duhem concedes that theoretical laws grow ever closer to perfection, he also asserts they already contain some true information beyond what can be captured by experimental laws. Lugg writes:

If anything Duhem espoused a version of what is nowadays called convergent realism. As we have seen he held that physics — left to its own devices — yields information about the nature of the world and that we are entirely justified in believing that its ontological claims are for the most part close to the truth . . . For him “physical theory confers on us a certain knowledge of the external world which is irreducible to merely empirical knowledge” and there is no avoiding the fact that a purely instrumentalistic physics would be of “meager importance” [Lugg1990, 417].

While Lugg and Needham take seriously Duhem’s claims about natural classifications, they make little effort to explicate his views on explanation. Needham suggests Duhem is referring to the formal elements of theories when he refers to explanation, but this fails to account for Duhem’s insistence that science does not go behind the veil. According to Lugg, ‘the butt of [Duhem’s] criticism is the view that physics provides ‘definitive explanations’; he was not against thinking of physics as directed towards the discovery of ‘provisional representations’ [Lugg1990, 4]. Still, this doesn’t account for Duhem’s strong insistence that science does not aim to produce an explanation.

More recently, authors have begun to acknowledge the difficulty of the tension in Duhem. According to Psillos, these difficulties are dissolved by attributing to Duhem an early version of structural realism. In his book *Scientific Realism*, Psillos alleges that the “essences beyond the veil” to which Duhem refers are, in fact, unobservable entities. Recounting the confusion about Duhem’s work, Psillos notes

Is Duhem’s position realist? It’s difficult to say, really. On the one hand, Duhem resisted to the very end, refusing to subscribe to atomism and other theories which posited unobservable entities. On the other hand, however, his adherence to natural classifications may be plausibly seen as a realist-enough position, given that Duhem understands ‘natural classification’ as revealing real relations among *unobservable entities* [Psillos1999, 38, emphasis his].

Psillos cashes out tension in Duhem as follows: To reject that science provides a **DE** is to deny the existence of the unobservable entities in our best scientific theories, while to posit that theory converges on a **NC** is to admit that the relations between the unobservables in our theory reflect the actual relations between entities. Psillos suggests this position is similar to John Worrall’s formulation of structural realism.

Worrall insists that to prove science successfully progresses, we must find a continuous thread between the differing models a theory uses over time. He claims to locate such thread through a view he calls “structural realism”. On this picture, even though our scientific theories are strictly false, and cannot be expected to posit the right entities in an ontology, they experience predictive success because they approximate the structure of the natural world. For example, Worrall comments about quantum mechanics:

The structural realist simply asserts, in other words, that, in view of the theory’s enormous empirical success, the structure of the universe is (probably) something like quantum mechanical [Worrall1989, 123]

In Duhemian terms, we assert that the classification provided in scientific theory is something like a natural classification, even though the theory is strictly false. To cast Duhem as a structural realist is to accept both Cartwright’s claim that Duhem took theoretical laws to be false and Lugg’s assertion that Duhem affirms scientific progress.

But ascribing structural realism to Duhem breeds an additional set of problems. From an exegetical standpoint, it’s a stretch to suggest, as Psillos does, that when Duhem refers to the “essential nature of bodies” he is denoting *unobservables*. The observable/unobservable

distinction, which has become a central point of contention in contemporary philosophy of science, does not appropriately map on to Duhem's disavowals of explanation and endorsements of natural classifications. This is a point I will illustrate fully in §4. Furthermore, to call Duhem a structural realist gives us only nominal insight into his view. Structural realists differ drastically concerning their definition of "structure," and to say Duhem is a realist about structure is not informative. In many ways, describing Duhem as a structural realist only returns us to our original problem: Duhem appears to be realist about something and antirealist about something else, though we haven't a clue how to clarify the "somethings."

The tension, then, is that Duhem denies that science aims at **DE**, where a **DE** is a peek "beyond the veil" at the true nature of physical bodies, and simultaneously insists that science aims to reflect a **NC**, where a **NC** is a peek "beyond the veil" at the classification of physical bodies according to their nature. But isn't a **NC** a type of **DE**? What is the difference between a scientific theory which aims at a **DE**, and a theory which aims at a **NC**?

The confusion is further intensified by Duhem's attitude towards atomism. Insisting that science ought to be autonomous from metaphysics, much of Duhem's work is devoted to the separation of physics and metaphysics.

[Physics] must be able to constitute itself through a proper method independent of any metaphysics. This method, which permits the study of physical phenomenon and the discovery of laws that connect them, without recourse to metaphysics, is the *experimental method* . . . *The experimental method rests on principles evident in themselves and independent of any metaphysics* [Duhem, Ariew, and Barker1996, 34, emphasis his].

Duhem repeatedly argues that our metaphysical opinions must not constrain the development of physics in any way. Additionally, he often criticizes *atomism* — the view that matter is ultimately reducible to rigid atoms and the interactions between them. He claims that atomism reflects a metaphysical opinion that ought not influence the development of scientific theory. However, many of Duhem's contemporaries came to ground their belief in atomism on experimental results and advancements in theoretical physics.

One example of experimental confirmation of atomism is the explanation of Brownian motion provided by Albert Einstein. In the early 19th century, Robert Brown used a microscope to observe the way molecules found in pollen grains traveled through water. No

existing scientific account could predict the trajectories of the molecules. However, in 1905, Albert Einstein published a paper titled “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen⁷” which explains the trajectory of the molecule on account of the motion of the atoms composing the fluid [Einstein1905]. For most scientists, this served as sufficient proof for the existence of atoms⁸.

Still, Duhem maintains his anti-atomist position in the 1906 publication of *Aim & Structure*, and did not rescind his injunctions against atomism in any publication before his death in 1916. Why did Duhem, an otherwise respected member of the scientific community, refuse to accept what his contemporaries welcomed?

Some philosophers, such as Cartwright, rely on Duhem’s separation of physics and metaphysics to explain his spurning of atomism. Cartwright claims Duhem rejects atomism because he denies the validity of inference to the best explanation. Therefore, even though Einstein explained Brownian motion by reference to atoms, we should not infer that such atoms exist and cause this motion.

However, not only does Duhem deny atomism, he often suggests that science should pursue a “peripatetic mechanics” — that is, a neo-Aristotelian approach to science which includes the Aristotelian distinction of “quality and quantity.”

To attempt to reduce all the properties of substances to shape and motion seems a chimerical enterprise, either because such a reduction would be obtained at the price of complications that would scare our imagination away, or even because it would be in contradiction with the nature of material things.

And so we are now obliged to accept into our Physics something other than the purely quantitative elements treated by geometers, to admit that matter has *qualities*; . . . we have to refasten our theories to the most essential notions of Peripatetic Physics [Duhem1980, 105, emphasis his].

Duhem certainly rejects the tenant of atomism that all natural phenomena can be reduced to “shape and motion.” But this stems from a technical worry, which I address in Chapter Five. Although Duhem resists atomism, at no point, which I am aware of, does he

⁷“On the movement of small particles suspended in a stationary liquid demanded by the molecular-kinetic theory of heat”

⁸For a further discussion of Brownian motion, see [Renn2005].

explicitly deny *the existence* of atomic particles. Instead, he denies that the whole of natural phenomena can be explained by a theory of atoms and their trajectories through space.

Specifically, Duhem worries that atoms won't be able to account for thermal phenomena. This is why he argues that we need a physics including *qualities*, or non-additive properties of matter.

Between the magnitude of a quantity and the intensity of a quality there exists an essential, profound distinction. Every quantity of specified magnitude can be obtained by adding one of other various quantities of the same kind and of smaller magnitude, which are its parts. There is nothing similar in the category of quality; ... juxtapose as many bodies as you will of which the heat intensity is that of boiling water — you will not make a body whose heat intensity is that of red-hot iron; heap up snowballs as you may, said Diderot, you will not be able to heat up an oven [Duhem1980, 3].

The distinction to which Duhem points does a great deal of work in classical thermodynamics albeit under a different name — “qualities” are called intensive properties “quantities” are called extensive properties. Extensive properties increase or decrease depending on the size of a particular system. A well known example of an extensive property is mass. The mass of a system increases as the size of a system increases. If you add a cup of water to a container of water, the mass of the system will increase. By contrast, intensive properties are properties of systems that do not depend on a system's size. A good example of an intensive property, which is mentioned by Duhem, is the property of temperature. If you add a cup of boiling water to a container which already contains boiling water, the temperature of the system will not increase.

Duhem asserts that fundamental physical theories must include qualities in order to properly account for the phenomena. In light of the importance of these qualities, he suggests that scientific theory hints at a modified Aristotelian metaphysics. In fact, contrary to Cartwright's suggestion, it appears that Duhem *invokes* some form of inference to the best explanation.

If we rid the physics of Aristotle and of Scholasticism of the outworn and demodded scientific clothing covering it ... we would be struck by its resemblance to our modern physical theory; we recognize in these two doctrines two pictures of the same ontological order, distinct because they are taken from a different point of view, but in no way discordant [Duhem1954, 310].

But if Duhem thinks physical theory provides evidence for an Aristotelian picture of

metaphysics, then he is not, as Cartwright suggests, opposed to inference to the best explanation. Duhem thinks there is significant metaphysical import to be found within the confines of scientific theory — but he denies that a theory provides us with a metaphysical system. It is evident that Duhem is trying to convey an important insight about the scientific method and its relationship to metaphysics, but nailing down just what he’s trying to point towards proves quite difficult.

This difficulty leads to a kind of puzzle: what does it mean to say that science does not aim at explanation but does converge on a natural classification? What is the difference between an explanation and a natural classification? The varied interpretations of Duhem focus on either his disavowals of explanation or his discussion of natural classifications, but fail to give an adequate account of both.

1.1 TOWARDS A BROADER CONCEPTION OF PHYSICAL THEORY

I suspect that contemporary difficulty in interpreting Duhem stems from current biases about how scientific theories evolve. Many contemporary philosophers assume that in order for a theory to be fundamental and/or explanatory, it must provide us with continuous trajectories of the entities which it describes. That is, the general structure of physical theory is to find differential equations which govern how an object evolves over time, provided that we input the appropriate initial conditions.

But, as Duhem was right to point out, much of science, even much of *physical* science, does not work this way. This *dynamic laws* picture of science ignores important technical worries about how to treat friction: the force which is caused by one surface coming into contact with another surface. For example, when a pendulum swings through the air, the contact between the particles in the pendulum and the air causes the pendulum to slow down. Likewise, if a ball is sliding on a surface, the friction caused between the ball and the table effect the velocity of the ball.

While friction is thought to result from electromagnetic forces in the particles that interact, these alleged interactions are far too complex to calculate. Therefore, to appropriately

describe phenomena, scientists are forced to rely on empirical methods, based in observation, in order to predict how friction will effect a system (rather than calculate this based on the fundamental principles of electromagnetic reactions). One way to get around the worries caused by friction, is to construct theories based on observations of systems in static states rather than systems in motion. This involves rather complicated methods of *ignoring* time, and further ignoring the continuous, dynamic motion of a system over time. Duhem refers to these methods as “virtual modifications” — and he believes they are at the core of scientific theory.

In this dissertation, I will illuminate these complicated methods and use them to shed light on debates in contemporary philosophy. In the next chapter, I connect Duhem’s notion of explanation with prevalent views in analytic metaphysics and philosophy of science. I provide a brief survey of the literature regarding scientific explanation, focusing on a thread that traces from Russell’s notion of causal processes to the causal mechanical model of explanation.

First I argue that the explanations Duhem wants to reject are similar to “causal lines” employed by philosophers attempting to distinguish naturally occurring causal processes from non-causal processes. Such “causal lines,” which allegedly appear in physical theories, are time-evolution trajectories of various systems and their components. It is usually assumed that any good physical theory will provide laws that generate time-evolution trajectories; I refer to such laws as dynamic laws.

Advocates of causal mechanical models of causation, such as Salmon and Dowe, claim that an explanation of some event E will describe the causal processes (or causal lines) that lead up to and constitute that event. In other words, it will cite the dynamic laws that provide a story of how the system ended up in its current state.

I argue that the kind of explanations Duhem worries about are explanations that invoke dynamic laws, such as those advanced in the causal mechanical model. That is, Duhem denies that science’s primary aim is to uncover explanations of the sort Salmon and Dowe suggest. Instead, Duhem argues, science aims to “classify” phenomena by uncovering delicate relationships between important, natural quantities that are often obscured in the search for dynamic laws.

Duhem's frustration with the dynamic-laws strategy, along with his distinction between explanations and natural classifications, is best illuminated by looking at several case studies in the history of physical theory. First I will consider the evolution of classical mechanics and Lagrange's use of a statics-first approach, and then turn to the development of thermodynamics and the discovery of entropy.

Lagrange's formulation of classical mechanics provides a detailed example of how variational principles and virtual modifications play a role in a physical theory. Prior to Lagrange, classical mechanics relied on a vectorial representation of classical mechanics (still commonly taught today). The most common sketch of the vectorial approach relies on Boskovichian point-particles. The location of a point particle is represented by a position vector extending between the particle and the origins of the coordinate system. The motion of the particle is defined by a specific equation, and the velocity of the particle is the time derivative of the function, and the acceleration is the second time derivative of the function. This is commonly referred to as the vectorial approach because the position, velocity, and acceleration are all represented by vectors.

Were we to look for causal mechanical explanations in classical mechanics, they would be found by tracing the causal lines depicted in vectorial mechanics. That is, we can use these vectors to define a detailed trajectory for each particle in our system, tracing out the motion of the particle which is caused by the applied force. Or so it seems. In practice, however, the application of this equation is surprisingly tricky due to worries that arise from boundary conditions and friction.

Lagrange provides a way out of these complications by applying a statics-first methodology that invokes hypothetical displacements rather than dynamic laws. This statics-first approach allows us to latch on to the important quantity of action. The derivatives of action are conjugate variable pairs, which allow us to track the conserved quantities in our system. These pairs actually allow us a peek into the "true affinities of things" without providing a detailed time-evolution trajectory of the particles in our system. In this sense, Lagrange's classical mechanics classifies without explaining.

Another helpful example of how the search for dynamic laws can obscure the discovery of important physical quantities is the history of entropy.

Entropy was discovered through the study of the Carnot cycle — a cycle which is traversed by a theoretical engine called a Carnot engine. The Carnot cycle is intended to represent the way that heat is converted into mechanical energy. The state space in which the Carnot cycle takes place is a dense series of equilibrium states. This is one reason why the cycle is considered an idealization: were a real engine to traverse a Carnot cycle, it must remain in equilibrium throughout the cycle.

Equilibrium states are states where a system is at rest. Therefore, it is impossible for a system to move, continuously, from one state of rest to another. In the same way I must interrupt my personal equilibrium to move from my couch to my bed, there will be a period of flux in any physical system as it moves from one state of rest to another. Such “intermediate states” are disregarded in the Carnot cycle — which is why the cycle is discontinuous with respect to time.

Because of such discontinuity, descriptions of the Carnot cycle fail to meet the criteria of the dynamic laws strategy and the demands of the causal mechanical model. The state space of thermodynamics does not reveal any “causal lines,” nor does it provide us with continuous time evolutions of how heat disperses. Nevertheless, it was through a careful study and revision of the Carnot cycle that engineers discovered the concept of entropy — which has proven to be an important physical quantity.

In the same way that the derivatives of action provide us with conjugate variable pairs in Lagrange’s formulation of classical mechanics, entropy couples together with temperature to define the internal energy of a system. Temperature plays the role of a generalized force while entropy plays the role of a generalized displacement, and the product of the two is used to determine the internal energy. This is structurally similar to the use of virtual work.

Once again we see how problematic frictional forces can be avoided by relying on a statics-first method which unveils important, previously unrecognized quantities such as entropy. These quantities allow us to classify physical phenomena, but, because they do not describe our system in terms of continuous time-evolutions, fail to provide “explanations” of the sort process theorists demand.

After a careful consideration of these two case studies, I return to Duhem’s philosophical writing and resolve the suggested tension in his work.

I argue that Duhem's disavowals of explanation are not intended as antirealist fodder, but instead signify his criticism of the dynamic laws approach. I unpack several passages where Duhem describes explanation in science, and reinterpret his seemingly vague and poetic prose as pointing towards important technical concerns such as problems that arise in coupling and the obscuring nature of friction.

Finally, I explore the implications of Duhem's distinction for contemporary philosophy of science.

I consider Cartwright's arguments that the use of such "idealization" in science merely reveals that "the laws of physics lie" and science does not provide us with knowledge about the natural world. I suggest, instead, that such laws are true and explanatory because they reveal delicate relationships between quantities, such as the conservation of energy, which *do* cause the observable phenomena to behave as it does, even if they do not provide detailed descriptions of this kind of behavior.

Duhem's project is important because it provides an alternative model for science. It is quite plausible that science can advance, progress, and provide important information about the natural world without producing fundamental dynamic laws. If this is right, as I indeed argue, philosophers ought to have an open mind about the direction in which science is likely to progress and what counts as a fundamental theory.

2.0 ON EXPLANATION

In the previous chapter, I present the following puzzle: Duhem denies that physical theory explains — specifically eschewing the view that science provides us with a complete metaphysical theory, while maintaining that science provides us with natural classifications that reveal peeks at the fundamental nature of reality.

In this chapter I have two goals — to connect Duhem’s criticisms to particular views in contemporary philosophy of science and metaphysics, and to provide a more detailed account of the “explanations” he rejects. I will direct Duhem’s criticisms towards “ideal DN texts” and “causal process theories,” arguing that both are versions of what I call a *dynamic laws strategy*. In spelling out the assumptions at work in these views of explanation, I will also articulate the aspects of such accounts that Duhem finds problematic.

2.1 DEDUCTIVE-NOMOLOGICAL ACCOUNTS OF EXPLANATION

Contemporary discussion of scientific explanation originates with Hempel’s deductive nomological (DN) model. On DN models of explanation, a particular phenomena (called the *explanandum*) is deduced from a series of premises (called the *explanans*)¹. The *explanans* are general scientific laws and the initial conditions of the system. For example, a DN explanation might proceed as follows:

Explanans: Water freezes at 273.16 K.

Explanans: This water reached a temperature of 273.16 K.

Explanandum: This water is frozen.

¹For a full account of the DN model, see [[Hempel1965](#)].

For a DN explanation to be valid, the sentences containing the *explanans* must be true and the *explanandum* must be a logical consequence of the *explanans* [Hempel1965]. The form of valid DN explanations is intended to mirror a valid derivations in deductive logic. These explanations are nomological, because they depend on the invocation of true natural laws. Hempel managed to formalize the naïve view of scientific explanation: it is the job of science to uncover true, natural laws and the phenomena we observe is grounded in some combination of these laws. That is, each event can be explained by an appeal to some subset of the natural laws and a system’s initial conditions.

While Hempel’s account effectively captures philosophical intuitions about science, it lacks the precision required to answer many philosophical questions. Specifically, Hempel’s account fails to define what counts as a “true natural law.” Scientific practice operates at a plurality of scales: psychologists are interested in laws and regularities of human behavior, while neuroscientists are interested in the laws which govern neuron firings in the brain. This same hierarchy of scales can be seen in physics: scientists working on fluid mechanics tend to ignore the behavior of individual atoms, while atomic physicists often ignore the motion and behaviors in the nucleus of the atom — leaving this work to the nuclear physicists. Physical theories at each scale provide laws about the behavior of the material studied. Are all of these natural laws? Are any of them?

A popular trend in philosophy is the assumption that the true, natural laws are the most “fundamental ones” — that is, the laws “at the bottom” are the only laws which count. The reasoning for such a view is rarely made explicit, but the argument runs roughly as follows:

1. A true, natural law must be without exception and inviolable.
2. If the behavior of matter at the “bottom level” of the universe is determined by natural laws, then this behavior is independent of the apparent laws at any other level of observation².
3. The “bottom level” includes basic pieces of matter and whatever forces or fields these pieces interact with.
4. Everything is composed of these basic materials.

²(1) is thought to entail (2) by means of the following reductio: Suppose (2) is false. It then follows that the behavior of matter at the “bottom level” is contingent upon some higher level laws. But if this is true, then the laws at the fundamental are either are not inviolable. Therefore, (2) must be false.

5. Therefore, knowledge of the basic materials and the laws which govern them provides us with a complete, physical story of the evolution of the universe.

To those who accept this argument, or some similar version of it, it follows that the only true, natural laws are those which govern the basic elements of the universe. For example, Douglas Kutach argues that all phenomena is ultimately derivable from the laws of fundamental physics.

...there are some fundamental laws of physics that govern the behavior of all particles and fields and ... ordinary macroscopic objects are merely aggregates of these fundamental microscopic parts [[Kutach2013](#), 20].

But if philosophers like Kutach are right to think that the only candidates for natural laws are the laws of fundamental physics, what about the wide variety of apparent explanations that science provides at a much higher level? Or, what about ordinary explanations that don't seem to include natural laws whatsoever?

For example, consider the following explanation:

Explanation Pittsburgh (EP): The water froze because I left it in my car overnight during the winter in Pittsburgh.

This seems like a reasonable explanation, but it fails to exemplify the DN model because it doesn't explicitly state any true laws of nature. One attempt to get around these worries is the suggestion that such apparent explanations are only placeholders for true explanations.

For example, on Hempel's view, explanations like this are mere sketches of more accurate DN explanations that employ natural laws in their formulation³. Peter Railton argues that so-called explanations, such as EP, are actually summaries of an "ideal DN text" — a text that contains a complete DN explanation for each particular phenomena⁴ ([[Railton1978](#)], [[Railton1981](#)]). In fact, almost all scientific explanations lack the detail needed to exemplify a DN explanation, and are therefore "summaries." Railton's ideal text contains a complete

³The first example of an explanation is also a sketch, albeit a better one. 'Water freezes at 273.16K' still lacks certain information (e.g. how we define water, how homogenous the temperature of the water is) that a complete explanation would include. These allegedly "incomplete" laws, like 'water freezes at 273.16' are often referred to, rather flippantly, as *ceteris paribus* laws.

⁴Railton actually refers to an ideal "ideal D-N-P text," where the P includes explanations that rely on probabilistic laws. I will set this aspect of his ideal text aside here, as the only thing it adds to this discussion is difficulty.

listing of the laws and initial conditions that contribute to the water freezing in the car (e.g. facts about the particles in the water, the insulation provided by the bottle and car, the thermal exchanges between the car and its environment, etc.). According to Railton, science is successful insofar as it approximates these ideal explanations. The closer that a particular explanation gets to a complete DN explanation, the more “explanatory” it is.

Railton’s account is not systematically *reductive* — it is not the case that macro laws can clearly be deduced from micro laws. Instead, he thinks that these macro level explanations convey “partial information.”

It is hardly novel to speak of sentences providing information about complete texts in this way: presumably we employ such a notion whenever we speak of a piece of writing as a summary, paraphrase, gloss, condensation, or partial description of an actual text such as a novel. Unfortunately, I know of no satisfactory account of this familiar and highly general notion . . . nor can I begin to provide an account of my own making⁵ [Railton1981, 240].

According to Railton, Hempel’s account of explanation appropriately characterizes complete, fundamental explanations. In contrast, the explanations provided through scientific research reveal only summaries or paraphrases of these texts. Moreover, these fundamental texts take on a certain kind of structure: the laws they provide are, presumably, initial value problems — differential equations which take initial conditions of a system as an input, and output the system’s trajectory through time. The suspected structure of this fundamental DN text can be further illuminated by consideration of the popular “causal process theories” in philosophy of science.

2.2 CAUSAL PROCESS THEORIES

Causal process theorists profess to find clues about the nature of causality from “the actual practice of science,” rather than raw conceptual analysis. Sundry forms of causal process theories are united by the hypothesis that physical theories detail processes that occur in nature. That is, physical theories tell the story of the path a system travels through time.

⁵For an argument as to why it is actually impossible to provide this sort of account, see [Wilson2011]

These paths are thought to be either *causal* or non-causal.

Bertrand Russell is considered a forefather of causal process theories. Russell built his account of causation in science on a notion of “causal lines⁶.” Causal lines are series of events which are connected in such a way that we can infer information about the events at the end of the series from information about events at the beginning of the series. Moreover, Russell argued that such lines are *continuous*, because they reflect the persistence of a particular object through time.

A causal line may always be regarded as a persistence of something, a person, a table, a photon, or what not. Throughout a given causal line, there may be constancy of quality, constancy of structure, or gradual changes in either, but not sudden change of any considerable magnitude [Russell1948, 455-457].

Recent defenders of causal process theories, such as Salmon and Dowe, think that the causal lines Russell describes are *world lines* through Minkowski space time. The world lines of interest — that is, those Salmon and Dowe consider to be causal processes — are those in which some important quantity can be tracked.

The Salmon-Dowe account of causation reduces cause to the transmission and conservation of a quantity (such as energy) from one system to another. To illuminate how such quantities are tracked, Salmon uses an example about a baseball breaking a window. The linear momentum of the baseball just before it impacts the window is equal to the total linear momentum of the shattered pieces of glass after the window is broken. Because momentum is conserved between the baseball and the shattered glass we can conclude that the baseball caused the window to shatter. Salmon articulates how we can be sure it was the baseball that causes the window to shatter and not a nitrogen atom that hits the window at the same time:

“ . . . the interaction constituted by the nitrogen molecule and the shattering window, momentum is not conserved. Take the window to be at rest; its linear momentum is zero. The linear momentum of the nitrogen molecule when it strikes the window is not zero, but fairly small. The total linear momentum of the pieces of the shattered window after the collision is enormously greater than that of the incoming molecule. In contrast, the total linear momentum of the baseball as it strikes the window is about equal to the momentum of the pieces of glass and the baseball after the collision. So if we talk about causes and

⁶The view I describe here is Russell’s later view, concerning what he takes to be a scientific notion of cause which is something like a distant cousin of the philosophical notion of cause that he criticizes in [Russell1912].

effects, we are justified in saying that the window was broken by the collision with the baseball, not by the collision with the nitrogen atom” [Salmon1994, 304].

Roughly, the Salmon-Dowe account claims that X causes Y if and only if some relevant quantity is conserved in the interaction of X and . While Salmon and Dowe provide a more precise definition of the notion of cause, I am primarily interested in the general paradigm which the view assumes, rather than the specific details of the account.

According to causal process theorists, our best scientific theories provide us with *processes* — or, in more specific terms, time-evolution trajectories⁷. On Salmon and Dowe’s picture, scientific theory tell us the continuous paths that conserved quantities travel.

As it turns out, the supposition that physical theories provide continuous trajectories places a rather rigorous constraint on what scientific theories must do — a constraint so rigorous that it *rules out* a large number of theories that are currently in use in scientific practice. Physicists often use methods for predicting the behavior of systems which do not provide us with continuous trajectories. Mark Wilson provides many examples of these alternative methods in his paper *Physics Avoidance*. Wilson argues that we are quite frequently able to get a hold on certain philosophical notions (like cause) by appealing to different techniques invoked in applied mathematics. Moreover, many of these useful mathematical treatments of physical systems do not provide us with *continuous trajectories*, nor do they rely on “laws” which come in the form of initial value boundary problems.

For example, Wilson discusses the strategy frequently employed in scientific practice of searching for a system’s equilibrium conditions. I will delve into this example in much more detail over the next two chapters, but the short story is that practitioners of science often aim to describe *where a system might come to rest* rather than describing *the path the system will take to rest*. To illustrate this idea in ordinary thinking, Wilson uses the example of Jack and Jill falling down a hill. If asked to predict where they will land, the answer is often “at the bottom of the hill.”

This type of reasoning reflects formal mathematical tools used to determine a system’s end state — a practice often used in science to *avoid* trying to describe a system’s evolution.

⁷Salmon actually remains agnostic about whether these theories include trajectories over time: he thinks it’s possible that time just is a reflection of a causal ordering. Nevertheless, these trajectories must be continuous.

Although such strategies are often employed in the practice of science, they fail to provide the “causal lines” demanded by process theories of causation, and also elude the expectations for scientific theory outlined by an ideal DN text. Both Railton and Salmon are making a similar mistake — they ignore a wide subset of scientific practice, ruling it superficial, because it fails to conform to their *a priori* expectations of what scientific theories must be like.

I suggested in the first chapter that Duhem defines an explanation as something like a “metaphysical system.” Ideal DN texts and continuous time evolutions fit quite nicely into Duhem’s picture of a metaphysical system. To understand how, it’s helpful to think about Duhem’s own notion of an ideal theory.

2.3 AN EXPLANATION IS A PERFECT THEORY

Duhem’s concept of an explanation is heavily influenced by his views on the relationship between physics and metaphysics. In *Physics and Metaphysics*, he defines physics as the experimental study of inorganic matter and metaphysics as an investigation of its essences.

To conform to contemporary usage, we give the name *physics* to the experimental study of inanimate things, considered in three phases: the observation of facts, the discovery of laws, and the construction of theories. We regard the investigation of the essence of material things, insofar as they are causes of physical phenomena, as a subdivision of *metaphysics*. This subdivision, together with the study of living matter, forms *cosmology* [Duhem, Ariew, and Barker1996, 30].

By this distinction, the study of material phenomena is the business of physics, while the study of the *causes* of this phenomena belongs to metaphysics. Duhem also thinks there is a separate field of study involving organic matter and its essence. This, in conjunction with metaphysics, is what he calls cosmology.

For Duhem, physics consists in the observation of facts, the discovery of laws and the construction of theories. It’s important to note that he is referring to “experimental laws” rather than the more metaphysical notion of a “natural law.” Experimental laws describe regularities which are observed in experimental contexts.

He further elucidates the distinction between physics and cosmology in *Physics of a Believer*.

If the physicist and the cosmologist study at the same time the laws of chemical combination, the physicist will wish to know very exactly what the proportion is among the masses of the bodies entering into combination, under what conditions of temperature and pressure the reaction may take place, and how much heat is involved. The preoccupation of the cosmologist will be quite different: observation shows him that certain bodies, viz., the elements in the combination have at least apparently ceased to be, and that a new body, viz., the chemical compound, has appeared; the philosopher will strive to conceive what this change of mode of existence really consists in. Do the elements really subsist in the compound? Or do they persist in it only potentially? Such are the questions he will wish to answer [Duhem1954, 300].

For Duhem, the process of science is largely descriptive: scientists seek to provide an increasingly detailed account of how physical changes occur and the precise conditions under which they happen. These accounts are composed of experimental laws — laws which describe the regularities observed in experiment. Cosmologists, on the other hand, are interested in far-reaching general truths. Rather than know precisely what proportion of elements make up a specific mixture, they wish to know if this mixture is a unity or an aggregate, if it is one thing or many things.

Cosmologists seek an *explanation*: a detailed account of general natural laws and the objects they govern. Duhem refers to such a theory as a “perfect theory” which “would be the complete and adequate metaphysical explanation of material things”. It is this type of explanation that Duhem denies is provided by science [Duhem, Ariew, and Barker1996, 68]. Were it the case that the experimental laws of science paralleled the cosmologist’s idea of a natural law, then certainly science could straightforwardly answer the questions of the metaphysician. More specifically, were science to furnish us with an ideal DN text, which contained a detailed picture of all the basic elements and the deterministic paths they follow, physical theory would make a clear cut contribution to metaphysics — from such a theory, we would be able to derive metaphysical truths. Duhem discusses such a derivation as it pertains to cosmological questions about freedom of the will.

How would one go about deriving from the principle of the conservation of energy and from other analogous principles the corollary, “Free will is impossible”? We should observe that these various principles are equivalent to a system of differential equations ruling the changes of state of the bodies subject to them; that if the state and motion of these

bodies are given at a certain instant, their state and motion would then be determined unambiguously for the whole course of time; and we should conclude from this that no free movement can be produced among these bodies, since a free movement would be essentially a movement not determined by previous states and motions [Duhem1954, 286].

Were we to have an ideal DN text, which could tell us the direct path of any particular body given any special set of initial conditions, then free will could be ruled out because all the states of the system could be determined by the differential equations which govern it⁸. Essentially, these “law-like” differential equations would demand that a system behave in a certain way. If that were the case, then how could we consider the system to be free? Its behavior would be fully determined by the laws. However, Duhem denies that the scientific theories we possess can be used to derive these conclusions of grandeur.

Now, what is such an argument worth? We selected our differential equations or, what comes to the same thing, the principles they translate, because we wished to construct a mathematical representation of a group of phenomena; in seeking to represent these phenomena with the aid of a system of differential equations, we were presupposing from the very start that they were subject to a strict determinism; we were well aware, in fact, that a phenomenon whose peculiarities did not in the least result from the initial data would rebel at any representation by such a system of equations. We were therefore certain in advance that no place was reserved for free actions in the classification we had arranged. When we note afterwards that a free action cannot be included in our classification, we should be very naïve to be astonished by it and very foolish to conclude that free will is impossible [Duhem1954, 286].

In this example, Duhem emphasizes that the law of the conservation of energy represents observed data — and this data follows determinate patterns. Therefore, this data can appropriately be represented by a system of differential equations. But this does not mean these differential equations are *metaphysical* laws — that they dictate, rather than describe, the behavior of inorganic objects. Because they are not metaphysical laws, we can’t think of them as explanatory. Rather, they describe the behavior of a chosen set of objects *within certain limits*.

Duhem is not denying that these laws are true — he merely denies that we can understand them as metaphysical laws. In the literature on natural laws, many philosophers have argued that the laws of science are not “real” laws because they are not universal, they do not have

⁸Duhem is setting aside the possibility of compatibilism — the view that free will is compatible with a deterministic universe.

unrestricted scope, etc. For example, Nancy Cartwright has argued that the laws of physics “lie” because they appear general but are, in fact, highly contingent laws which depend on the fixing and holding of numerous external conditions⁹. But, on Duhem’s view, this presupposes a metaphysical picture of what a law must be like. It’s true that metaphysicians wish for broad, universal laws — laws which will allow us to deduce particular instances of phenomena when given the correct initial conditions. However, Duhem thinks this is an unrealistic and inaccurate portrayal of how the laws in science work: moreover, just because experimental laws are not metaphysical laws, it does not mean they are “false” or “inaccurate representations of the world.”

A related point about entropy also helps clarify Duhem’s point.

In the middle of the last century, Clausius, after profoundly transforming Carnot’s principle, drew from it the following famous corollary: The entropy of the universe tends toward a maximum. From this theorem many a philosopher maintained the conclusion of the impossibility of a world in which physical and chemical changes would go on being produced forever; it pleased them to think that these changes had had a beginning and would have an end; creation in time, if not of matter, at least of its aptitude for change, and the establishment in a more or less remote future of a state of absolute rest and universal death were for these thinkers inevitable consequences of the principles of thermodynamics.

The deduction here in wishing to pass from the premises to these conclusions is marred in more than one place by fallacies. First of all, it implicitly assumes the assimilation of the universe to a finite collection of bodies isolated in a space absolutely void of matter; and this assimilation exposes one to many doubts. Once this assimilation is admitted, it is true that the entropy of the universe has to increase endlessly, but it does not impose any lower or upper limit on this entropy; nothing then would stop this magnitude from varying from $-\infty$ to $+\infty$ while the time itself varied from $-\infty$ to $+\infty$; then the allegedly demonstrated impossibilities regarding an eternal life for the universe would vanish. [Duhem1954, 287-288].

Clausius managed to prove that entropy increases and tends towards a maximum. Many philosophers take this to mean that the universe must reach a definite end. Duhem first argues, above, that such a conclusion is too hasty and does not properly take into account the scientific findings. However, this is not the real substance of his concern. His main objection to this reasoning is that it assumes the experimental laws of science have an unrestricted scope — that the experimental laws scientists discover should be considered true without qualification.

⁹For more on Cartwright’s view, see [Cartwright1983].

But let us confess these [above] criticisms wrong; they prove that the demonstration taken as an example is not conclusive, but do not prove the radical impossibility of constructing a conclusive example which would tend toward an analogous end. The objection we shall make against it is quite different in nature and import: basing our argument on the very essence of physical theory, we shall show that it is absurd to question this theory for information concerning events which might have happened in an extremely remote past, and absurd to demand of it predictions of events a very long way off [Duhem1954, 288].

According to Duhem’s understanding of physical theory, experimental laws all have limits: they are true within a limited scope. While we might be able to get a handle on how a system will evolve locally, it is often very hard (if not impossible) to make sense of how such a system will behave over a prolonged period of time. This worry is especially clear in the context of chaotic systems such as weather. While chaotic systems are still considered to be deterministic, the slightest change in initial conditions can cause an exponential difference in the state of the system as it evolves through time. Hence, we can’t make any predictions about how the weather is going to behave three months from now — even if we think there is some kind of determinate answer to it.

Of course, someone such as Railton would argue that an ideal DN text would provide us with precisely the detail needed to specify the initial conditions in a way that we could track the evolution of even the most chaotic systems. Duhem does not deny such a theory could exist or even be true: he merely denies that this sort of theory is what science provides us with. Rather, science works to describe, in as much detail as possible, the real regularities observed in experimental contexts. Science aims to describe and classify experimental laws.

2.4 DYNAMIC LAWS VS. CLASSIFICATIONS

I began this chapter by presenting two popular views of scientific explanation at work in contemporary philosophy of science — Railton’s “ideal text” hypothesis, and world line analysis via causal process theorists. Both of these accounts presuppose that the goal of science is to provide us with differential equations which, when paired with the appropriate initial conditions, will reveal how quantities evolve over time. To Salmon, Dowe and Railton, such differential equations are “true natural laws,” meeting Hempel’s stipulations on acceptable

premises in DN explanations. I will refer to accounts which presuppose this view of scientific theory as *dynamic laws strategies*. According to the dynamic laws strategists, science seeks to uncover true, dynamic laws of the sort described above.

However, Duhem denies that science aims to provide us with explanations — where explanations are theories that furnish us with these explanatory, dynamic laws. I argue that Duhem’s disavowals of explanation are really arguments that science does not attempt to furnish us with metaphysical laws, such as these dynamic evolutions, nor does it proceed towards an ideal DN text.

His disavowals of ideal DN explanations notwithstanding, Duhem insists that science does more than represent the data. Good theories order the results of observation in a way that hints at the real affinities of things, tending towards “natural classifications.” Questions about how these orderings organize data and what Duhem means by a natural classification have caused a great deal of exegetical confusion. In the next chapters, I will provide two case studies which illuminate this distinction. For now, I merely wish to foreshadow the results, and demonstrate the contrast between DN explanations and natural classifications through a discussion about the pendulum.

2.5 THE PENDULUM

A pendulum is a simple machine that consists of a mass, called a bob, on a string attached to a pivoting point. When the bob hangs directly beneath the pivoting point, it remains at rest unless disturbed. This rest position is known as the pendulum’s equilibrium state. When the bob is moved to one side or the other, a restoring force presses the bob back towards its rest position.

In elementary physics class, pendulums are often used to illustrate the conservation of mechanical energy via exchanges between kinetic energy and potential energy. Kinetic energy is the energy associated with the motion of the bob (i.e. its speed), while potential energy is the energy associated with the bob’s disposition towards movement (i.e. its height). As the pendulum moves away from its equilibrium position, the potential energy increases and

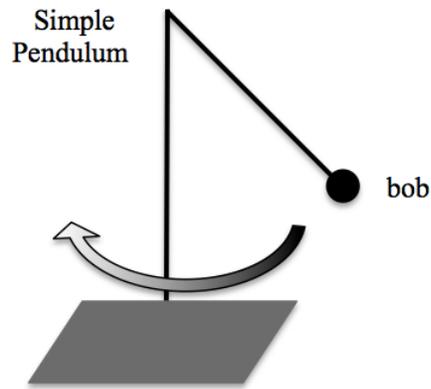


Figure 1: An example of a simple pendulum.

the kinetic energy decreases. In contrast, when the pendulum moves towards equilibrium, the kinetic energy increases and the potential energy decreases. The pendulum is used to emphasize that kinetic energy is continuously exchanged with potential energy.

But, in truth, we witness this posited interchange only approximately in any real life set up, for numerous frictional effects continually drain away kinetic energy before it can be completely reconverted back into potential energy. In fact, this apparent loss explains why the conservation of energy, in its modern understanding, wasn't recognized as valid until the 1840s (earlier writers thought that the energetic losses needed to be carried away as the kinetic energy of small particles, which isn't true). Actual losses of potential energy are caused by frictional forces such as air resistance. This causes the overall mechanical energy of the system to decrease, showing that the energy is not perfectly conserved in such exchanges.

There are two ways a scientist might approach the pendulum, which correspond to Duhem's distinction between explanations and natural classifications.

2.5.1 The dynamic laws picture of the pendulum

To provide an adequate DN explanation — that is, give a detailed account of everything that occurs over time in our system —, we will need to carry out a fully adequate dynamics laws modeling. On this picture, we try to straightforwardly write down differential equations that will capture the exact motions of the bob over time. But such a feat is not easily obtained. We cannot write down differential equations that perfectly describe the motion of the bob because pendulums, like all real-world systems, face the effects of friction in its various forms (e.g. air resistance, the extension and mass of the cord on which the bob swings, three dimensional motion, contact between the string and the pivot, etc.). To uncover an ideal DN explanation, we must quantify all of these forces and explain how they interact with our bob¹⁰.

Even if we were *able* to mathematically describe the motion of the bob, the equations used would obscure the exchanges between potential and kinetic energy. From a dynamic point of view, it appears as if nothing is actually being conserved as our pendulum moves through time. The pendulum loses mechanical energy — this is why it eventually comes to rest¹¹. As Duhem puts it, potential energy is a very abstract concept — identifying it proved no small task. We must be careful not to underestimate the subtleties of the notion.

To deal with the complications of describing the actual motion of the pendulum, scientists generally disregard the effects of friction and employ “simplifying assumptions” about the system. They assume that the pendulum is completely isolated, the string the bob hangs from is massless and remains taut, motion can only occur in two dimensions, and the system will not lose any energy due to friction. This simplified pendulum can then be described by the differential equation

$$\frac{d^2\theta}{dt^2} + \frac{g}{l}\sin\theta = 0 \quad (2.1)$$

where g is the acceleration due to gravity, l is the length of the rod connected to the mass, and θ is the angle between the rod and its initial position.

¹⁰Another problem that arises in this context concerns constraint forces. I’ve decided not to address constraint forces here, for the purpose of simplicity. However, it should be noted that some of the forces we have to take into consideration are co additionally problematic because they function as unknown quantities in our problem rather than as determinant forces. For a deeper discussion of constraint forces, see [Wilson2011].

¹¹This is also why perpetual motion machines are impossible.

Philosophers refer to such pendulums as “idealized,” but that deflationary terminology doesn’t render the exchanges between kinetic and potential energy storage their proper due. We now believe that those back-and-forth exchanges capture nature’s most intimate workings, whereas the “corruption” of friction merely represent the inevitable “noise” that intrudes as soon as we insert any physical system into a natural environment. Duhem’s focus on “natural classifications” represents an insistence that science’s principle task is to uncover these deeper patterns of exchange, rather than embark upon a quixotic voyage to model the friction-laden bobs of unfettered nature with wholly accurate equations, such as a dynamic laws picture suggests.

But what kind of investigative methodology is likely to result in “natural classifications” of the sort we seek? Duhem proposes a novel answer, based upon his deep historical studies of prior discoveries of abstract quantities similar to potential energy.

2.5.2 Duhem’s approach to the pendulum

Rather than generate differential equations which describe the trajectory of the bob, Duhem suggests we consider “freeze-frames” of the pendulum in each of the states it will pass through. We can experimentally determine the static force needed to hold the mass in place at each point in a trajectory, and its speed as it passes through that point. We can then create a series of static “balances” between kinetic and potential energy along the general path of the pendulum.

The kinetic energy and potential energy vary in such a way that the total mechanical energy stays the same, making the total energy a conserved quantity. Scientists experimentally determine the balances of potential and kinetic energy for any particular position of the pendulum. These static results are strung together to determine a frictionless trajectory for the pendulum. Imagine creating a beaded necklace: each bead represents one state of the pendulum and the string is the index of time on which the beads are organized.

Frictional forces arise when a system is in motion. Duhem suggests we create an imaginary frictionless history for the pendulum in which frictional effects get suppressed through gentle “quasi-static” shifts between one stage and another. Such shifts are called quasi-static

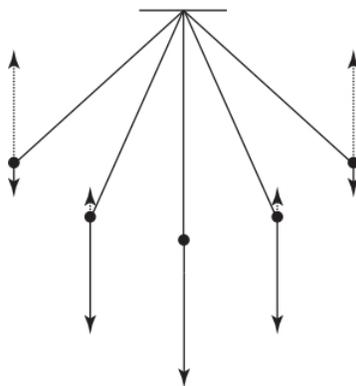


Figure 2: An illustration of Duhem’s perspective of the pendulum. In each state of the pendulum, the dotted line represents the potential energy and the solid line represents the kinetic energy. The total energy remains the same.

because we imagine the system moves from one state of rest (or equilibrium) to another. We imagine the pendulum moves through a series of static states. This strategy for depicting the trajectory of the pendulum allows us to capture the motion of a pendulum without considering the energy lost to friction — opening our eyes to critical changes between quantities like potential and kinetic energy. By contrast, if we aim to provide DN explanations of the pendulum’s trajectory, crucial relationships between conserved quantities become obscured. This is because a DN model attempts to trace the *actual* path of the pendulum, along which energy is not conserved.

The continuous motion of the pendulum described by equation (1) is a *virtual* motion: it does not represent the actual trajectory a pendulum follows. Duhem refers to these virtual motions as *virtual modifications*. Rather than write such virtual modifications as “idealizations” which can only “approximate” truth, Duhem argues that such modifications help physical theory to capture the deeper tendencies of nature.

In the next two chapters, I provide a more detailed examination of how these quasi-static processes play a significant role in the development of physical theory.

3.0 LAGRANGIAN MECHANICS

3.1 MOTIVATION FOR LAGRANGE'S MECHANICS: CONSTRAINTS IN THE POINT MASS SETTING

To understand Lagrange's formulation of classical mechanics, it is helpful to first consider the point particle formulation of Newton's laws (often referred to as Newtonian mechanics or the vectorial approach). The basic equation of Newtonian mechanics states

$$m_i \ddot{r}_i = F_i^{(e)} + \sum_j F_{ji} \quad (3.1)$$

where m_i is the mass of some particle i , \ddot{r}_i is the second time derivative (or acceleration) of i , $F_i^{(e)}$ is the external force on i , and F_{ji} is the force of another particle in the system, j , on i . Theoretically, this set of differential equations allows us to calculate the behavior of a system of particles if we know the forces applied to them (and the masses of the particles). The physical idea is that the force acting on a particle in any definite direction is equal to the product of the particle's mass and acceleration in that direction. The framework these calculations occur in is referred to as the point-mass setting because particles are represented as *point masses*: rigid, distinct, extensionless mathematical points.

On this view, the Newtonian formulation of classical mechanics appears to be a successful and relatively simple dynamic laws approach to describing the behavior of mechanical systems. Just find out what forces are acting on your particle, and you will be able to predict the behavior of a particular point mass¹. Moreover, because this is an initial value boundary

¹In practice, however, when the system is larger than three particles, the differential equations become unsolvable and no trajectory can be predicted

problem, you will have a special kind of knowledge about the behavior of the particle: you can predict the trajectory of the particle as it evolves through time.

The ability to predict the particle's trajectory through time is related to the signature of the ordinary differential equations involved in representing Newton's laws. The solutions to these equations are well-defined initial value problems. All this means is that when you specify a point in the domain of the function (in this case the position and velocity of the particle), the equation provides the evolution of the particle over time.

Due to the signature of the ODEs involved in the point mass setting, and the simplistic nature of the primary equation used to represent the motion of a system of particles, this treatment of Newton's laws is initially enticing to philosophers. However, to capture and describe the motion of most systems of particles in the actual world involves complicated methods which are more involved than straightforward mathematical solutions to ODEs. One of the first things the scientist attempting to apply Newton's laws to any specified system must consider is the effect on the motion of the particle that will occur as a result of the systems constraints.

A constraint is another physical system that interacts with the system of particles that we are interested in. For example, consider a bead whose motion is constrained to a slanted surface, a gas in a container, or a system of molecules bound together as a rigid rod. In all of these cases, the otherwise "pure" trajectory of each particle is at least partially determined by external constraints on the system. The naïve way of thinking about constraint forces suggests that they are just another external force on the system. Their presence doesn't do anything to negate the previously suggested equation of the system; they merely indicate a new set of external forces. However, a careful examination of how such forces are calculated shows this picture is inaccurate.

One classic textbook summarizes the situation by saying

... one might obtain the impression that all problems in mechanics have been reduced to solving [a] set of differential equations. One merely substitutes the various forces acting upon the particles of the system, turns the mathematical crank, and grinds out the answers! Even from a purely physical standpoint, however, this view is oversimplified. For example, it may be necessary to take into this account *constraints* that limit the motion of the system [Goldstein, Poole, and Safko2002, 12].

To handle constraints in the point mass setting, additional equations must be introduced. For example, if a mass is constrained to move around in a circle, we would introduce the equation

$$x^2 + y^2 = r^2, \tag{3.2}$$

which describes the spherical shape of the surface the mass is constrained to. This equation describes the boundaries of motion for our mass. Unfortunately, such equations present themselves as additional information which needs to be solved for rather than determinate forces whose sum dictates the trajectory of each particle. As Goldstein writes,

[The equations of constraint forces] are among the unknowns of the problem and must be obtained from the solution we seek. Indeed, imposing constraints on the system is simply another method of stating that there are forces present in the problem that cannot be specified directly but are rather known in terms of their effect on the motion of the system [Goldstein, Poole, and Safko2002, 13].

Part of the reason why these constraint forces are indeterminate is their reactionary nature; the way in which they interact with the system depends in part on how the system interacts with the constraints. They are not determinate forces whose value is always the same, but forces of reaction which vary depending on the force that is exerted by the system. Mark Wilson discusses this problem in his paper “What is Classical Mechanics Anyway?”

[P]oint mass mechanics only tolerates forces that *approximately implement* standard constraints. Why? For starters, no array of point particles exerting plausible forces can hold our “bead” to any fixed geometrical contour such as [a wire] \mathbf{W} . . . Now suppose that we were able to arrange a schedule of forces arising from the wire \mathbf{W} that could bind bead \mathbf{b} traveling at velocity \mathbf{v} perfectly to \mathbf{W} ’s unyielding surface. Can these same forces perform the same chore if the velocity of the bead had instead been \mathbf{v}^* ? To do that, the forces will need to be *velocity sensitive* (because \mathbf{W} must exert stronger forces upon a faster \mathbf{b} to pull \mathbf{b} to its requisite destination) [Wilson2012, 25].

Wilson points out that whatever these constraint forces are that bind the bead to a string, they operate in a way that adapts to the changing velocity of the bead. No matter how quickly the bead goes, it will still follow the contours of the wire exactly. This is why Goldstein suggests that constraint forces are only known in terms of their effect on the motion of the system, and as a result provide difficulties for mathematic descriptions of the motions of systems of particles.

Not only does the vectorial approach make problems that invoke constraints difficult to solve, it relies on a notion of force that is inconsistent with the Newtonian approach. Regarding the “velocity sensitive” constraint forces that hold a bead on a wire, Wilson writes

[C]ustomary point mass interpretations of Newton’s third law forbid forces of this character . . . Forces of a permissible type can hold \mathbf{b} near to \mathbf{W} , but only at the cost of some complex wobbling [Wilson2012, 25].

In addition, Wilson points out that such point-mass approaches to constraint forces don’t take into account the friction that occurs between the mass and the constraint, e.g. the bead and the wire. In light of such worries about constraints, the point-mass setting becomes increasingly unsatisfactory.

To surmount the . . . difficulty . . . that the forces of constraint are unknown a priori, we should like to so formulate the mechanics that the forces of constraint disappear. We need then deal only with the known applied forces [Goldstein, Poole, and Safko2002, 16].

The numerous difficulties that constraint forces cause for the point-mass setting calls for a different approach to classical mechanics — an approach that can do away with all of the microscopic constraint forces between individual point masses and instead focus on the general forces which are applied to the system.

Lagrange’s mechanics does precisely this: it allows us to derive predictions based on the external forces alone. However, when we move to Lagrange’s mechanics we abandon a dynamic laws approach in search of a more abstract and general classification of physical phenomena. In the next two sections, I will spell out precisely how this works.

3.1.1 Generalized Coordinates and Configuration Space

Before introducing the principle of virtual work, it is important to consider the mathematical structure in which Lagrange’s theory takes place. Lagrange relies on *generalized coordinates*. Generalized coordinates are an abstraction from the geometrical concept of Cartesian coordinates. On the standard picture of Cartesian coordinates, we can specify a point in space according to its position along three different axes (x, y, z) . The position of a particle on any

given axis can often be described as a function of time.

$$\begin{aligned}x &= f(t) \\y &= g(t) \\z &= h(t)\end{aligned}\tag{3.3}$$

However, sometimes it is useful to describe the trajectory of a point in spherical coordinates (r, θ, ϕ) . We can then express the relationship between the spherical coordinates and the rectangular coordinates by means of a coordinate transformation, which depicts each rectangular coordinate in terms of the spherical coordinates:

$$\begin{aligned}x &= r \sin \theta \cos \phi \\y &= r \sin \theta \sin \phi \\z &= r \cos \theta\end{aligned}\tag{3.4}$$

We can rely on our geometrical intuitions about the relationship between spherical coordinates and rectangular coordinates in order to point us towards the more abstract concept of a generalized coordinate. Generalized coordinates are *any set of coordinates by which we can uniquely specify the state of the system*. They needn't correspond to a particle's physical location, as is the case with Cartesian coordinates and spherical coordinates. The idea of a generalized coordinate just is a coordinate by which we specify our system. We then consider each generalized coordinate (q_n) to be expressible by a function of time so that

$$\begin{aligned}q_1 &= q_1(t) \\q_2 &= q_2(t) \\&\dots \\q_n &= q_n(t)\end{aligned}\tag{3.5}$$

Any parameters which specify our system can then be chosen as the general coordinates. For example, if we are interested in the configuration of a rigid body, we can choose just six general coordinates to specify it (three pertaining to location and three pertaining to rotation), ignoring all of the particles which compose it. Notice that the microstructure is irrelevant with respect to the coordinates we choose to specify the system.

In the same way that a triple of Cartesian coordinates specifies a point in Euclidean space, generalized coordinates specify a point in a generalized configuration space. The dimension of the configuration space corresponds to the number of generalized coordinates used to specify the system. In Lagrange’s mechanics, the number of coordinates needed corresponds to the number of degrees of freedom for a specific system. A degree of freedom is just a way in which the system is free to move (e.g. a system constrained to non-rotational movement on a flat surface has two degrees of freedom; it can move freely about a plane).

The generalized coordinates used to specify a system are represented in a generalized configuration space as a single point (because the dimensions of the configuration space correspond to the number of coordinates). However, it should be evident that the configuration space we refer to here no longer maps directly to physical space; it is a mathematical abstraction used for the purpose of computing the evolution of a system. In describing the generalized configuration space, Cornelius Lanczos writes

[T]he position of a rigid body — with all the infinity of mass points which for it — is symbolized as a single point of a 6-dimensional space. The 6-dimensional space has, to be sure, nothing to do with the physical reality of the rigid body. It is merely *correlated* to the rigid body in the sense of a one-to-one correspondence [Lanczos1964, 13].

3.2 LAGRANGE’S STATICS AND THE PRINCIPAL OF VIRTUAL WORK

Lagrange’s formulation of mechanics begins with statics: the study of the equilibrium conditions for a system. According to the vectorial approach, a system is in equilibrium when each individual point-mass is in equilibrium. That is, a system is in equilibrium when the sum of the forces on each individual point-mass is zero. It’s not difficult to see why this understanding of statics, though intuitively simple, is quite problematic in practice. If one aims to discover the conditions of when some rigid body (e.g. a lever) is in equilibrium, one needs to be able to calculate the forces on *each individual particle in the rod* in order to find the conditions of equilibrium. Lagrange provides a much clearer way of formulating a system’s equations of equilibrium, based on the *principle of virtual work* sometimes called the *principle of virtual velocities*. Lagrange expresses the principle as follows:

If an arbitrary system of any number of bodies or mass points, each acted upon by arbitrary forces, is in equilibrium and if an infinitesimal displacement is given to this system, in which each mass point traverses an infinitesimal distance which expresses its virtual velocity, then the sum of the forces, each multiplied by the distance that the individual mass point traverses in the direction of this force, will always be equal to zero. Furthermore, the small distances traversed in the direction of the forces are considered positive and the distances traversed in the opposite direction are considered negative. [Lagrange, Boissonnade, and Vagliente2010, 23]

The intuitive idea is that a system is in equilibrium when any infinitesimal displacement of the system in a direction that is consistent with its constraints (e.g. moving a lever up or down rather than breaking it in half) will be arrested by the other forces in the system. To grasp some physical idea of what this suggests, consider the fact that if you lightly tap one side of a balanced scale, the scale will return to its originally balanced state because the other forces will restore the system to a state of equilibrium. However, the displacements of interest to Lagrange are not actual displacements but virtual displacements; these displacements are *virtual* because they do not occur in time, a fact I will explain momentarily.

To provide a sketch of the difference between the vectorial and Lagrangian approaches to equilibrium problems, it helps to consider an example. A seesaw exemplifies a typical system considered in the context of constraints: the system of a class one lever and fulcrum. The seesaw is subject to a variety of macroscopic constraints. The board that lies across the fulcrum is rigid. A long, flat plank (the lever) is pinned to the fulcrum at the center so it cannot slide along the horizontal axis. A standard problem in statics asks us to suppose we place a weight, m_1 at some distance l_1 from the fulcrum. If we want to balance the seesaw with a different weight, m_2 , how far from the fulcrum shall we place the weight?

The vectorial approach tells us the system will be in equilibrium when the sums of the forces on *on each and every point mass composing the lever* are equal to zero. In contrast, the method of virtual work defines the equilibrium state for the seesaw according to the macroscopic properties of the system. Specifically, if you apply an infinitesimal displacement (consistent with the constraints of the system) to some part of the lever, work will be done by the other forces to compensate for this displacement bringing the total work to zero. The *principle of virtual work* relies only on the forces applied to the system in order to determine the state of equilibrium, rather than taking into consideration the forces acting on each

individual point mass.

For example, suppose Jack and Jill are sitting on this seesaw and we displace Jack ever so slightly. If the seesaw is in equilibrium, we imagine an equal and opposite displacement of Jill to occur.

If Jack and Jill are *not* in equilibrium, then when some infinitesimal displacement occurs other forces will not compensate and keep the system in equilibrium. Hence, we know that the system is in equilibrium when the sum of the virtual work for any infinitesimal virtual displacement is 0. If we want to find the forces that will keep a system in equilibrium, we look for the forces that, for any virtual displacement, the sum of the virtual work vanishes.

At this point it's important to discuss what makes the displacements of interest "virtual." To explicate this idea, I will rely on the system of the lever discussed in the previous example.

Although I have only considered the approaches to the lever that are provided by vectorial mechanics and Lagrangian mechanics, it is important to note that the question of how to balance a lever long pre-dates either of these theories. The law of the lever was first formulated by Archimedes, who noticed the relationship between the effective force applied to the lever and the distance from the fulcrum at which the force was applied [Dijksterhuis1956]. The law of the lever states that our seesaw is in balance when the product of Jill's weight and distance from the fulcrum is equal to the product of Jack's weight and distance from the fulcrum.

$$m_1gl_1 = m_2gl_2. \tag{3.6}$$

The law of the lever depends on the effective force of turning moments. A moment is a force with which the lever turns about the fulcrum. Moments are calculated by taking the product of the magnitude and direction of the force applied to the lever and the distance from the fulcrum at which the force is applied. In the case of Jill, the force applied to the lever is the force of her weight. However, the magnitude of the force of her weight depends on the state of the seesaw. When she is in equilibrium with Jack, the force of her weight (and the force of Jack's weight) are applied along the y-axis in the direction of gravity. If the seesaw moves out of equilibrium, the direction of the force applied by Jill (and, consequently, the magnitude of the force) will change to reflect Jill's position on the seesaw. This is because the force of gravity is reduced when a mass is on an incline, as opposed to parallel with the

surface of the earth.

Now what does this have to do with virtual work? Remember, we calculate the virtual work by looking for the location on the seesaw Jill needs to sit so that the work done to create any infinitesimal displacement of her will be compensated for by the work done by Jack's weight, bringing the sum of the effective forces on the system to zero. Now if we *actually* displace Jill ever so slightly, causing the lever to be at an incline, part of the force of Jill's weight will be used to keep her pressed against the lever, and the full gravitational force of Jill's weight will no longer be applied to keep Jack in balance². But we are interested in the point on the seesaw Jill ought to sit at when the full force of her weight is applied in the direction of gravity. Therefore, we stipulate that the displacement is *virtual* because we ignore the loss of force that would be caused by an actual displacement.

In some ways, virtual displacements are similar to extremely isolated experiments — or the behavior of a system in a vacuum. However, they go beyond physical isolation, and instead mathematically isolate individual variables — allowing us to change them without changing other variables which depend on them. Although we can mathematically create these displacements, such changes in our system could never occur in time simply because of the complexities introduced by the changes in surrounding forces.

3.2.1 The Principle of Virtual Work

To introduce a rigorous formulation of the principle of virtual work, it is helpful to introduce the language of vectorial mechanics and assume that the external forces F_1, F_2, \dots, F_n act on points of the system specified by P_1, P_2, \dots, P_n . The virtual displacement of these forces will be denoted by

$$\delta R_1, \delta R_2, \dots, \delta R_n \tag{3.7}$$

.

R_1 is the displacement caused by the application of F_1 to P_1 . These displacements must be consistent with the constraints of the system. The constraints are factored into the equations of virtual work due to our choice of coordinates. We are able to reduce the number

²I am forever indebted for Mark Wilson for elucidating these principles to me on multiple occasions.

of coordinates we need to represent our system according to the ways our system is free to move in terms of the given constraints.

The principle of virtual work states “*the given mechanical system will be in equilibrium if, and only if, the total virtual work of all the impressed forces vanishes* [Lanczos1964, 75]:

$$\delta W = \sum F_n \cdot \delta R_n = 0. \quad (3.8)$$

Translating this into the language of generalized coordinates, it becomes

$$\delta W = \sum F_n \cdot \delta q_n = 0. \quad (3.9)$$

The general idea is similar to the suggestion from the point-mass setting that a particle is in equilibrium whenever the forces acting on that particle vanish. However, it generalizes this equilibrium idea for a single particle to a system of particles. The system is in equilibrium whenever the forces impressed on the system are equal to zero. Since we know, for example, that the particles in a rigid rod are already in some kind of internal equilibrium — that is, we expect the rod to remain rigid and not suddenly break in half — we ignore the smaller scale behavior of the particles which compose the rod. We can ignore all of the interactions between the individual point masses that compose the system.

As Lanczos writes,

We now come to the physical interpretation of the principle of virtual work. According to Newtonian mechanics, the state of equilibrium requires that the resultant force acting on *any* particle of the system shall vanish. This resultant force is the sum of the impressed force and the forces which maintain the given constraints. These latter forces are usually called the “forces of reaction.” Since the principle of equilibrium requires that “impressed force plus resultant force of reaction equals zero,” we see that the virtual work of the impressed forces can be replaced by the negative virtual work of the forces of reaction [Lanczos1964, 76].

What Lanczos calls the “forces of reaction” are what we have called the forces of constraint. The principle of virtual work tells us that whatever force we impress upon a system will bring that system into equilibrium when it is equal to and consistent with the forces caused by the constraints on the system. Reconsider the issue that Wilson suggests with the bead on a wire. The point-mass formulation of the bead on a wire requires us to posit

theoretically impossible velocity-sensitive forces that adjust depending on the velocity of the bead³. Here, however, the forces of reaction are built into the equilibrium conditions as the “forces of reaction.” They are not taken as determinate forces, but as the forces which balance out the forces applied to the system.

To summarize, Lagrange supplies a treatment for static systems that resolves the problems constraint equations pose for vectorial mechanics. Lagrangian mechanics is formulated in a generalized configuration space, which allows an entire system to be treated as a single point. The foundation of Lagrange’s treatment of static problems is the principle of virtual work, which states that a system is in equilibrium when the sum of the work done by the effective forces on the system is equal to zero. However, this work is only virtual because the displacements are virtual displacements; any actual displacement to the system would create a loss of effective force, making the calculations of the equilibrium conditions significantly more complex.

3.3 A CLASSIFICATION OF STATICS

At this point, I want to take a slight detour and discuss the historical significance of Lagrange’s statics, before moving on to Lagrangian dynamics.

Lagrange did not invent the principle of virtual work — it was originally formulated by Bernoulli⁴. Lagrange did not discover the law of the lever — it was uncovered thousands of years ago by Archimedes. Lagrange did not postulate that a point is in equilibrium when the sum of its effective forces are equal to zero; this had been rigorously formulated by Newton. What makes Lagrange’s work so important is his ability to create a unified treatment of statics: he used the principle of virtual work to unite the previously disparate discoveries in equilibrium problems such as the law of the lever and the composition of forces by classifying them as different instances of a single principle.

Scientists already possessed experimental facts about how particular systems are kept in

³The reason why they are theoretically impossible is that such forces violate Newton’s Third Law. For a clear explication of this problem see [Wilson2012].

⁴For a detailed discussion of the evolution of the Principle of Virtual Work, see [Capecchi2012].

equilibrium. They knew the experimental truth of the law of the lever, they knew that a system was in equilibrium when the sum of effective forces applied to the system was zero, etc. Additionally, there were a great number of facts about fluids in equilibrium that applied to specific conditions but could not be unified.

In *Analytic Mechanics*, Lagrange is able to generalize the principle of work to continuous bodies simply by exchanging the summation of the virtual work for an integral. Concerning Lagrange's work in fluid mechanics, Duhem writes

From this the laws of equilibrium of filaments and flexible membranes take a singular clarity and generality; but above all it is the study of the equilibrium of liquids which tests the breadth and penetration of Lagrange's mechanics [Duhem1980, 30].

For example, scientists had long noticed that when a fluid is at rest, the force it exerts at any point in any direction has a single magnitude. Pascal was the first to attempt to bring some rigorous definition to this notion of hydrostatic pressure. Lagrange not only gave a rigorous treatment of the laws of hydrostatic pressure⁵, he also managed to unify the treatment of hydrostatic pressure with the treatment of other problems concerning equilibria.

Hydrostatic pressure is the pressure on a particular section of a fluid when all of its parts have ceased to flow. Therefore, if you take any cubic section in the fluid, that particular section has also ceased to flow. We can think of the hydrostatic pressure as the force that keeps this cube in place. The force is applied to all sides of the cube, as the cube does not flow in any direction. While mechanicians were aware of the laws of fluid flow for some time, a rigorous formulation of these laws eluded them prior to *Analytical Mechanics*.

...everything was not clear and rigorous in the theory of the equilibrium of fluids; the nature of hydrostatic pressure remained quite obscure; it was admitted that this pressure existed, that it was always normal to the surface element to which it referred, that its magnitude did not vary when this surface element turned about one of its points; but of these propositions one had not proof, and even no precise definition of the pressure [Duhem1980, 31].

As Duhem points out, although mechanicians recognized the existence of hydrostatic pressure, and were aware of its properties, they were unable to unify this knowledge with their knowledge of, say, the law of the lever. Lagrange's amazing feat was to bring these seemingly

⁵It turns out Lagrange's work could not actually handle hydrostatic pressure, but Duhem (reasonably) thought it did. The confusion has to do with the proper understanding of infinitesimals.

distinct areas of experimental knowledge together under a unified, rigorous mathematical principle.

Here we see the beginnings of Duhem’s distinction between an explanation and a natural classification. Lagrange’s mechanics do not provide us with some “behind the veil” description of the behavior of bodies “in themselves.” Nor do they provide us with information about the “essence” of bodies or the “essential structure” of a fluid or a rod. In fact, it intentionally disregards the microstructure of the systems in order to focus instead on the role played by effective force. However, this ignorance allows Lagrange to unify and classify multiple areas of scientific research beneath a single principle.

Not only did Lagrange introduce a way to classify disparate experimental laws, he also provided a unified method for the treatment of equilibrium problems. Lagrange describes this aspect of his work in *Analytical Mechanics*.

The authors who have written on the Principle of Virtual Velocities in the past have concentrated on proving the veracity of this principle by demonstrating the congruity between solutions obtained using this principle with those obtained from the ordinary principles of statics rather than to demonstrate its application to solve directly the problems of this science. We propose to fill this latter task with all possible generality and to deduce from this principle analytical formulas which contain the solution of all the problems of the equilibrium of bodies [Lagrange, Boissonnade, and Vagliente2010, 60].

Lagrange was not satisfied with showing that the principle of virtual work was often equivalent to alternative methods used to derive the equilibrium conditions of varied systems; instead he develops a method for deriving the equations of equilibrium for any system directly from the general principle of virtual work. He refers to this technique as the “Method of Multipliers,” and these multipliers have since become known as *Lagrange Multipliers*.

Lagrange’s method of multipliers relies on the basic principles of calculus of variations — a branch of mathematics to which Lagrange was a significant contributor. The calculus of variations deals with minimizing and maximizing specific functions and their constraints. For example, suppose we want to find the lowest point of a well. What do we know about this point? Imagine two marbles: one is sitting at the lowest point in the well, and the other positioned halfway down one side of the well. One major difference between these two marbles pertains to the regions surrounding their location. If we slightly displace the marble sliding down one side of the well in a direction consistent with the constraints of motion,

we know the marble will either move to a higher location or a lower location. By contrast, consider the *tiniest* displacement of the marble resting at the lowest point. If we think of the well as being “continuous,” then the height of the marble will not increase or decrease. The bottom of the well is flat.

Translating this into the language of mathematics, the well is a well-defined, continuous curve such as $f(x) = x^2$. Think of our lowest-point-marble. We discovered that when our marble is infinitesimally displaced, the rate of change (with respect to height) that it experiences is zero. The same would be true, say, if our marble were at the top of a mountain. Therefore, if we want to find the minima or maxima of our function, we should look for the points where the rate of change of our function is equal to zero. To do this we set derivative of our function, $f'(x) = 2x$, to zero, and solve. It turns out that the only stationary point of our function is (0,0). Next, we need to determine if this point is actually a maxima, a minima, or a *saddle point*. A saddle point is a point where the maxima of one curve intersects with the minima of another curve, similar to what is displayed at the center of a saddle. In order to test and see if a stationary point is a saddle point, we take the second derivative of the function and run a secondary test. This second test tells us about the rate of change of the rate of change. If it is consistently increasing or consistently decreasing we have discovered a true extremum point. If the second derivative increases in some directions and decreases in others, we have discovered a saddle point. However, in problems of mechanics, rarely does it matter if a point is an actual minima or maxima:⁶ our primary interest is in stationary states.

Let’s reconsider these “tiny” displacements of our marble, referred to as “explorations of the infinitesimal neighborhood of a point” [Lanczos1964, 38]. Such explorations are called *variations*, and the principles that depend on them are *variational principles* (e.g. the principle of virtual work). The variations are *infinitesimal* displacements — displacements that are so small they cannot be measured, but are not equal to zero. One example of an infinitesimal displacement is the variation of Jill on the seesaw. Another common use of infinitesimals in calculus are the *ds* used to denote the width of the tiny rectangles that divide the area under a curve. However, the infinitesimals represented by the d-process in

⁶Except when looking for the stability of equilibria.

calculus are thought of as actual divisions of the area, while the displacements considered in variational principles are *virtual* displacements — mathematical experiments. We can hypothesize such displacements according to the laws of mathematics, but they are often physically impossible. Lagrange focused on the virtual character of these displacements, and was the first to create a unique symbol to denote the process of variation: the symbol δ .

Now we can be a bit more rigorous in defining Lagrange’s method of multipliers ⁷. Suppose we begin with a continuous, differentiable function of arbitrary variables

$$F = F(u_1, u_2, \dots, u_n) \tag{3.10}$$

and apply infinitesimal displacements to each coordinate of F. The transformed coordinates become

$$\delta u_1, \delta u_2, \dots, \delta u_n \tag{3.11}$$

and the change in the function F is

$$\delta F = \frac{\partial F}{\partial u_1} \delta u_1 + \frac{\partial F}{\partial u_2} \delta u_2 + \dots + \frac{\partial F}{\partial u_n} \delta u_n. \tag{3.12}$$

Equation (12) is known as the “first variation” of our function. In order to operate with finite quantities, we can introduce ϵ as a parameter that approaches zero and then let

$$\delta u_1 = \epsilon a_1, \delta u_2 = \epsilon a_2, \dots, \delta u_n = \epsilon a_n \tag{3.13}$$

where a_n corresponds to the direction cosine of the virtual direction each parameter was displaced in. Equation (12) for the specified direction now becomes

$$\frac{\delta F}{\epsilon} = \frac{\partial F}{\partial u_1} a_1 + \frac{\partial F}{\partial u_2} a_2 + \dots + \frac{\partial F}{\partial u_n} a_n. \tag{3.14}$$

If F has a stationary value, then the sum of the first variations for any arbitrary displacement must vanish so that

$$\sum_{k=1}^n \frac{\partial F}{\partial u_k} a_k = 0. \tag{3.15}$$

⁷The derivation of the multiplier method follows Lanczos [[Lanczos1964](#), 39-42].

Here a_k is an index of all the possible directions in which we might evoke a virtual displacement. However, because the displacements are virtual, they are arbitrary, and we can invoke them in whichever direction we wish. Therefore

$$\frac{\partial F}{\partial u_k} = 0, (k = 1, 2, \dots, n). \quad (3.16)$$

To summarize, if a function F of n variables has a stationary value at some point, P , it follows that the partial derivatives of F with respect to each of the n variables will vanish at P .

The variation that we have just discussed is called a “free variation,” because it describes how a system behaves without constraints. But suppose we evaluate the variation of the function

$$F = F(u_1, u_2, \dots, u_n) \quad (3.17)$$

along with some constraint

$$f(u_1, u_2, \dots, u_n) = 0 \quad (3.18)$$

The intuitive idea is that at least one of our u_k is constrained in such a way that it cannot be displaced. Although we can handle this by trying to get rid of the constrained variable, that is, rewrite it in terms of the other variables, that procedure can become difficult when multiple constraints are in place. This is where Lagrange’s method of multipliers shows its power. Suppose that we apply a variation to the constraint equation (18) in order to obtain

$$\delta f = \frac{\partial f}{\partial u_1} \delta u_1 + \frac{\partial f}{\partial u_2} \delta u_2 + \dots + \frac{\partial f}{\partial u_n} \delta u_n = 0. \quad (3.19)$$

We can simultaneously suppose that F has a stationary value where

$$\delta F = \frac{\partial F}{\partial u_1} \delta u_1 + \frac{\partial F}{\partial u_2} \delta u_2 + \dots + \frac{\partial F}{\partial u_n} \delta u_n = 0. \quad (3.20)$$

We know that in the case where the variables in F are independent, $\frac{\partial F}{\partial u_k}$ vanishes; however, the constraint equation makes the variables dependent on one another. Lagrange’s idea was to introduce a multiplier of some undetermined factor, λ , to δf and then add it to δF .

Because we have set $\delta f = 0$ in (19), multiplying δf by λ and adding then adding it to δF will not change the value of δF .

$$\frac{\partial F}{\partial u_1} \delta u_1 + \frac{\partial F}{\partial u_2} \delta u_2 + \dots + \frac{\partial F}{\partial u_n} \delta u_n + \lambda \left(\frac{\partial f}{\partial u_1} \delta u_1 + \frac{\partial f}{\partial u_2} \delta u_2 + \dots + \frac{\partial f}{\partial u_n} \delta u_n \right) = 0 \quad (3.21)$$

On the one hand, we have added 0 to 0 — a rather trivial move. On the other hand, the move is not at all trivial because we have added a sum. Just because the sum is equal to zero does not mean that each individual term equals zero. In fact — we know there is some term that does not equal zero, the term we have constrained.

We can then rewrite (21) in the form

$$\sum_{k=1}^n \left(\frac{\partial F}{\partial u_k} + \lambda \frac{\partial f}{\partial u_k} \right) \delta u_k = 0 \quad (3.22)$$

Remember, there is some term we want to eliminate: the term that is constrained. Let's name this term u_n . We can choose λ to be the factor such that it causes our problematic term to vanish. We just solve for λ so that

$$\frac{\partial F}{\partial u_n} + \lambda \frac{\partial f}{\partial u_n} = 0 \quad (3.23)$$

We can now treat our problem as a free variation with one less term. For all of our other terms, we know that the coefficient of each δu_k will vanish, so multiplying it by λ will not effect the overall sum.

I have here shown how Lagrange multipliers work for a problem with a single constraint, but the method can be generalized to handle an arbitrary number of constraints. This displays the further power of Lagrange's mechanics. Not only did Lagrange develop a principle which unified and classified experimental laws, he developed a systematic way of solving problems that involve variational principles. By showing that the equations for any equilibrium problem can be derived from a variational principle, he managed to unify in both principle and practice the whole of statics.

It is helpful to pause for a moment and consider just how different the structure of Lagrange's mechanics is from the dynamic laws picture and the vectorial approach to equilibrium. Lagrange provides us with a general principle from which we can derive some

equations that need to be solved, along with an algebraic method for handling the constraint equations. What he does not tell us is anything about the nature of matter itself. He relies on generalized coordinates and a generalized notion of force — none of which need to correspond to the “essential properties” of matter. In contrast, the point-mass formulation suggests we ought to understand equilibrium from the bottom up: a system is in equilibrium when each individual point mass is in equilibrium. There is no talk of a general, unifying principle; instead the point-mass setting focuses on “fundamental bits of matter” and the “laws” that apply to them.

Lagrange does not give us a dynamic law of nature so much as he provides a principle. A principle is a general, fundamental truth from which an experimental law can be derived. These principles allow us to unify our theory, and even tell us something about how the world is classified, but they don’t reveal the sort of information appropriate for admission into an ideal DN text.

The difference in structure between the use of variational principles and the point mass setting will continue to crystallize as we turn to consider Lagrange’s approach to dynamics.

3.4 LAGRANGE’S DYNAMICS AND D’ALEMBERT’S PRINCIPLE

While statics is the treatment of systems in equilibrium, dynamics is the study of objects in motion. One of the most famous laws of classical dynamics is Newton’s second law of motion

$$F = ma. \tag{3.24}$$

The mathematician and philosopher d’Alembert uncovered a way to unite this equation and the principle of virtual work. The first step in d’Alembert’s reasoning was to point out that Newton’s second law is equivalent to

$$F - ma = 0. \tag{3.25}$$

Next he substituted ma with something he called the *force of inertia*, I . Given that we define I as a force vector equal to $-ma$, it follows that

$$F + I = 0 \tag{3.26}$$

While this appears to be a trivial variable substitution of Newton’s second law, what makes d’Alembert’s discovery a work of genius is that it allows us to treat dynamic problems as a special class of equilibrium problems. If we treat $-ma$ as an effective force, even a system in motion maintains a type of “equilibrium” — the sum of the work done by the “effective forces” is equal to zero. Concerning this discovery, Lanczos writes

Apparently nothing is gained, since the intermediate step gives merely a new name to the negative product of mass times acceleration. It is exactly this apparent triviality which makes d’Alembert’s principle such an ingenious invention and at the same time so open to distortion and misunderstanding.

The importance of the equation [3.26] lies in the fact that it is *more* than a reformulation of Newton’s equation. It is the expression of a *principle*. We know that the vanishing of a force in Newtonian mechanics means equilibrium. Hence, the equation [3.26] says that the addition of the force of inertia to the other acting forces produces equilibrium. But this means that if we have any criterion for the equilibrium of a mechanical system, we can immediately extend that criterion to a system which is in motion. All we have to do is to add the new “force of inertia” to the previous forces. By this device *dynamics is reduced to statics* [Lanczos1964, 89].

As Lanczos points out, by treating $-ma$ as an effective force, the general principle behind the equations of dynamics now bears a strong resemblance to the principle behind the equations of statics. We are able to extend the criteria we have developed for equilibria problems to the problems of motion. To determine the path of motion that a particle is going to follow, find the path where the sum of the work done by the “effective forces” on the particle (including the force of inertia, I) is equal to zero. Moreover, because Lagrange had developed a rigorous, unified, algebraic treatment of equilibria problems by means of the method of multipliers, these techniques could further be used on the problems of dynamics.

D’Alembert’s insight that $-ma$ can be treated as an effective force, and that even the problems of dynamics can be treated as problems of statics, led Lagrange to develop the following formulation of d’Alembert’s principle:

$$\sum (F_i - m_i a_i) \cdot \delta r_i = 0 \tag{3.27}$$

When we compare this equation to equation (8), it is clear that d’Alembert’s principle just is a special application of the *principle of virtual work* in the case of dynamics. Instead of summing only the effective forces, we also subtract from the forces ma , or, in d’Alembert’s terms, we add the *inertial force* $-ma$. Because d’Alembert showed that even systems in motion can be thought of as instantiating a special kind of equilibria, the rigorous analysis Lagrange developed for the treatment of statics can be extended to dynamics.

The “distortion and misunderstanding” to which Lanczos refers has to do with the way d’Alembert’s principle is used to generate the equations of motion. First, it should be clear that d’Alembert’s principle isn’t used to *solve* the equations of motion, but to provide differential equations which also must be solved to discover the actual motion of a system. D’Alembert’s principle is not itself the equation of motion for a system, but merely a principle by which we can generate the equations of motion for a system. D’Alembert’s principle provides us with a way to use the general principle of statics to derive the equations of motion for a variety of systems.

This does not mean we can actually *solve* a dynamical problem by statical methods. The resulting equations are *differential equations* which have to be solved. We have merely *deduced* these differential equations by statical considerations. The addition of the force of inertia I to the acting force F changes the problem of motion to a problem of equilibrium [Lanczos1964, 89].

One application of d’Alembert’s principle that Lanczos explores is the conservation of energy. The conservation of energy can be derived as a special consequence of d’Alembert’s law. To do this, Lanczos begins by expressing d’Alembert’s principle as

$$\sum_{k=1}^N (F_k - m_k A_k) \cdot \delta R_k = 0 \quad (3.28)$$

where A is the acceleration and R is the displacement. If we assume that the effective forces can be derived from the potential energy function, then the work of these forces will be equal to the negative variation of the potential energy, so that we can rewrite d’Alembert’s principle as

$$\delta V + \sum m_k A_k \cdot \delta R_k = 0 \quad (3.29)$$

Because we can treat the work of the effective forces as variations in the potential energy, the only work which remains to be considered is the work done by the inertial force. Therefore, if

we combine the virtual work done by the potential energy with the work done by the inertial force (as is the case in equation (3.29)), the sum of the total virtual work is zero.

Remember δ denotes an arbitrary virtual displacement of some quantity which is consistent with the system's constraints. Now, we can consider a *special case* of equation (3.29) where we let these *arbitrary* displacements be equal to actual displacements occurring during the interval dt so that

$$\delta R_k = dR_k. \quad (3.30)$$

Because we are now considering a special application of equation (3.29) where the virtual displacements coincide with actual displacements over an infinitesimal time interval, it follows that the virtual change δV of the potential energy will also coincide with the actual displacement dV that occurs during dt . We can then replace the acceleration A_k with the second derivative of R_k , so that

$$\sum m_k \ddot{R}_k \cdot dR_k = \sum m_k \ddot{R}_k \cdot \dot{R}_k dt = \frac{d}{dt} \left(\frac{1}{2} \sum m_k \dot{R}_k^2 \right) dt = dT \quad (3.31)$$

where

$$\frac{1}{2} \sum m_k \dot{R}_k^2 = \frac{1}{2} \sum m_k v_k^2 = T. \quad (3.32)$$

T is the kinetic energy of the system. Therefore, in this special case where we set the virtual displacements to the actual displacements, equation (3.29) becomes

$$dV + dT = d(V + T) = 0 \quad (3.33)$$

which when integrated yields the equation

$$T + V = \text{constant} = E. \quad (3.34)$$

As this example reveals, d'Alembert's principle is not a method for solving the equations of mechanics but instead for deriving them. Special cases of the principle can help yield the basic laws of mechanics such as the conservation of mechanical energy. D'Alembert's principle is a *principle* — it is certainly not the kind of *metaphysical law* which can be fit into a DN explanation.

⁸This derivation follows Lanczos pp. 94-96.

Recall that dynamic laws strategies expect the most basic laws in a theory to be the laws which describe how individual objects evolve over time. But Lagrange describes the laws of mechanics in a completely different form. Rather than stipulate dynamic laws for the individual particles in our theory, analytical mechanics looks at the entire path traveled and determines which way a particle will go. This way of describing a system's trajectory is not based in dynamic laws, but instead, these dynamic laws can be derived from general, classificatory principles.

Lagrange's mechanics allows us to classify numerous equations of mechanics that, prior to classification by a principle, appear distinct. For example, not only the conservation of energy but also the action of apparent forces (such as rotational forces) and Gauss's principle of least constraint can be shown to be consequences of d'Alembert's principle.

As Duhem writes

The invention of this general principle, devoted to putting into equations all the problems of Dynamics, was the object of long and powerful efforts, of which Lagrange has retraced for us the history; these efforts resulted in the discovery of d'Alembert's principle . . .

D'Alembert's Principle reduced the reduction of equations of any problem of Dynamics to the reduction to equations of any problem in Statics . . .

All the essential ideas introduced by Lagrange in the study of Statics were thus taken over into the study of Dynamics, and their fertility was thereby increased immensely [Duhem1980, 32-33].

Duhem highly praised the Lagrange's use of d'Alembert's principle; he saw it as a primary example of the success of science. Many students of science, however, find it needlessly confusing to claim that we can treat bodies in motion as if they were in equilibrium. Lanczos points out that our ability to use the methods of statics for the the treatment of dynamic systems is intrinsically linked to the notion of a *virtual* displacement.

The criterion for the equilibrium of an arbitrary system of forces is that the total virtual work of all forces vanishes. This criterion involves *virtual*, not actual displacements, and is thus equally applicable to masses at rest and to masses in motion. Since the virtual displacement involves a possible, but purely mathematical experiment, it can be applied *at a certain definite time* (even if such a displacement would involve physically infinite velocities). At that instant the actual motion of the body does not enter into account [Lanczos1964, 89-90].

Virtual displacements do not (and could not) occur in time. Because the idea of a virtual displacement involves a static framing of the system, the actual motion of the body and the

dynamical change of the effective forces do not interfere with the virtual displacements. These displacements happen outside of time, so to speak, so the actual motion of the body is irrelevant. Instead the inertial force is considered to be the resistance of the body to the effective forces applied to it.

3.5 THE ROLE OF LAGRANGE

Another important aspect of Lagrange’s mechanics is its algebraic nature. As Lagrange points out in his introduction to *Analytical Mechanics*:

The reader will find no figures in the work. The methods which I set forth do not require either constructions or geometrical or mechanical reasonings: but only algebraic operations, subject to a regular and uniform rule of procedure.

In *Analytical Mechanics*. Lagrange was able to combine contributions to mechanics by his predecessors and contemporaries, such as Bernoulli, Euler, Newton and d’Alembert, into a unified, systematic, algebraic way of solving the majority of problems in mechanics. Duhem quotes Fourier’s praises of Lagrange for being able to achieve such simplicity and unity in his work.

“Lagrange,” said Fourier, “was born to invent and enlarge all the sciences of calculation . . . The distinctive feature of his genius consisted in the unity and breadth of his vision. In everything he applied himself to simple thought, correct and very deep. His principal work, *Mecanique analytique*, could be called Mecanique philosophique, for it reduced all the laws of equilibrium and motion to a single principle; and what is no less admirable, he submitted to them a single method of calculation of which he, himself, was the inventor.” [Duhem1980, 23]

Lagrange’s formulation of mechanics allows a deep unification of the discipline, and is in many ways more fundamental to classical mechanics than the point-mass approach. However, it fails to reveal any special information about the nature of matter or the microlevel behavior of atoms. Instead it arranges scientific theory from the top down — using a principle to derive and classify the various experimental laws of mechanics.

To summarize, d’Alembert’s insight was that the quantity, $-ma$, can be treated as an effective force, and by doing so one could transform the problems of dynamics into the problems

of statics. Lagrange managed to unify the treatments of statics and dynamics by extending the principal of virtual work to dynamic systems with his formulation of d'Alembert's principle. Because virtual displacements are time-independent, the actual motion of the body does not effect their applicability. His work is unique because it brought together seemingly disparate areas of research into a unified, rigorous, algebraic treatment.

3.6 THE NEW MECHANICS: ENERGETICS

In the previous section, I have introduced the reader to the content of Lagrange's mechanics and some aspects of its historical development. As I mentioned previously, much of Duhem's work (both philosophical and scientific) was influenced by his perception of the significance of Lagrange's treatise. In this section, I will explicate Duhem's supposition that a final, unified scientific theory, in the form of *Energetics*, would in many ways mimic the structure of Lagrange's theory of classical mechanics.

3.6.1 Virtual Modifications

At the heart of Lagrange's unified approach to mechanics lies the principle of virtual velocities, now known as the principle of virtual work. The intuitive idea, as expressed previously, is that a system is in equilibrium when the virtual work vanishes. That is, whenever infinitesimal modifications to the system — which are consistent with the system's constraints — balance each other out, the system is in equilibrium.

Duhem suspects that Lagrange's work in *Analytical Mechanics* only begins to tap the potential held by the principle of virtual work. The heart of Duhem's project of energetics, which he refers to in *Evolution of Mechanics* as “the New Mechanics,” is the concept of a virtual modification — a generalized conception of virtual work. A virtual modification is an *unrealizable*, infinitesimal displacement of some quantity in a system. One of the most precise formulations Duhem gives of how a virtual modification works is provided in his *Commentary on the Principles of Thermodynamics*.

He begins by stipulating that a system is composed of two different collections of magnitudes. One set of magnitudes, which he refers to as the “nature” of the system, and denotes with capital sentence letters ‘A, B, . . . L’, are magnitudes which do not vary as the system varies. These are to represent the unchanged properties of a system (e.g. the mass of a system remains fixed while it changes position). In contrast to these fixed magnitudes, we have magnitudes which define what Duhem calls the *state* of a system. The magnitudes used to describe the state of a system are denoted by $\alpha, \beta \dots \lambda$. We can imagine that the magnitudes which represent the state of the system are independent — that is, the value of one does not influence the value of another. Moreover, we can suppose each variable can be increased or decreased continuously.

Regarding the notion of a virtual modification, Duhem writes

Let us therefore imagine a continuous series of different states of the system, that is to say, a continuous series of groups of values of the quantities $\alpha, \beta, \dots \lambda$. Let us successively fix our attention on these various states, in an order which allows continuous passage from one state to another. By designating this a *purely intellectual operation*, we are saying that we impose a *virtual change* on the system [Duhem2011, 41].

One example of a virtual modification is a virtual displacement used in virtual work. In the seesaw example, some magnitudes (e.g. Jill’s weight) are held fixed, while other magnitudes (e.g. Jill’s position) are allowed to “vary in a continuous manner.” In the case of Jill, this variation is an infinitesimal displacement in some direction consistent with constraints. This change is “virtual” because we could not actually change Jill’s position without also changing the force due to her weight.

Duhem considers another example, which involves the ideal gas law

$$PV = nRT, \tag{3.35}$$

where P is the pressure, V is the volume, n is the number of moles of the given gas, R is the ideal gas constant and T is the temperature. Both R and n remain constant for any particular system *by definition*, but the pressure, volume, and temperature can vary. Duhem considers this to be a mathematical property of the system. Our mathematical representation of an ideal gas allows us to vary the volume from 0 to $+\infty$, despite our inability to physically vary

the temperature of any system over such a wide range. These mathematical experiments are what Duhem refers to as “purely intellectual operations.” Duhem goes on to say,

All changes realisable by a system correspond to the variations of the quantities $\alpha, \beta, \dots, \lambda$ compatible with the definitions of these quantities. The sequence of states through which the system passes therefore constitutes a virtual change of the system.

Conversely, can a virtual change always be regarded of the sequence of states that a system traverses during a real change? Remember that the variables $\alpha, \beta, \dots, \lambda$ which can, by their definitions, take arbitrary values, maybe be connected to one another by physical laws, it can be seen that a virtual change maybe be compatible with the definitions of the variables appropriate for representing the various states of the system but conflict with certain physical laws, and, consequently, not be physically realisable [Duhem2011, 41].

Even though our mathematical representation permits us to vary the values of the temperature, volume, and pressure in an unrestricted fashion, this freedom may not fully transfer to physical variations of these quantities due to interference from other physical laws and restrictions. For example, as was previously discussed in the case of virtual work, if we actually displace Jill ever so slightly it will change the quantity of the effective force due to gravity and so the displacements we consider in virtual work cannot be actualized. In the same manner, many virtual modifications are not realizable. Duhem further discusses this in *Evolution of Mechanics*.

J. Willard Gibbs has studied theoretically the dissociation of a perfect gaseous composition into its elements, each regarded as a perfect gas. A formula has been obtained that expresses the law of chemical equilibrium within the body of such a system. I propose to discuss this formula. To this end, leaving unchanged the pressure that supports the gaseous mixture, I consider the absolute temperature that appears in the formula, and I make it vary from 0 to $+\infty$.

If to this mathematical operation one wishes to attribute a physical meaning one is presented with a veritable mob of objections and difficulties. No thermometer can be made to recognise temperatures below a certain limit, none can determine sufficiently high temperatures; this symbol which we have called absolute temperature cannot, by the measuring processes at our disposal, be translated into something that has a concrete meaning, unless its numerical values stays between a certain minimum and a certain maximum only [Duhem1980, 113].

In formulating the law for ideal gasses, the mathematical reasoning used allows us to vary the temperature independently of consideration of other physical constraints. That is, we can represent the temperature as an independent variable, capable of taking on any value between 0 to $+\infty$. Not only is this implausible because of other physical laws that are

not taken into account, as was mentioned previously, but it is difficult for us to even give a physical definition of temperature above or below a certain range. While our mathematical formalism allows us to vary the temperature indiscriminately, it is not clear what this even means in terms of an actual modification. It's true that I can heat up water on the stove and increase its temperature, but what does it mean to increase the temperature of the water to infinity?

However, Duhem does not think these questions are actually problematic for a scientific theory because we assume the modifications are a virtual, mathematical experiment in the same way we imagine the modifications in Lagrange's mechanics to be virtual displacements. What this means for Duhem's take on realism I will discuss in the next chapter. Here I merely want to explain what a virtual modification is, and highlight its importance in Duhem's *New Mechanics*.

As we see in the case of classical mechanics, sometimes these quantities are restricted (mathematically) by constraint conditions such as inequalities. Therefore, even if the temperature is normally allowed to vary, we can add an equation to our representation that fixes the temperature at a constant value, or limits its variations to a small range. As was the case in Lagrange's mechanics, virtual modifications must be consistent with these constraints.

To create a virtual modification is

To impress upon the variable quantities that characterise the state of a system some infinitesimal alterations allowed by the constraints is to impose upon the material system a virtual modification [Duhem1980].

Duhem's generalized notion of a virtual modification is that it is an infinitesimal displacement in one of the variable quantities of a representation of a system that is consistent with the system's mathematical constraints. Again, such a modification can (and almost always does) ignore some *actual* physical constraints such as we saw in the displacement of Jill on the seesaw, but these are not the same as constraint forces. Virtual modifications help us deal with constraint forces, but they are virtual because they ignore physical changes that effect the system when it evolves through time.

3.6.2 The Structure of Energetics

Virtual modifications are at the core of Duhem's reformulation of physical theory. As he writes,

The notion of virtual modification was at the root of the Mechanics of Lagrange as it is at the base of the New Mechanics, but how more general it is in the latter than in the former! The only virtual alterations that were known to the Mechanics of Lagrange were the alterations of shape and position of the various parts of the system; many other alterations are considered by the New Mechanics [Duhem1980, 116].

Lagrange's demonstrated that the whole of statics (and eventually dynamics) can be derived from variational principles, based on virtual modifications. Duhem claims that the same structure applies to the New Mechanics. The theory will be built upon a general variational principles which depend on virtual modifications of the various qualities associated with matter. However, rather than make the foundational principle of the New Mechanics the principle of virtual work, Duhem claims all of the New Mechanics is rooted in the principle of the conservation of energy.

In an explicit comparison of the development of the New Mechanics to the growth of a tree, Duhem claims that trunk from which all of the New Mechanics grows is the principle of the conservation of energy. Different branches of science arise from the trunk of the tree. The first that he considers is the theory of thermodynamics, or what he often refers to as the theory of reversible modifications.

In his discussion of reversible modifications, Duhem considers the work of Gibbs in thermodynamics.

The fundamental principle of the New Statics is therefore presented exactly in the form that Lagrange gave to the principle of the Old Statics; the quantity whose existence has been revealed to us by the axioms of Thermodynamics plays in the former the role that the potential of the internal forces played in the latter; from this there comes the name of Internal Thermodynamical Potential which we have attributed to this quantity [Duhem1980, 135].

The "potential of the internal forces" in Lagrange's statics is what we have previously called the "forces of reaction." Remember, the sum of the internal forces are equal to the sum of the impressed forces for any system in equilibrium. Duhem asserts that the same general idea is true for thermodynamic equilibrium; however, instead of focusing on the sum

of the forces we focus on the sum energy exchanged. The internal energy must be equal to the impressed energy in order for the system to be in equilibrium. Duhem introduces the term “Internal Thermodynamic Potential” to refer to what Gibbs called fundamental functions.

The New Mechanics divides into four different branches: reversible processes, frictional systems, permanent alterations and electric currents. Each of these branches also fork into branches for statics and dynamics. This is Duhem’s picture of energetics, a theory he takes to converge on a natural classification.

Beside the principal trunk of Thermodynamics, besides the Mechanics of systems without friction or hysteresis [systems who converge on reversible modifications], we have seen rise up two other stems, still young and the development of which is a very long way from being accomplished: the Mechanics of frictional systems and the Mechanics of systems with hysteresis. These two stems are not distinguished, first of all, from the principal trunk; up to a certain height they remain knit together, identical with it. All that which precedes the use of this notion, all that appeals to the principle of the Conservation of Energy alone is common to the three mechanics.

Issued from the same roots, a fourth stem arises, born a long time ago and robust already; it treats the Mechanics of the electric currents ... [Duhem1980, 178].

Although Duhem only presents four branches of mechanics, he is open to the idea that someday a new branch might arise. He thinks New Mechanics provides space for this — it is not, as it stands, the ultimate and final physical theory. It merely provides a structure within which science can continue to flourish.

To summarize, Duhem thought the New Mechanics would classify all of physical theory in a structure quite similar to the structure Lagrange provided for classical mechanics. The New Mechanics, often referred to by other scholars as Energetics, would have its roots in the variational principle of the conservation of energy. The different branches of science are all to be seen as the study of different means of energy transformations.

3.7 CONCLUSION

Dynamic laws strategists, such as Railton, Salmon, and Dowe, presuppose that science advances by providing us with laws that describe the time evolutions of objects in a set domain.

Duhem provides us with an interesting counterexample to this approach to theory building in his discussion of Lagrangian mechanics.

Lagrange manages to describe Newtonian laws of motion by appealing to macroscopic rather than microscopic entities. Additionally, the principles at the heart of Lagrange's theory are general mathematical principles rather than dynamic laws. These principles make use of *virtual modifications* — mathematical experiments which cannot be realized in the natural world. Nevertheless, these mathematical experiments allow us to rigorously fix the context of a system in such a way that we can ignore the problematic influence of *friction*.

In the next chapter, I explore the role virtual modifications play in the foundations of thermodynamics, focusing on the discovery of entropy through the Carnot cycle.

4.0 THE CARNOT CYCLE

4.1 A BRIEF OVERVIEW OF THERMODYNAMICS

Thermodynamics is a scientific theory about exchanges of heat and work. Using only a few macroscopic parameters (pressure, temperature, and volume), thermodynamics is able to describe stable behaviors of a wide variety of physical systems *independent* of the microscopic composition of these systems. For example, there are certain relationships between the temperature of a system and its pressure that are quite independent of the atomic make-up of the system — these relationships persist in a variety of different materials which vary in their microstructure.

Despite its deceptive name, Thermodynamics is not a dynamic theory. It does not contain *dynamic laws*, because the only well-defined states in classical thermodynamics are equilibrium states. Equilibrium states are states where the thermodynamic properties of the system are at rest; we consider a system to be in equilibrium when none of its thermodynamic properties are changing. For example, suppose we put a few ice cubes in a glass of water that is at room temperature. The temperature of the water will be in flux while the ice cubes melt: therefore, the system is not in equilibrium. However, temperature of the water will eventually cease changing, the pressure and the volume will also remain the same, and the system will be in a state of equilibrium.

In thermodynamics, the system is not well-defined while it is trying to reach equilibrium (e.g. when the ice is melting), but can only be properly defined when these macroscopic parameters have a definite value. For this reason, the configuration space of thermodynamics contains only equilibrium states. This raises a problem for certain conceptions of scientific theory, such as the one held by dynamic laws strategists. For example, as you may recall,

causal process theorists argue that “processes” are basic to scientific theory. These processes are the continuous paths of systems. However, thermodynamic theory does not tell the story of the continuous evolutions of thermal systems. Instead, it describes *quasi-static processes*.

To understand both the configuration space of thermodynamics and quasi-static processes, it helps to consider an example. Suppose we are interested in how ice melts when we add it to a cup of water. Now, if we add two large ice cubes, it will take some time for the ice cubes to melt. Because we can only define the state of the system when it is in equilibrium, it seems like we will have no information about what happens to the system while the ice melts. But suppose, rather than adding two large ice cubes, we add several small ice cubes one at a time. We can then measure and define the state of the system when each ice cube has melted. Now suppose that the ice cubes are *infinitesimally* small — and after each one melts we measure the thermodynamic parameters of the system. This would create a series of equilibrium states that seem to approximate a continuous process. We can assume that a dense series of these states will approximate the actual process the water goes through when the large ice cubes melt.

Every process described in thermodynamics is such a series of equilibrium states. Of course, this is metaphysically problematic because an equilibrium state is (by definition) a state where the system ceases to change. This means that if the system ever entered into a state of equilibrium it would stay there: it cannot continuously traverse a series of equilibrium states. Moreover, our system “jumps” from one state of equilibrium to another and thermodynamics gives us no information about how it behaves in between these jumps.

One classic textbook describes quasi-static processes as follows:

A quasi-static process is thus defined in terms of a dense succession of *equilibrium* states. It is to be stressed that a quasi-static process therefore is an idealized concept, quite distinct from a real physical process, for a real process always involves nonequilibrium intermediate states having no representation in the thermodynamic configuration space [Callen1985, 96].

This idealization involved in a quasi-static processes is often smoothed over by assuming such processes proceed infinitely slowly — or, in our example, that each ice cube is infinitely small — so that each change in the system is so minimal, it’s as if the system remains in equilibrium throughout the process. Another example of such a process is the isothermal expansion of an ideal gas contained in a cylinder. Suppose such a cylinder is in an infinite heat

bath, so that the temperature of the gas remains fixed (even though, under normal conditions, the temperature decreases when the gas expands). Now, suppose our gas is compressed by a piston that covers the cylinder. In order for this process to be quasi-static, we suppose that increment the piston moves is infinitesimally small and that the gas immediately moves from one equilibrium state to the next. Because such changes are infinitesimally small, it will take an infinite amount of time for the gas to expand a measurable amount. Moreover, when a gas actually expands in the natural world, it surely does not pass through this series of equilibrium states (for if it entered into one of these states its thermodynamic properties would cease changing).

There are several *prima facie* reasons to believe that classical thermodynamics is an impoverished theory. It doesn't tell us anything about the microstructure of thermal systems, it relies on only a few, basic, observable quantities, the only definable states of the system are equilibrium states, and the processes it describes are unrealizable. Nevertheless, thermodynamics manages to describe an impressive number of constraints that effect the behavior of a wide range of distinct systems — it seems to latch on to the world, and actual regularities in the world — in a rather robust way. One of the clearest ways in which thermodynamics provides us with novel, important information about the structure of the natural world is apparent when considering the discovery of entropy.

4.2 THE DISCOVERY OF ENTROPY

While the full history of entropy is too delicate and detailed to flesh out here¹, the basic idea is that entropy is a quantity developed to explain the loss of energy observed in heat engines. In the seventeenth century, several engineers noticed that steam has motive power — it is able to power the mechanical motions of machines. The earliest steam engines were incredibly inefficient, able to do only a little mechanical work for any given amount of heat. James Watt noticed that much of this wastefulness was due to steam lost to condensation. By minimizing the condensation, Watt was able to increase the efficiency

¹For an excellent discussion of the history of entropy see [Uffink2001].

of heat engines significantly [Müller2007, 48]. The trend of building increasingly efficient heat engines continued. Attempts to make machines more efficient were works of engineering genius, but did not reflect an increased understanding of the nature of heat. These discoveries did depend on the discovery of a complex, behind the veil picture of heat transfers. As Ingo Müller writes,

None of the engineers who invented or improved the steam engine or the air engine was in any way distracted by any soul-searching about the nature of heat, or whether or not there was a caloric. They proved that heat could produce work by doing it, and doing it better and better as time went on [Müller2007, 51].

The drive of engineers to produce increasingly efficient heat engines culminated in the work of Sadi Carnot. Carnot wondered about the limits of heat engine efficiency. He decided to approach the problem of engine efficiency from the perspective of mathematics rather than applied engineering. Recognizing that any actual steam engine would lose heat due to frictional forces and complications such as vaporization, he constructed an idealized heat engine, called a Carnot engine, which is not subject to such effects. The cycle the Carnot engine undergoes is called a Carnot cycle.

A Carnot cycle is traversed by an imaginary system (such as a piston and cylinder filled with fluid) that is heated and cooled. The cycle is composed of four stages. In the first stage, the system is placed in a heat bath where the volume of the system increases while the temperature of the substance remains fixed. Eventually the system is removed from the heat bath. The volume of the system continues to expand, only now the temperature of the system decreases because it can no longer absorb heat from the heat bath.

The temperature continues to decrease until it reaches the temperature of a second, cooler, heat-bath known as the refrigerator. When the system reaches this lower temperature, it is placed in the refrigerator. The volume of the system decreases while the temperature is fixed at the cooler temperature. Then the system is removed from the refrigerator and the volume continues to decrease until the system is restored to its initial volume and temperature. Because the system begins and ends in the same state, we call this process a cycle.

Although this description of the Carnot cycle sounds like a dynamical model of a system evolving through time, it is actually represented across a *dense sequence of equilibrium*

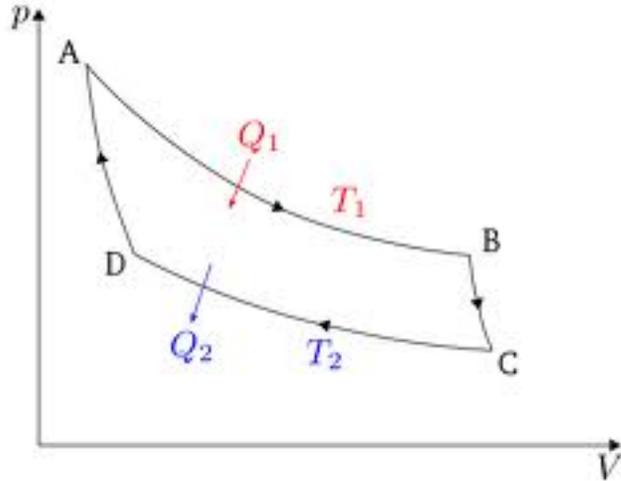


Figure 3: Pressure-volume diagram for a Carnot cycle.

states. Remember, the configuration space used to describe processes such as the Carnot cycle consists of *only* equilibrium states. In effect, Carnot thinks of his machine's shifting conditions in the quasi-static manner that we approached the pendulum: as a series of gently induced pushes from one state to another.

One way to represent the Carnot cycle is on a pressure-volume diagram such as the one in Figure 3². Points such as A and B, and all the points between them, represent *equilibrium* states of the system. These are states at which the pressure and volume will eventually come to rest. Now surely there is a time lapse in between each point on our diagram where the pressure or the volume of the system fluctuates. *These are not represented*. Again we see science employing Duhem's concept of a virtual modification. The Carnot engine can only virtually pass through these states.

The quasi-static processes which compose the Carnot cycle are *reversible* processes. Because the system remains in equilibrium the whole time, we can imagine the same process happening in reverse (taking the final state as the initial state, and the initial state as the final state). Duhem addresses reversible processes in his text *Thermodynamics and Chemistry*.

²Diagram by [Derkleinebauer2008].

Between states 1 and 2 arrange a series of *equilibrium positions* $\alpha, \beta, \gamma, \delta, \dots$ following each other in a continuous manner. The system, placed in any of these positions, would remain there indefinitely. This series of positions of equilibrium cannot then be passed through by the system either in one direction or in the other; *it does not correspond to a realizable transformation of the system.* ... This series of equilibrium states $\alpha, \beta, \gamma, \delta \dots$ which is passed over by no modification of the system is, in some sort, the common boundary of the real transformations that bring the system from state 1 to the state 2 and of the real transformations that bring the system from state 2 to state 1 [Duhem2009, 69-70].

Duhem recognizes that, in representations like the Carnot cycle, the actual process is a deviation from the process described by thermodynamic theory. He thinks of these quasi-static processes as the boundaries between two different kinds of actual processes. For example, consider the expansion and compression of a gas. In practice, the expansion of the gas will not actually be reversible, because it will happen so quickly that the system does not remain in equilibrium the whole time. The same is true for the compression of a gas. However, if we slowed these two processes down infinitely, so that they were quasi-static, they would each approach the same series of equilibrium states, as Duhem describes.

These static states can be partially determined experimentally, much like the case of potential and kinetic energy. We can perform experiments to see how the pressure and volume balance at any particular point of the system, then string these static states together to form a quasi-static path for the system to traverse. This path is importantly devoid of friction, which obscures the conservation of important thermodynamic quantities. *Because* the state space of thermodynamics contains equilibrium states, and represents ideal, limiting processes rather than actual processes, thermodynamic models such as the Carnot cycle ignore the intervening and complicating effects of friction. What's especially interesting to the project at hand is that by reducing the effects of friction in this rather ingenious way, we are able to uncover quantities and natural laws which are otherwise obscured.

Carnot managed to prove a theorem which states that no heat engine can be more efficient than a Carnot engine, however, he *couldn't* quantify just how efficient a Carnot engine actually was. As previously stated, the complete path of the Carnot engine is a cycle. For any thermodynamic system to be a process, it must be the case that all of the state variables — thermodynamic parameters — have the same value at the beginning and the end of the cycle. For example, the temperature, volume and pressure must all be the same.

But to understand the efficiency of the Carnot engine required the introduction of a fourth state variable — one that Carnot suspected was heat.

Because Carnot subscribed to the caloric theory of heat, he thought that heat was a substance (much like a fluid) which entered and exited bodies. His metaphysical presumptions caused him to believe that the quantity of heat was conserved when a mechanical engine traverses a cycle. However, it turns out that every time a Carnot engine completes a cycle it *gains* heat. In fact, this additional heat is precisely what is converted into work, allowing the Carnot engine to effect its environment.

It wasn't until almost thirty years later that Clausius determined the actual efficiency of the Carnot engine. Clausius discovered a hidden state variable of thermodynamic systems — the quantity S (which was later called the entropy). The change in entropy as a Carnot engine travels through its cycle is equal to 0.

$$\Delta S = \frac{dQ}{T}. \quad (4.1)$$

This means that the value for entropy is the same when the Carnot engine begins and ends its cycle — indicating that entropy satisfies the definition of a state variable. What's of special interest to us, is that this ever so important and hidden state variable, entropy, was discovered through a careful analysis of the virtual modifications at the heart of thermodynamic theory. It did not become evident through the production of dynamic laws.

As it turns out, entropy couples together with the quantity of temperature to define the internal energy of a system. Much the same way the total energy of a system can be understood as its kinetic energy plus its potential energy, the change in the internal heat energy of a thermodynamic system is temperature times the change in entropy (TdS). Together, temperature and entropy form a thermodynamic *conjugate variable pair*. In thermodynamics, the changes in energy are understood by infinitesimal modifications within such conjugate variable pairs. In the case of temperature and entropy, a change in the temperature creates a change in entropy in a manner analogous to how a force applied to a point causes a spatial displacement. In order to understand how energy is exchanged between a system and its environment, thermodynamics appeals to changes in conjugate variable pairs. That is, these

exchanges explain why the internal energy increases or decreases, and provide constraints on the ways in which the energy can be transformed.

These exchanges between conjugate variable pairs are structurally similar to the modifications at play in the method of virtual work. Remember, in the context of virtual work, generalized forces cause virtual displacements, and the work is the product of the force times the displacement. In thermodynamics, conjugate variable pairs are composed of two types of properties: intensive and extensive. Intensive properties are such that their value is independent of the system's size. For example, temperature is an intensive property. A glass of water and a swimming pool might have the same temperature, even though they are quite different in volume. By contrast, extensive properties do depend on the size of the system. These are properties like mass. The internal energy of a thermodynamic system can be calculated by multiplying an intensive property like temperature (which plays a role analogous to force) with an extensive property like entropy (which plays a role analogous to displacements).

Not only does entropy form a conjugate variable pair with temperature, it also serves as one of the most foundational quantities in thermodynamic theory — grounding and shaping the earliest thermodynamic laws. The first law of thermodynamics, which is a special case of the law of conservation of energy for thermodynamic systems, was proven by the introduction of entropy. It was only by recognizing the relationship between temperature, entropy, and the total internal energy of a system that Clausius was able to show energy was being conserved. Additionally, the second law of thermodynamics states that the entropy of a system almost always increases. It is this entropy increase that helps us to understand the irreversibility of most thermodynamic processes, as well as explaining the observation that heat always flows from warmer bodies to cooler bodies.

4.3 THERMODYNAMICS AS A NATURAL CLASSIFICATION

Like potential energy, entropy has proven an important, albeit abstract, scientific quantity. For example, we understand how heat is exchanged by looking at how temperature drives

changes in entropy. On a grander, philosophical scale, we understand why certain processes have a direction (e.g. why a system tends towards equilibrium rather than away from it) via the notion of entropy. Of course, we still lack the complete dynamical story of how heat is lost and transferred when a heat engine traverses through a real-world cycle, and our discovery about temperature and entropy does not bring us any closer to an ideal DN text. But this does not undermine the importance of entropy or suggest that we should view explanations which reference entropy as “explanatory sketches” which only hint towards some grander dynamic laws picture. As we begin to understand the delicate relationship between entropy and temperature, we are able to classify experimental data and see “hints” of the structure of the natural world, just as Duhem suggested.

Thermodynamics, like Lagrange’s mechanics, does not rest on dynamic laws but, rather, upon general principles. From these principles, we can derive a large number of stable, general constraints that govern the behavior of thermodynamic systems. But rather than “explain” how these systems work — e.g., provide some initial conditions for the system and the laws which will move it from one state to the next, — we are able to classify different physical phenomena according to the categories revealed by our theory. For example, we learn that the processes in nature which are approximately reversible are those where the system stays close to equilibrium and there is no change in entropy. We can also classify the various states of a system according to the state variables, which include entropy. Moreover, we can even classify different thermodynamic systems according to what kind of heat transfers are possible. We know that systems which are, for example, isothermal (have a fixed temperature) will also have a fixed entropy.

Duhem claims this classification approximates a natural classification because the groups that it picks out are *natural kinds*. That is, isothermal processes are actually a unique set of processes that are structurally similar. On Duhem’s view, if we *did* have access to some true, ideal DN text (hand delivered to us by an angel), this perfect theory would classify isothermal processes together in virtue of their structural similarity. However, because our theory gives us only a classification and not an explanation, we don’t have a handle on why these processes form a natural kind. In some ways, Duhem thinks that to be an isothermal process is a symptom rather than a cause: perhaps there is some fundamental, dynamic story

to be told about how these processes are similar. But that story is a metaphysical story, and not a scientific story, on Duhem's view — moreover, it's clearly not what thermodynamics presents us with.

Nevertheless, Duhem insists that thermodynamics helps us learn more about the structure and relatedness of the natural world. For example, thermodynamics helps us to reliably predict the final state of thermal processes. That is, we can't be sure how the heat is actually distributed through a glass of water while the ice is melting, but we can make reliable predictions about what temperature the water in will be after the ice melts.

Thermodynamic constraints cannot be easily dissolved into a unified mechanical picture. Rather than ignore these macroscopic properties when focusing on, for instance, the behavior of classical systems, we have to find ways to couple together thermal properties and mechanical properties. This is why Duhem emphasizes that we need to include qualities and not just quantities in our scientific theories. However, I will cover this issue in more detail in the next chapter.

4.4 STATISTICAL MECHANICS

In the previous sections I have argued, on behalf of Duhem, that thermodynamics manages to reveal novel insights about the structure of the natural world without providing any microscopic dynamic laws. One might react to this assertion by pointing out thermodynamics is now considered to be grounded in statistical mechanics — a probabilistic theory about dynamic laws which govern micro-constituents of thermodynamic systems.

Statistical mechanics is used to patch together the higher level theory of thermodynamics with mechanical, dynamic laws that reign over the microscopic level. Because statistical mechanics defines rigorous relationships between the macroscopic properties of thermodynamics, such as temperature, with microscopic properties of mechanical systems, such as the mean kinetic energy, the relationship between statistical mechanics and thermodynamics is often put forth as a paradigmatic case of theory reduction.

If thermodynamics can be reduced to statistical mechanics, does thermodynamics actu-

ally provide us with any important insight? Can it stand as a counterexample to dynamic laws strategists, if it has already been reduced to a set of dynamic laws? There are two ways of handling this objection — both of which follow. First I will sketch Duhem’s argument that the mechanical theory of heat is not actually a reduction of thermodynamics. Next, I will argue that even if it were the case that this reduction went through, there is still an important lesson to be learned about dynamic laws strategies.

4.4.1 A Brief Sketch of Statistical Mechanics

Before going much further, I want to provide a brief sketch of how thermodynamic phenomena can allegedly be reduced to dynamic laws. As mentioned before, the central notion of thermodynamics is the idea of equilibrium: all systems tend towards thermodynamic equilibrium when left in isolation. But this is quite the opposite of, for example, Newton’s law of inertia which states that an object (e.g. atom) in motion will remain in motion (unless interrupted).

Boltzmann and Clausius both suggested a similar connection between observable thermodynamic equilibrium and microlevel dynamic laws: the temperature of a substance is equal to the average kinetic energy of the particles which compose it. Therefore, if a substance is in thermal equilibrium, it doesn’t mean that the particles are also at rest, it simply means that they are, on average, not increasing or decreasing in velocity. Similarly, the pressure of a system is understood as the impact of these molecules on the sides of the container.

Determining the precise state of a thermal system from the standpoint of mechanics requires us to know the exact position and velocity of each individual particle³. Because thermodynamic systems contain large numbers of molecules, it is practically impossible to determine the actual mechanical state of the system. Rather than work with predictions about individual states of the system based on the dynamic laws which govern the particles, statistical mechanics focuses on ensembles: probability distributions over all possible states of the system. As each state of the system evolves through time, the ensemble itself also

³I have chosen to address the reduction of statistical mechanics to classical mechanics for reasons of consistency and clarity. Similar problems arise in the case of quantum mechanics.

evolves through time⁴.

The central problem with reducing thermodynamics to statistical mechanics was and continues to be the arrow of time. The second law of thermodynamics asserts that the entropy of a system almost always increases. This law is used to account for our perceived arrow of time — the fact that certain events, such as ice cubes melting in a glass of water, seem to occur in a specified direction: it is never the case that ice cubes develop in a glass of water. Dynamic laws alone appear incapable of representing this directionality; dynamic processes are always (mathematically, at least) reversible. Therefore, how can this directional behavior emerge from purely dynamic laws?

Various solutions have been suggested to this problem, and it is still an active question in contemporary philosophy of physics⁵. One way to resolve this worry is to argue there is a yet to be formulated dynamic law which is time-asymmetric that will be able to account for these worries. Another is to suggest that these apparently irreversible processes can really be reduced to a certain kind of sensitivity to initial conditions. Still others have suggested it's not a problem with the dynamic laws but instead with the state of the world — the world is in a certain state such that the dynamic laws consistently produce these seemingly asymmetric processes, but it is quite possible, were the world in a different state, that they would produce symmetric processes.

4.4.2 A Reduction or an Illustration?

Duhem was familiar with attempts to reduce thermodynamics to statistical mechanics, and addresses mechanical theories of heat in *Evolution of Mechanics*, claiming that said theories illustrate rather than explain.

... all that is logically permissible to affirm is that it is, nevertheless, possible to construct mechanically, at least to define by certain algebraic conditions, some ensembles of bodies whose stationary motions are governed by formulae analogous to the equations of Thermodynamics. To go back to a few words which Boltzmann borrowed from Maxwell, *the Mechanical Theory of Heat does not furnish a mechanical explication of Thermodynamics, it gives only a dynamical illustration of it* [Duhem1980, 67].

⁴While there is an interesting philosophical question concerning how to interpret these ensembles, I have chosen to ignore that question here. For more information see [Sklar1993].

⁵See [Albert2003] or [Wallace2010]

While we might be able to model thermodynamic observations using microlevel terminology, this is quite different from deriving the laws of thermodynamics from basic, classical laws. This is Duhem’s distinction between an illustration and an explanation: an illustration can show how certain regularities might “look” from a dynamic laws point of view, but a true explanation must prove these phenomena are consequences of dynamic laws. Rather than derive thermodynamics from the dynamic laws, statistical mechanics introduced thermodynamic phenomena as constraints on dynamic systems.

Duhem considers Helmholtz’s work on monocyclic systems as an object lesson for how statistical mechanics illustrates rather than explains thermodynamics. This is a reference to Helmholtz’s 1884 paper “Principien der Statik monocyclischer Systeme⁶” where he provides a mechanical model for what was known as the *heat theorem*⁷

$$dS = \frac{dE + PdV}{T} \quad (4.2)$$

Helmholtz provides the following definition of the class of systems he studies:

I understand the term *monocyclic systems* to mean those mechanical systems in whose interior one or more stationary, closed motions are present, but which, when there are several, have velocities that depend upon only one parameter. I further assume that only conservative forces act between the individual bodies that define the system, which consist of relatively fixed constraints, while the external forces that must be added in do not necessarily need to be conservative. I refer to the problems that I will treat as static whenever it is assumed that the variations that the state of the system experiences come about with such slight velocities that the system never deviates noticeably from those states in which it can continually abide under them [von Helmholtz, 1].

To get a clearer picture of Helmholtz’s idea, it’s useful to invoke some of the concepts from Lagrange’s analytic approach to mechanics. Essentially, Helmholtz is interested in systems which have one macroscopic degree of freedom, even if these systems consists of multiple particles. Additionally, the particles which compose the system are also constrained to one dimension. For any specified energy level, each particle has one precise trajectory or cycle it will follow. Hence the name “monocyclic.”

⁶Principles of Static, Monocyclic Systems

⁷For more information on the heat theorem, and Helmholtz’s contributions to thermodynamics in general, see [Gallavotti1999].

Duhem illustrates monocyclic systems by — quite literally — a “toy model” of a top [Duhem1980, 62-63]. Imagine, for a moment, a rapidly spinning top. Although it is rotating, the rotations happen so quickly that the top appears to be in a stationary state. These imperceptible rotations are an example of the motions of the atoms in the system. Now, suppose we slowly rotate the top about some axis (while the spinning continues). This second motion would be observable. This corresponds to the one-dimensional motion of a monocyclic system as a whole.

In a contemporary paper on Helmholtz’s work, Michele Campisi describes monocyclic systems as “one-dimensional conservative systems in a confined potential where there is only one periodic trajectory per energy level” [Campisi2005, 281]. He further uses the example of a particle in a box to illustrate a monocyclic system, but a particle inside a double well potential is not monocyclic because the particle might travel one of two different cycles at each different energy level.

Helmholtz was successful in finding a class of monocyclic systems that could model thermodynamic properties. In fact, he was able to prove the following theorem ⁸ :

Helmholtz Theorem *Let $H(p, q; V) = p^2/2m + \varphi(q; V)$ be the Hamiltonian of a one-dimensional monocyclic system. Let a state be characterized by the set of quantities:*

$E = \text{total energy} = \langle K \rangle + \varphi,$

$T = \text{temperature} = \text{twice the time average of the kinetic energy} = 2 \langle K \rangle_t,$

$V = \text{volume} = \text{the external parameter},$

$P = \text{pressure} = \text{minus the time average of } \frac{\delta\varphi}{\delta V} = \left\langle -\frac{\delta\varphi}{\delta V} \right\rangle_t,$

then the differential

$$\frac{dE + PdV}{T} \tag{4.3}$$

is exact and $S_H(E, V)$, defined as

$$S_H(E, V) = \log 2 \int_{x-(E,V)}^{x+(E,V)} \frac{dx}{h} \sqrt{2m(E - \varphi(x, V))} \tag{4.4}$$

is the generating function, i.e.,

$$dS_H = \frac{dE + PdV}{T} \tag{4.5}$$

where the symbols $x \pm (E, V)$ denote the turning points of the trajectory, i.e., the roots of the equation $E - \varphi(x, V) = 0$.

⁸I have modeled the presentation of this Theorem after Campisi, [Campisi2005, 282].

What Helmholtz is able to prove is that a certain class of mechanical models actually mimics thermodynamic behavior. However, as Duhem points out, this is quite different from deriving the laws of thermodynamics from the laws of classical mechanics. He uses two key aspects of this example, one technical and one philosophical, as arguments against a reductionist picture of the relationship between thermodynamics and statistical mechanics.

The technical point that Duhem makes is that Helmholtz was forced to use an analytic (non-vectorial) approach to classical mechanics in his formulation. As Duhem writes,

Thus, to define the monocyclic systems whose properties are capable of imitating the thermodynamical relations, Helmholtz was obliged to submit them to conditions that expressed certain analytic characteristics of the functions used; it is quite difficult to translate them into mechanical language, and even more difficult to draw from them some precise information about the assumptions that it would be useful to make, touching upon the structure of atoms or the nature of calorific motion. From there it is permissible to ask if this analogy between the laws of monocyclic systems and the equations of Thermodynamics has any foundation in the nature of things [Duhem1980, 64].

Hamiltonian mechanics is an analytic version of mechanics quite similar to Lagrange's approach. Although it focuses on a slightly different quantity, it still makes use of variational principles and does not track the individual motion of each particle. It is also built upon variational principles, much the same as Lagrange's theory. In short, Hamiltonian mechanics is no more a dynamic laws approach to classical mechanics than Lagrangian mechanics is. Duhem's point then is — while Helmholtz is able to find a class of models that represents thermodynamic phenomena, it requires the use of a non-dynamic laws formulation of classical mechanics. If this is the case, we have good reason to doubt that Helmholtz has actually connected thermodynamic phenomena to microscopic, dynamic behavior.

For example, Helmholtz's theorem relies on using Volume as a constraint — and we have already discussed the problems that constraints create for vectorial approaches to dynamics. In fact, on Duhem's picture, thermodynamic laws as a whole interact with vectorial mechanics in just this way: as constraints. But if they are constraints, they are additional information to be solved for which alters the course predicted by the dynamic laws, rather than being derivable from the dynamic laws.

Besides this technical point, Duhem makes a philosophical one. That is, the class of models picked out by Helmholtz's theory is totally arbitrary. Helmholtz theorem does not

even hold for all monocyclic systems, but a special class of them. But what is so special about these monocyclic systems *from a particularly dynamic perspective?*

One can prove the existence of such an integrant factor on the condition of restricting the generality of the monocyclic systems studied; unfortunately, it is difficult to interpret, within the meaning of the Mechanical Theory of Heat, the restrictive conditions to which one must appeal [Duhem1980, 63].

Duhem points out that the mechanical theory of heat does not *classify* systems in such a way that these apparently thermodynamic models of monocyclic systems form any kind of a unified class. It is only the classification that is provided by thermodynamic theory that suggests we should be interested in models which behave in this way. In this sense, our knowledge of macroscopic phenomena informs the way we pick out a relevant class of solutions to the equations of classical mechanics. But if this is the case, it wasn't the dynamic laws of the vectorial approach which helped us uncover this class of models. Instead they are found by relying on information encoded in the higher level theory of thermodynamics. It is *thermodynamics* that provides us with a more natural classification than the dynamic laws.

4.5 CONCLUSION AND SUMMARY

In the previous chapter, I focused on the way in which physical theories classify rather than explain through a case study of Lagrangian mechanics. In this chapter, I have attempted to clarify how these classifications can be natural: that is, how they provide us with insight into the natural world.

Entropy, which is a critical quantity in physical theory, was not discovered through the search for dynamic laws, but rather through an abstract idealization of a physical process. Nevertheless, it was only in this frictionless environment that the delicate relationship between temperature and entropy can be discovered. Once recognized, the role of entropy in natural processes is overwhelming: it helps to explain the irreversibility of processes like ice melting or stirring cream into coffee. Although there have been many attempts to recreate

the role entropy plays in these lower level laws, *modeling* entropy is quite different than deriving it. Moreover, for dynamic laws to appropriately describe thermodynamic phenomena, we have to factor in thermodynamic constraints that point to highly selective classes of models within the possible worlds created by dynamic laws.

These dynamic laws do not “pick out” the relevant classes for a natural classification. Instead, it is thermodynamics which shows us important ways to classify this lower level phenomena. In this way, although thermodynamics is a macroscopic theory that does not provide us with dynamic laws, it clearly plays a fundamental role in the structure of physical theory.

5.0 REFRAMING DUHEM

In the previous chapters, I have provided two different case studies which undergird the general principles Duhem espouses in his philosophy of science. I have also provided a rough sketch of Duhem’s alternative model of a “theory of everything” — his picture of the New Mechanics.

In this chapter, I will return to some of the most difficult passages in Duhem, and use the previous case studies to shed light on Duhem’s conception of natural classifications and metaphysical explanation.

To most effectively execute the project at hand, I first consider Duhem’s anti-atomism. As I mentioned in Chapter 1, Duhem often rails against atomism at a time when experimental evidence for atomism was strong enough to convince most scientists of its veracity. I have shown that Duhem held a top-down, non-atomic picture of scientific theory, but I have not motivated his view. Why did Duhem reject atomism? Why did he suppose that our “final theory” would mirror Lagrange’s mechanics and not a vectorial approach?

5.1 THE COMPLEXITY OF COUPLING

One consistent theme in Duhem’s work is the rejection of atomism. In *Logical Examination of Physical Theory* he writes

The school of the neo-atomists, the doctrines of which center on the concept of the electron, have taken up with supreme confidence the method we refuse to follow. This school thinks its hypotheses attain at last the inner structure of matter, that they make us see the elements as if some extraordinary ultra-microscope were to enlarge them until they were made perceptible to us.

We do not share this confidence. We are not able to recognize in these hypotheses a clairvoyant vision of what there is beyond sensible things; we regard them only as *models*[Duhem, Ariew, and Barker1996, 238].

Some philosophers, such as Nancy Cartwright, interpret Duhem’s rejection of atomism as a symptom of his antirealism. However, Mark Wilson has recently argued, I think correctly, that Duhem’s anti-atomism stems from his worries about coupling. Coupling occurs when a system is perturbed by two different effects that join together and create a unique effect on the system — an effect that is more complicated than the application of each independent perturbation to the system applied successively¹. One example of coupling, common to engineers, is that of thermomechanical coupling, the interaction of the mechanical features of a system with the temperature-related features of a system. Wilson describes one example of thermomechanical coupling in “Two Cheers for Anti-Atomism”.

Consider an iron bar. If we *strike* one end with a mallet, we will send a pulse of compressive stress through its interior, a process that is governed, to first approximation, by the familiar wave equation. Likewise, if we *heat* an extremity, we will send a parcel of heat across the bar, in rough accordance with Fourier’s celebrated heat equation. But, surely, these two effects will couple to each other greatly complicating the detailed flow, because the compressive effort supplied to the bar will gradually elevate the temperature of the bar beyond our simple Fourier’s law expectations. Likewise, locally heightened temperatures will dilate the bar’s length, spoiling the simple patterns of the standard wave equations. In many industrial settings, coupling effects are sufficiently strong that one needs to find a framework in which they can be jointly treated [Wilson2013, 5].

As can be seen from Wilson’s example, even if we know how an iron bar will behave when we strike it with a hammer or if we heat it, we cannot necessarily predict (without any additional mathematical machinery) what is going to happen to the bar if we do both of these things at the same time. To understand how these two interventions on the system will couple together, it is often important to introduce an entirely new theoretical framework.

Wilson suggests that the pursuit of such a framework, where coupled systems can be treated jointly, is the real motivation for Duhem’s Energetics. Much disagreement arose between Duhem and scientists such as Maxwell and Kelvin concerning the treatment of coupled systems. Maxwell also accepted the use of the virtual modifications that Duhem outlined in his *New Mechanics*, but he believed these variational techniques could ultimately

¹This is meant as an intuitive definition of coupling. The formal definition states that coupling occurs when a system stores energy by at least two different means.

be reduced to merely mechanical knowledge. That is, if we had all the information about the tiniest bits of matter (say, atoms), and we knew all the rules that the atoms obeyed, then we would be able to directly deduce the behavior of coupled systems. In fact, were such speculations true, on this “most fundamental” level, the idea of coupling itself would disintegrate, as there would not be two different kinds of stored energy working together but only the motion of one kind of substance. Thermal notions, such as temperature or pressure, would no longer be applicable on this fundamental level. Instead, we would be able to say all that needs to be said using only the notions of classical mechanics².

Although Duhem adamantly denies such a reduction is possible, what he does not say, Wilson points out, is that atoms do not (or cannot) exist.

Observe that Duhem's animus towards “molecular hypotheses” (understood as efforts to escape into thermally free realms) needn't coincide with a general opposition to the postulation of minute entities below the scale of, e.g., microscopic observation. There is nothing in his writings or interests that precludes (insofar as I am aware) the study of, e.g., suspensions of very minute particles, such as contaminants in the atmosphere. It is merely that one should sometimes expect to employ thermal tools in these settings as well [Wilson2013, 6].

Wilson's point is that Duhem was not against the existence of atoms, but against the idea that all phenomena, e.g. coupled systems, could be appropriately formally represented without any reliance on the concepts of theories such as thermodynamics. A key point in the structure of thermodynamics is that it does not attempt to reduce all of the various ways that energy can be observed into a single mechanical picture.

Not only does Wilson charitably illuminate Duhem's anti-atomistic rants, he also provides important observations concerning the history of science that provide credibility to Duhem's view. One such observation pertains to problems of conceptual closure that occur if we attempt to eradicate thermal (or other non-mechanical) notions from our scientific theories. Wilson uses the example of shockwaves to illustrate this problem.

Shockwaves are high pressure wave fronts that can develop when the source of a sound is moving as fast as (or faster than) the speed of sound. The effects of shockwaves were observed in WWII when fighter pilots noticed additional drag when they descended at close to sonic speeds. The build up of the sound waves near the sound source create a high

²Wilson refers to classical mechanics as the “old mechanics,” in contrast to Duhem's New Mechanics.

pressure front known as a “sound barrier.” When the sound barrier is crossed by the sound source, a famous “sonic boom” is emitted. An important feature of shockwaves is, due to their compressed wave front, their behavior is nonlinear.

Shockwaves can also be created by injecting a high pressure pulse to a moving gas. The example Wilson uses concerns a blast of air traveling in a tube. If we administer a high pressure pulse to this gas, we will create a barrier much like the sound barrier created by planes traveling at sonic speed. Now, if we apply the standard gas dynamics for classical systems (that is, systems that can be modeled in terms of position, velocity, and density) the dynamics of the gas should be describable by the inviscid Burger’s equation

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = 0 \tag{5.1}$$

where u is the velocity and x is the position.

However, Wilson points out that our pulse will cause a “piled up singularity,” which is the shock front, “that leaves us with no ‘mechanical’ criterion for deciding how the gas will distribute itself over a fan-like region lying behind the shock front” [Wilson2013, 8]. In order to understand how the gas will distribute, we have to import the thermal notion of entropy. Thermodynamics tells us that the entropy of the gas will expand, and by appeal to this thermodynamic principle we are able to predict how the gas will distribute across the region. Hence, our solely mechanical model fails and *must* rely on the principles and concepts of thermodynamics.

Escaping to a more “fundamental” approach, in this case, will not help either. That is, one might try to avoid the introduction of entropy by moving to a smaller-scale framework than the one at which the gas was initially modeled. However, empirical reasons suggest that we must model these smaller-scale “atoms” as non-rigid, flexible bodies. If that is the case, then we will still make use of continuum modeling (such as the inviscid Burger’s equation) to describe how these molecules behave when they are impacted. But pulses in the molecules will still be transmitted in a manner that is non-linear, thereby allowing shock fronts to build up inside of our atoms just like the initial shock fronts that we were trying to avoid by moving to this smaller scale length. Should we go “lower” to escape *these* shock fronts, the same thing will happen. Eventually we will have to introduce thermal notions in order to

model the distribution of the gas caused by the shockwave. If this is the case, then the “old mechanics” with its purely mechanical notions are not sufficient to describe all observable phenomena.

As Wilson points out, worries like this are what motivate Duhem to search for a framework where thermal notions, mechanical notions, magnetic notions, etc. can co-exist in peaceful harmony.

These difficulties, and plenty of others it would take too long to enumerate, tell us that it is time to stop; that it is not necessary to follow any further these attempts made to decrease more and more the number of primary notions upon which Physics rest [Duhem1980, 94].

Duhem’s opposition to “atomism” is really an opposition to attempts to eradicate non-mechanical notions (e.g. thermal, magnetic, electric notions) from physical theory.

To attempt to reduce all the properties of substances to the properties of shape and motion seems a chimerical enterprise, either because such a reduction would be obtained at the price of complications that would scare our imagination away, or even because it would be in contradiction with the nature of material things.

And so we are now obliged to accept into our Physics something other than the purely quantitative elements treated by geometers, to admit that matter has *qualities* [Duhem1980, 103].

Duhem’s pursuit of energetics was an attempt to unify disparate areas of research according to their grounding in the conservation of energy. Whether or not atoms exist was not the question of interest to Duhem; likewise, the success of atomic modeling to solve problems like Brownian motion was not problematic on his view. Duhem was perfectly comfortable with admitting mechanical notions — and even atomic ones. What he was uncomfortable with was the elimination of these non-mechanical notions from the realm of physical theory. Duhem thought that mechanical systems exist, but most real systems are coupled systems and their behavior cannot be captured by the old mechanics alone.

5.2 DISGUISED APPEALS TO VARIATIONAL PRINCIPLES

So far I have provided some of the scientific background which influenced Duhem, along with sketching both his reasons for rejecting the atomism of the “old mechanics” and his

expectations for the New Mechanics. Now, in light of these developments, let us return to the previously discussed tension between natural classifications and explanations.

Duhem denies that science aims at explanation, where an explanation is a metaphysical theory about the ontology of the world. Simultaneously, Duhem insists science aims to reflect a natural classification — a “behind the veil” peek at the true ontological ordering of physical phenomena.

In this section I will argue that Duhem’s disavowals of explanation are best understood as resisting the claim that science permeates the appearances of a particular phenomena in order to give us a detailed description of the dynamic laws from which it can be derived,. Nevertheless, Duhem maintains the belief that scientific theory provides us with novel physical insight — suggesting he is *not* antirealist, as many have argued. In fact, many of the passages used to suggest that Duhem is antirealist can alternatively be read as accounts of virtual modifications in physical theory.

For example, in his description of how theories are formed, Duhem writes

Among physical properties which we set ourselves to represent we select those we regard as simple properties, so that the others will supposedly be groupings or combinations of them. We make them correspond to a certain group of mathematical symbols, numbers, and magnitudes, through appropriate methods of measurement. These mathematical symbols have no connection of an intrinsic nature with the properties they represent . . . Through methods of measurement we can make each state of a physical property correspond to a value of the representative symbol, and vice versa . . . The diverse principles or hypotheses are then combined together according to the rules of mathematical analysis . . . The magnitudes on which his calculations bear are not claimed to be physical realities, and the principles he employs in his deductions are not given as stating real relations among those realities; . . . The various consequences thus drawn from the hypotheses may be translated into as many judgments bearing on the physical properties of bodies [Duhem1954, 19-20].

At face value, Duhem paints a pragmatic picture of the philosophy of science. We pick a few measurable quantities and stipulate that they correspond to specific mathematical symbols. We then exercise mathematical freedom in combining these quantities, bound only to the laws of mathematical analysis and without any consideration of physical laws. He even specifically states that these magnitudes are not “physical realities” and these principles are not real relationships.

However, there is a strong similarity between what Duhem describes here and the method used for forming theories based on variational principles, specifically the role played by

generalized coordinates. As was discussed in Chapter 3, handling constraint forces in classical mechanics requires the introduction of generalized coordinates. These coordinates form a generalized configuration space, which cannot be said to directly map on to the physical world. However, these quantities are still believed to *exist*, despite the fact that the machinery we use to specify our system involves the introduction of purely mathematical coordinates.

There is a strong resemblance between the aforementioned excerpt from *Aim and Structure* Lanczos' text on the role of variational principles in mechanics.

Analytical mechanics is a completely mathematical science. everything is done by calculations in the abstract realm of quantities. The physical world is translated into mathematical relations. This translation occurs with the help of coordinates. The coordinates establish a one-to-one correspondence between the points of physical space and numbers. After establishing this correspondence, we can operate with the coordinates as algebraic quantities and forget about their physical meaning. The end result of our calculations is then finally translated back into the world of physical realities [Lanczos1964, 7].

Perhaps Duhem's claim that magnitudes are not physical realities is a reference to formal features associated with the use of generalized coordinates in mathematized physical theories, rather than a defense of antirealism. When we employ generalized coordinates, we are bound to an abstract configuration space that does not directly correspond to the physical world. Likewise, many of the mathematical procedures we perform are not physically conceivable, let alone realizable. But this problem permeates both the observable and unobservable aspects of our theory, rather than privileging the former over the latter, in the antirealist vein of van Fraassen does,

The same can be said regarding Duhem's notion of truth in a theory. Just as Lanczos points out that the end result of our calculations is translated back to the physical world in order to compare our predictions by experiment, Duhem writes

Thus a true theory is not a theory which gives an explanation of physical appearances in conformity with reality; it is a theory which represents in a satisfactory manner a group of experimental laws. A false theory is not an attempt at an explanation based on assumptions contrary to reality; it is a group of propositions which do not agree with the experimental laws. *Agreement with experiment is the sole criterion of truth for a theory* [Duhem1954, 20-21, emphasis his].

In contemporary contexts, the aforementioned quote appears to assert that the unobservable entities in our theory don't really exist but are mere instruments of predictive success.

However, with respect to virtual modifications, this quote takes on an entirely different meaning. Duhem is arguing that theories that rely on virtual modifications are in fact *true*, even though the calculational machinery they rely on describes unrealizable modifications.

As I pointed out in Chapter 1, Cartwright accuses Duhem of antirealism about theoretical laws. I am suggesting the *opposite* is true: Duhem *defends* physical theories even when they include models that are not physically realizable. When Duhem claims that agreement with experiment is the sole criterion for the truth of a theory, he is not suggesting that theories *merely* agree with experiment. Instead, he is claiming that *true* theories may contain highly abstract models that describe unrealizable processes. Nevertheless, we shouldn't take theories built on these types of principles to be any less true or "fundamental" than theories that attempt to give microscopic modelings of particles.

As for theoretical laws, Duhem states that they are neither true or false, strictly speaking, but always approximate. To unpack this, he compares physical laws with what he calls "common-sense laws" such as "All men are mortal." Duhem suggests that the labels true or false make sense in the case of common-sense laws, because the interpretation of such laws is effectively fixed by the terms they use and the scope of such laws is always universal. However, he does not think physical laws can be simply classified as either true or false.

Such is not the case with the laws that a physical science, come to full maturity, states in the form of mathematical propositions; such laws are always symbolic. Now, a symbol is not, properly speaking, either true or false; it is, rather, something more or less well selected to stand for the reality it represents, and pictures that reality in a more or less precise, a more or less detailed manner. But applied to the a symbol the words "truth" and "error" no longer have any meaning; so, the logician who is concerned about the strict meaning of words will have to answer anyone who asks whether physics is true or false, "I do not understand your question" [Duhem1954, 168].

To Duhem, a physical theory is a kind of representation — similar to a statue such as Michelangelo's David. It does not make sense to ask if Michelangelo's representation of David is "true," because statues are not the sort of things that can be true or false. Occasionally, we refer to a work of art as being "true to life," but what we mean by this idiom is that the artwork very closely approximates the aspect of the world it attempts to represent. The real means for evaluating representations is a comparison between the representation and the represented — and this comparison occurs along a continuum of similarity rather

than a binary function. Additionally, Duhem does not limit the scope of representations to “entities” in physical theories; he claims that physical laws are themselves representations. Mathematical formulated theoretical laws are actually representations of the regularities which govern the world in which we live.

Although Duhem attempts to clarify how the laws of physics can be more or less approximate, his explanation only serves to confuse the issue.

The experimental method, as practiced in physics, does not make a given fact correspond to only one symbolic judgment, but to an infinity of different symbolic judgments; the degree of symbolic indetermination is the degree of approximation of the experiment in question. Let us take a sequence of analogous facts; finding the law for these facts means to the physicist finding a formula which contains the symbolic representation of each of these facts. The symbolic indetermination corresponding to each fact consequently entails the indetermination of the formula which is to unite these symbols; we can make an infinity of different formulas or distinct physical laws correspond to the same group of facts. In order for each of these laws to be accepted, there should correspond to each fact not the symbol of this fact, but some one of the symbols, infinite in number, which can represent the fact; that is what is meant when the laws of physics are said to be only approximate [Duhem1954, 169]

. This difficult passage is best unpacked by drawing on some of the points about variational principles, specifically the example of virtual work.

First, the treatment of physical problems by the variational method is full of the indeterminacy to which Duhem alludes. This indeterminacy begins with the scientist’s choice of coordinates. Remember, Lagrange’s approach to mechanics requires any specific physical problem to be represented in generalized coordinates. These coordinates allow the scientist a freedom in his choice of representation, as the system does not directly determine which set of coordinates a scientist must use. Lanczos writes

We are allowed sovereign freedom in choosing our coordinates, since our processes and resulting equations remain valid for an arbitrary choice of coordinates. The mathematical and philosophical value of the variational method is firmly anchored in this freedom of choice and the corresponding freedom of arbitrary coordinate transformations [Lanczos1964, xxv].

Lagrange’s mechanics makes use of generalized coordinates, which allow a practitioner to choose from a variety of possible methods to specify the state of the system; given any particular system, we are free to represent it in a multitude of ways. Hence, as Duhem says, a given fact does not correspond to one symbolic judgment but to an infinity of them.

Another type of indeterminacy, which pertains to laws, concerns the accuracy of experimental measurements. After the aforementioned passage, Duhem describes an example of scientists trying to determine the path of the sun as it rotates around the Earth. First, they represent the sun as a perfect sphere (knowing this is partially inaccurate), and then attempt to construct the trajectory of the center of the sphere. Next, the scientists try to fill in details about the path of the sun by means of a great deal of measurements, many of which contain various degrees of inaccuracy. While experimenters intend to measure the location of the center of the sun as it travels along its trajectory, their instruments might only allow them to pinpoint the center of the sun within the range of an inch. But if this is the case, any law which describes the arc of the sun within an inch of its actual path will agree with our data. But surely there are a great number of laws which can agree with the data at this level of proximity. Hence, Duhem claims we can make an infinity of different formulas or physical laws to correspond to the same group of facts.

The concerns Duhem expresses here appear to be rather run of the mill worries about accuracy, with no specific realist or antirealist implications. His primary point is that our mathematized laws are actually representations of physical laws, and thereby more or less accurate depending on the precisions of measurement and the strength of the instruments we use. He emphasizes this point by claiming

Any physical law, being approximate, is at the mercy of the progress which, by increasing the precision of experiments, will make the degree of approximation of this law insufficient; the law is essentially provisional. The estimation of its value varies from one physicist to the next, depending on the means of observation at their disposal and the accuracy demanded by their investigation [[Duhem1954](#), 174].

Our laws are limited by the instruments we use to collect data. Science is heavily dependent on observation, and observation allows some degree of indeterminacy into our laws.

I think Duhem is also worried about another, deeper sort of indeterminacy that arises from the nature of the variational method. To understand this kind of indeterminacy, it is helpful to focus for a moment on the example of Jack and Jill on the seesaw.

Recall that the Principle of Virtual Work is formulated as

$$\delta W = \sum F_n \cdot \delta q_n = 0. \quad (5.2)$$

where W is the work done, F is the force, q is some generalized coordinate, and δ represents a virtual infinitesimal displacement.

A seesaw is in equilibrium when the virtual work vanishes — when any infinitesimal displacement consistent with the constraints is “balanced” by the other forces acting on the system. Recall that the relevant consequence of the principle of virtual work in this particular case is the law of the lever, which states that the seesaw will be in equilibrium when

$$m_1 g l_1 = m_2 g l_2. \quad (5.3)$$

Now suppose that just when Jill figures out where to sit on the lever to balance the seesaw, Jack gets distracted and decides to go play video games. Jill sinks to the ground moping, wondering how she can return the seesaw to a state of equilibrium. Lucky for Jill, Jack’s little brothers come outside, and Jill manages to convince them to sit on the seesaw with her. The two brothers can now shift their position on the beam to bring the seesaw back into a state of equilibrium. Although the actual forces on the seesaw are now quite different (there are two boys instead of one, they will have to sit at different locations along the lever), it still follows that for any infinitesimal, virtual displacement of Jill, the virtual work will vanish.

We replace Jack with his two brothers without causing a noticeable difference to the system, at least not from the standpoint of the coordinates used in virtual work, because we can represent the system in several different ways — either by Cartesian coordinates, or coordinates which represent the angle between the lever and the fulcrum and the length of the lever, etc. There are numerous generalized coordinates we might pick to symbolize our system. Notice the difference between this and the point-mass formulation, where it is crucial to know each specific force at work on the individual atoms which compose the system.

Because Lagrange's mechanics does not individuate particular forces at work in a system, the equations for equilibrium when Jack and Jill sit on the seesaw are identical to the equations when Jill and Jack's brothers sit on the seesaw. As Duhem points out,

It is therefore not too important to know in detail each of the forces applied to the various bodies of the system, its point of application, its magnitude, its direction; provided that the information given about the set of forces allows the calculation of the virtual work done over a virtual displacement, enough is known; all other additional data are superfluous . . . Thus it is that for the various forces applied to a solid body one can substitute a certain set of two forces, either a force and a couple, or yet other combinations of forces; all these combinations which appear distinct to a geometer, provide the same work for a virtual displacement of the solid body; a mechanician does not therefore distinguish the one from the other [Duhem1980, 25-26].

Not every geometric difference is an algebraic difference. While it might be clear to a geometer — or one interested in a detailed picture of our system — that different forces are at work when Jill sits with Jack than when Jill sits with Jack's brothers, from the standpoint of a mechanician — one who is concerned with the working structure of the system — these two systems are indistinguishable. It is perfectly fine to substitute a single force with three forces in the same direction as long as the total of the forces is the same.

To connect this with the language Duhem previously used to describe physical laws, we see a high degree of indeterminacy in the principle of virtual work because this principle does not represent a system with a high level of detail. Therefore, if we were to ask the question of how similar is the representation to the system (in the way we might ask how much a portrait reflects its subject), it seems that the principle provides only a chalk outline of the system it models. But this is not to say the law is false — only to say that it is not a complete representation of the seesaw.

It is important to point out that Duhem does not take the abstraction or indeterminacy of Lagrange's mechanics to be a flaw. In fact, he thinks this is just how physical theory works. Theorists aim to create a representation of a system that will provide the groundwork for developing mathematical laws which describe experimentally observed regularities. It's quite natural that the representation is incomplete, because it is only representative of a particular data set gathered under highly restricted experimental conditions. The world is complex, and so are its laws. We are, however, able to form symbolic laws which are approximate

to these within a restricted domain. As our experimental context grows more complex, and begins to reflect the complex interrelationships of systems in the real world, we see how our physical laws are incomplete.

This becomes apparent when we witness the phenomena of coupling as previously discussed. On the provisionality of symbolic laws, Duhem states

Physical law is provisional not only because it is approximate, but also because it is symbolic: there are always cases in which the symbols related by a law are no longer capable of representing reality in a satisfactory manner . . .

Let the physicist place some oxygen between the plates of a strongly charged electrical condenser; let him determine the density, temperature, and pressure of the gas; the values of these three elements will no longer verify the law of compressibility and expansion of oxygen. Is the physicist astonished to find his law at fault? Not at all. He realizes that the faulty relation is merely a symbolic one, that it did not bear on the real, concrete gas he manipulates but on a certain logical creature . . . too simple and too incomplete to represent the properties of the real gas placed in the conditions given now. He then seeks to complete this schematism and make it more representative of reality . . . [Duhem1954, 174].

As for the first portion of this quote, when Duhem claims that the law is provisional *because* it is symbolic, he is emphasizing that the law is a representation. Therefore, the approximate nature of representation plays a role in the provisionality of physical law.

To explain how this occurs, Duhem considers the effects of an electric field on the compression or expansion of an ideal gas. Under normal conditions, oxygen behaves similar to an ideal gas, close enough for its behavior to be described by the ideal gas law

$$PV = nRT \tag{5.4}$$

where P is the pressure, V is the volume, n is the amount (number of moles) of the gas, R is the ideal gas constant, and T is the temperature. However, when a gas such as oxygen is placed between the plates of a strongly charged electrical condenser, the electric field will affect the internal energy of the system. Therefore, the traditional means of tracking the expansion of the gas will fail due to the coupling of the thermodynamic system with an electric system.

But this “failing” of the ideal gas law will not trouble the physicist, insists Duhem, because he knows the “faulty relation is a symbolic one.” The model he was using to predict the behavior of the gas was too simple for this particular physical situation. Other properties of the gas are relevant to understanding its behavior in an electric field. Or, perhaps, other

environmental features of the system apart from the gas are relevant. Does this mean the ideal gas law is false? Not at all.

Consider, for a moment, facial composites that are created by criminal sketch artists. Victims and witnesses of crimes attempt to describe perpetrators, and sketch artists fill in the details with a sketch. These sketches (and the descriptions they correspond to) can be used to eliminate certain suspects. Imagine that a victim was assaulted in a night club. The night club has a large number of cameras, including cameras at every entrance and exit, so the police have a means to compile a list of who was in the bar that night. The assault did not happen on camera, but was witnessed by several patrons of the club (in addition to the victim). Suppose that when the victim describes the perpetrator, he is able to report to the criminal sketch artist that the assailant was a blonde, white male — and nothing else. The victim is in shock and the assault is a bit of a blur. The sketch artist attempts to draw a generic picture of a white male, emphasizing the blond hair color. Let's say that our sketch artist is careful and doesn't want to misrepresent the victim's testimony, so they leave the facial features (eyes, nose, mouth, facial shape) out of the drawing.

What can we say of this representation? Is it *false*? Let's assume that everything the victim said was true — the assailant was a blonde, white, male. It's not that the picture is false or even inaccurate, it simply doesn't represent enough information to sufficiently pick the suspect out on security tapes. Surely the police will be able to eliminate *some* suspects on the basis of this information, but it won't justify any search warrants.

Later, the sketch artist speaks to one of the witnesses. She is quite sure the assailant was in his forties, with a rounded face and strong nose. The specific type of strong nose is left to the imagination of the sketch artist. Even if the artist properly represents the perpetrator's nose, the sketch still lacks sufficient detail to pick out our assailant.

We can imagine this story wandering down several different paths. It's possible that our sketch artist will accurately represent all of the data she gathers, but the drawing still does not contain enough data to eliminate the majority of the bar patrons. Even worse, it's possible that she will create a sketch that eliminates *the perpetrator*, because she misrepresents some detail. Perhaps the nose she creates looks crooked to the detectives, and they eliminate a man with a strong straight nose. Even in this case, I don't believe we would

call the representation itself *false*. To say the drawing is false is, as Duhem points out, a category mistake. What we would judge instead is that the representation, however well it fits the data collected, lacks the appropriate detail to eliminate suspects in this particular context. Had the crime happened in a small convenience store, where there were only a few suspects, perhaps the drawing would successfully pinpoint the perpetrator. Unfortunately, the potential suspect pool was too large — there were too many white males in their forties at this particular bar for the criminal sketch to serve its purpose.

Of course, police detectives do not rely heavily on criminal sketches, and they certainly don't count as admissible evidence in court, in part for the reasons mentioned. Nevertheless, we do think such sketches can be accurate when they contain the right amount of detail for an appropriately sized pool of suspects. This is analogous to Duhem's point about the role of representation. Even when representations lack some detail that is relevant for making a particular prediction, it does not follow that they are false or inaccurate. They simply do not approximate our system closely enough in order to adequately describe its behavior. On the other hand, some representations that do lack a great deal of detail *are* able to predict the behavior of a system in certain contexts. This is what we witness by the success of Lagrangian mechanics.

While Duhem does express concerns about the truth of scientific theories or the inaccuracy and approximate nature of science, I have argued that such claims are actually references to the use of variational principles in the building of physical theory, rather than “antirealist” worries about unobservables, etc. Additionally, I have argued that Duhem's worries about “approximation” have more to do with representations lacking detail than representations being “false.” Duhem does not think representations can be true or false. To say physical law is a representation is to say that it models some regularity in the phenomenal world. This does not mean the law lacks truth or accuracy — it might only lack detail.

5.2.1 Dissolving the Tension

In the previous section I have presented some of the difficult passages from Duhem, in an effort to show how worries about variational principles underlie many of his comments about

the philosophy of science. In this section, I take up the question of how a theory can fail to explain but still be considered a natural classification. The clearest answer to this question comes from a look back at the examples of analytic mechanics and thermodynamics.

5.2.1.1 Physical Theories do not Explain First I will argue that neither Lagrangian mechanics nor thermodynamics provide an explanation, according to Duhem’s use of explanation:

Duhemian Explanation (DE): A scientific theory provides an explanation if it presents us with a non-sensual (“behind the veil”) description of the physical bodies which cause the phenomena we observe.

Duhem asserts that science does not aim to reach behind the veil and provide us with information about the essence of bodies. This certainly seems true in the case of Lagrangian mechanics. Instead it makes use of a generalized configuration space, where we select our coordinates according to the degrees of freedom of some macroscopic system. For example, to represent a rigid body, we are able to select six different coordinates to denote the state of the system. But these coordinates tell us about the body’s spatial location and degree of rotation. They tell us nothing about the nature of the body or its essence. On the contrary, generalized coordinates allow us to ignore the microscopic features of the system in order to make predictions. Reflecting on this aspect of Lagrange’s mechanics, and the role it will play in the New Mechanics, Duhem writes

The New Mechanics is organised, not for the speculative contemplation of the essence of things, but for the practical necessity of acting upon the bodies of the external world and modifying them according to our needs [Duhem1980, 118].

The same can be said for thermodynamics. Although thermodynamic theory describes very general and stable constraints that apply to a wide range of materials, it doesn’t give us insight into the essence of heat or the nature of matter. In fact, one of the most interesting things about thermodynamics is that it picks out classes of thermal phenomena that can be realized in an assortment of materials with varying microscopic structure. Thermodynamic laws are true *independently* of the microstructure of the systems which instantiate them.

We learn nothing about the essence of matter from the application of Lagrange’s mechanics or thermodynamics. We do, however, learn all about the behavior of matter. When

Duhem asserts that science is not for speculative contemplation but for practical necessity, I don't think he is intimating that science is solely pragmatic. Rather, he is thinking of concrete cases in physical theory in which we need not know a substance's metaphysical essence in order to understand the laws it obeys.

For example, bodies are represented in Lagrange's mechanics according to observations about their degrees of freedom rather than their essential nature. There is a great amount of detail in the system that we are able to ignore. In this sense, our theoretical laws are approximate because they do not reveal the underlying microstructure is in our system, or provide a dynamic laws type story of how bodies manage to evolve. Concerning the question of how bodies are able to move and how work is able to be done, Duhem comments

What is the nature of this contribution [of work done on a system] and how is it carried out? A difficult problem, whose clear solution seems to be quite beyond the bounds of human reasoning. But this problem of the *communication of substances* is the object of Metaphysics, not of Physics. Physics does not attempt to elucidate it . . . [Duhem1980, 119].

According to Duhem, physical theory does not reveal precisely how energy is exchanged — it only tells us that energy is exchanged. It gives us *some* information about the causal chains in physical processes, but fails to give us enough detail for this description to count as a metaphysical picture. Likewise, thermodynamics does not give us a clear story of the continuous, dynamic evolution of heat. Instead it reports that heat is gained in a Carnot cycle and converted into work. But *how* this process occurs remains a mystery. There is no “picture theory” of heat in classical thermodynamics, such as the porous sponge described by caloric theory.

Duhem further claims that

When a physical theory is taken as an explanation, its goal is not reached until every sensible appearance has been removed in order to grasp the physical reality . . .

Thus, it follows that in order to judge whether a set of propositions constitutes a physical theory or not, we must inquire whether the notions connecting these propositions express, in an abstract and general form, the elements which really go to make up material things, or merely represent the universal properties perceived [Duhem1954, 9].

An explanation, on Duhem's view, must represent *the elements which compose material things*. In contemporary terms, explanations are built upon microstructures. Explanations include the most “basic” components of the world, and describe their interactions. But

this is clearly *not* what analytic mechanics *or* thermodynamics aim to describe. Both fail Duhem’s criteria for an explanation because they do not include details about a system’s microstructure.

Interestingly, at times Duhem seems to think parts of science can achieve “explanations.” As I mentioned in Chapter 1, Duhem presents the theory of acoustics as the paradigmatic case of an explanation. Reflecting on the theory, he writes

The explanation which acoustic theories give of experimental laws governing sound claims to give us certainty; it can in a great many cases make us see with our own eyes the motions to which it attributes these phenomena, and feel them with our fingers . . .

Most often we find that physical theory cannot attain that degree of perfection; it cannot offer itself as a certain explanation of sensible appearances, for it cannot render accessible to the senses the reality it proclaims as residing underneath those appearances . . . [Duhem1954, 8].

This quote suggests that Duhem thinks the acoustic theory of sound *might* be an actual explanation of sound. Sound might *really behave* like a wave. However, he these instances of explanation provide us with the basic objects and laws unify all of the disparate phenomena we perceive. He denies that these explanations (such as the one for sound) are “fundamental.”

It is clear that neither Lagrange’s mechanics nor thermodynamics is an explanation in Duhem’s sense of the word. Lagrange’s theory does not tell us about the essence or true nature of matter. The New Mechanics, which Duhem believed was the final form of physics, was to be grounded in the laws of thermodynamics with a structure quite analogous to Lagrange’s mechanics. When Duhem claims “physical theory” does not explain, he is referring to a physical theory structured like the New Mechanics.

5.2.1.2 Convergence on a Natural Classification After arguing that science does not aim to give explanations, Duhem provides a new definition of a physical theory.

A physical theory is not an explanation. It is a system of mathematical propositions, deduced from a small number of principles, which aim to represent as simply, as completely, and as exactly as possible a set of experimental laws [Duhem1954, 19].

For Duhem, a theory begins with principles. Lagrange begins his theory with the Principle of Virtual Work, which is later generalized to d’Alembert’s Principle. From these two principles, we are able to deduce differential equations (or, in Duhem’s terms, mathematical

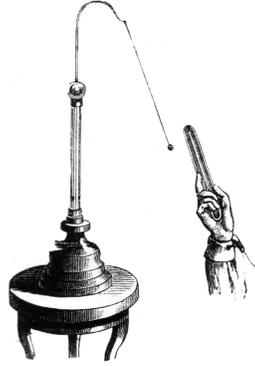


Figure 4: An example of a pith ball electroscope.

propositions) whose solutions accurately describe observable patterns of motion. Notice the contrast between this concept of a physical theory, which begins with general principles, and a theory which begins with a basic ontology and set of dynamic laws.

In order to understand Duhem's distinction between a principle and an experimental law, I will take a brief detour to explicate some of Duhem's terminology. Experimental laws are the generalizations of facts. Facts are propositions that describe the results of individual experiments.

Placed in contact with the external world in order to understand it, the human mind first encounters the domain of facts. It sees that a piece of amber, rubbed by a silk rag, attracts a pith ball suspended from a silk thread, at a distance; that a piece of glass, rubbed with a woollen rag, does the same thing; that a piece of copper, rubbed with the same woollen rag, also does the same thing, provided that the piece of copper and the woollen rag are both carried by a glass sleeve, etc. Each observation, each new experiment, presents a new fact [Duhem, Ariew, and Barker1996, 1].

Duhem considers a scientist performing experiments by means of a pith ball. A pith ball electroscope, illustrated in figure 4³, was one of the first instruments used to measure electricity. Pith balls are lightweight objects covered in metal that are easily moved by an electric charge. When a scientist observes that a piece of amber rubbed by a silk rag attracts a pith ball, she obtains particular knowledge about the results of the specified experiment.

³Drawing of the pith ball electroscope by Morten Bisgaard, [?].

The fact demonstrated in the experiment is that this particular piece of amber rubbed by this particular piece of cloth attracted that pith ball.

One might wonder how the scientist knows that the object is amber or that the rag is silk, as these terms imply classifications made available by scientific theory. Although Duhem notoriously worries about holism in experiments, he sets those issues aside in this context. I intend to follow suit and take for granted that when the scientist observes a particular experiment, she is warranted in believing a particular fact about what she has observed. Perhaps we can say, at the most basic level, that she observes “when that thing is rubbed by that thing it attracts that thing,” replacing terms like amber and silk with ostensive definitions. Nevertheless, Duhem is unconcerned with such details at this point.

After a scientist repeats an experiment (and variations of that experiment) over and over, she begins to generalize the facts she has observed by means of induction. These generalizations, which occur in a non-formal language, are what Duhem considers experimental laws. For example, a scientist might observe that a particular rock falls when dropped. She might drop the rock several times until she make the induction that this particular rock always falls when dropped. She might vary the experiment by dropping many different kinds of rocks and eventually come to the generalization that all rocks fall when dropped. These universal statements about particular data sets are what Duhem calls experimental laws.

The mind arrives at the understanding of *experimental laws* through induction, transforming the facts it has come to understand. Thus, the facts we have just mentioned, and other similar facts it is possible to observe, leave the mind, through induction, to this law: When similarly rubbed, all bodies become capable of attracting a pith ball suspended on a silk thread. Creating a new word to express the general property that this law asserts, the mind says: Through suitable rubbing, all bodies are electrified [Duhem, Ariew, and Barker1996, 1-2].

Duhem continues his example of the pith ball experiments, showing how scientists form experimental laws about electricity. Duhem suggests these laws are generalizations of the experimental facts formed by the use of induction. Facts are the first level of understanding the external world, and the generalizations of these experimental facts are the second.

In *Aim and Structure*, Duhem considers a third level of understanding. This third level of understanding is the condensation of experimental laws into *principles* — what Duhem con-

siders the formation of a physical theory. In the same way that experimental laws generalize facts, principles generalize experimental laws.

From these principles we can always, through regular and sure calculation, extract the law we wish to use. It is no longer necessary, therefore, to keep watch over the knowledge of all these laws; the knowledge of the principles on which they rest is sufficient [Duhem1954, 23].

According to Duhem, principles are the foundations of physical theories, and they provide us with a means of calculating experimental laws. We find a concrete example of this when we consider analytic mechanics and thermodynamics.

Lagrange found a way to derive the experimental laws of statics, such as the law of the lever, from the more general Principle of Virtual Work. The Principle of Virtual Work cannot be thought of as a physical law, because it requires a decent amount of abstraction and relies on mathematical experiments performed outside of time. But it does give us a way of *deriving* the experimental laws and the laws of experience. As I pointed out in §3.3, the amazing feat of Lagrange was to unify other experimental force laws (e.g. the law of the lever, or the force summation law) into separate applications of the same principle. In this way, the foundational principles of physical theory condense experimental laws.

The same can be said for the basic principles of thermodynamics. Thermodynamics does not rest on dynamic laws — which tell us how our system evolves — but instead provides general rules for how the motion of systems will be constrained. These principles help us create representations, and even approximate dynamic laws, for different systems. Still, the principles are not themselves dynamic laws — and the representations in thermodynamics are quasi-static processes rather than continuous time evolutions.

However, Duhem is quick to point out that physical theory doesn't only condense the laws, it also classifies them.

Theory is not solely an economical representation of experimental laws; it is also a *classification* of these laws ...

Theory gives, so to speak, the table of contents and chapter headings under which the science to be studied will be methodically divided, and it indicates the laws which are to be arranged under each of these chapters [Duhem1954, 24].

We can see an example of how theories classify by considering Duhem's picture of the New Mechanics. Recall that Duhem claimed the New Mechanics was like a tree rooted in the Principle of the Conservation of Energy, which then divides into systems with reversible modifications, frictional systems, systems with permanent alterations, and systems that contain currents. From these branches grow different experimental laws. These laws are classified according to the principles from which they can be derived.

We can also see how Lagrange's theory helps classify experimental laws. By use of d'Alembert's suggestion that we can define the quantity $-ma$ as the inertial force, and set the sum of the applied force and the inertial force to zero in classical dynamics, Lagrange showed that the laws of dynamics could be classified with the laws of statics. Once d'Alembert's principle proved that dynamic systems can be classified as special cases of equilibrium conditions, then the tools Lagrange had developed for the treatment of statics could be extended to treat dynamics. As Duhem writes,

These classifications make knowledge convenient to use and safe to apply. Consider those utility cabinets where tools for the same purpose lie side by side, and where partitions logically separate instruments not designed for the same task: the worker's hand quickly grasps, without fumbling or making a mistake, the tool needed. Thanks to theory, the physicist finds with certitude, and without omitting anything useful or using anything superfluous, the laws which may help him solve a given problem [Duhem1954, 24].

Lanczos echoes Duhem's thoughts concerning the use of d'Alembert's principle.

The importance of [d'Alembert's principle] lies in the fact that it is *more* than the expression of Newton's equation. It is the expression of a *principle*. We know that the vanishing of a force in Newtonian mechanics means equilibrium. Hence [d'Alembert's principle] says that the addition of the force of inertial to the other acting forces produces equilibrium. But this means that if we have any criterion for the equilibrium of a mechanical system, we can immediately extend that criterion to a system which is in motion [Lanczos1964, 89, emphasis his].

Scientific theories classify different experimental laws by showing what mathematical tools can be used to describe varying sorts of phenomena. By classifying dynamics along with statics, Lagrange shows that our methods for statics (in this case, the Lagrange multipliers) can also be applied to dynamic systems. Therefore, the fruitfulness of scientific theory goes beyond condensing experimental laws — it also classifies them.

Still, Duhem suggests the tree growing from physical theory produces more than pragmatic fruits. Duhem insists that these classifications converge on what he calls a *natural classification*. In an essay titled “The English School and Physical Theories,” Duhem introduces the idea of a natural classification by discussing what he calls a *perfect theory* as I discussed in the first Chapter. He uses this notion of a perfect theory to illustrate what makes a classification natural.

We must evidently judge the degree of perfection of a physical theory by the greater or lesser conformity which that theory offers to the ideal and perfect theory . . . [T]his ideal and perfect theory . . . would be the complete and adequate metaphysical explanation of material things. This theory, in fact, would classify physical laws in an order which would be the very expression of the metaphysical relations that the essences that cause the laws have among themselves. It would give us, in the true sense of the word, a *natural classification* of the laws [Duhem, Ariew, and Barker1996, 67-68].

A *natural classification* is an ordering of the laws that expresses the relations between the essences that cause the laws. For example, if all matter turned out to be wave-like, then a natural classification would place laws that are grounded in the amplitude of a wave increasing in one category, and laws that are grounded in the distance the wave travels in another category. It would classify together laws that reflect a common metaphysical cause, and separate them from laws that pertain to a different feature of matter.

However, Duhem thinks a perfect theory is out of our reach.

Such a theory, like everything that is perfect, infinitely surpasses the scope of the human mind. The theories which our methods permit us to construct are no more than a pale reflection of it. The metaphysical method gives us only information that is too general, too lacking in detail, and too paltry about the essence of material things to be able to serve in classifying physical laws [Duhem, Ariew, and Barker1996, 68].

A perfect theory is a metaphysical system. Duhem expresses deep concern about the move from metaphysics to physics — such as Descartes’ attempts to derive the laws of physics from the geometry of extended matter in Euclidean space [Duhem, Ariew, and Barker1996, 70]. Descartes’ concept of matter proved to be an insufficient explanation of many experimental laws. This is why Duhem thinks our ideas about metaphysics are “too paltry” to classify physical phenomena.

The experimental method, on the other hand, has problems of its own.

The experimental method, the only one to which we are to have recourse in pursuit of this goal, does not capture the essence of things, but only the phenomena through which things manifest themselves to us. It does not allow us to reconcile the laws with one another except through exterior and superficial analogies which translate the true affinities of the essences from which the laws emanate [and] perhaps frequently betray them [Duhem, Ariew, and Barker1996, 68].

Duhem claims that our only chance of capturing the physical laws is by means of the experimental method. However, this method does not reveal the true nature of bodies, as I argued in the previous section, but instead allows us to organize and classify the experimental laws. We do not classify these laws according to “the true affinities of the essences.” That is, our laws are not classified according to the true causal structure of their physical nature, because science does not fully reveal this information to us. Our classifications will never be natural classifications, because we do not ascertain the deep metaphysical structure that would allow us to see the true relations between laws.

Nevertheless, we are not without hope.

But however imperfect physical theories are, they can and they should tend toward perfection. No doubt they will never be anything but a classification, stating analogies between laws but not capturing the relations between essences. We can and should always seek to establish them in such a way that there would be some probability that the analogies brought to light by them would not be accidental agreement, but true relations, showing the connections that really exist among essences. In other words, we can and should seek to render these classifications as far from artificial, as *natural* as possible [Duhem1954, 68].

And so, to converge on a natural classification, a scientific theory ought to classify laws in a way that is parallel to the true “natural” classification of laws. This doesn’t mean that a theory will reveal the true essence of bodies, but it will be able to group together similar laws in the same way that similar laws are grouped together according to their metaphysical essences. This is best illustrated by an example.

Let’s assume (contrary to Duhem) that our knowledge of the periodic table of the elements is actually a “perfect theory,” and provides us with knowledge of the metaphysical nature of matter. Let’s further assume the following classification of matter is *the* natural classification provided by a perfect theory.

Suppose that scientists who are not in possession of this “perfect theory” view a set of substances for the first time — they are given samples of Hydrogen, Iron, Water, Ammonia,

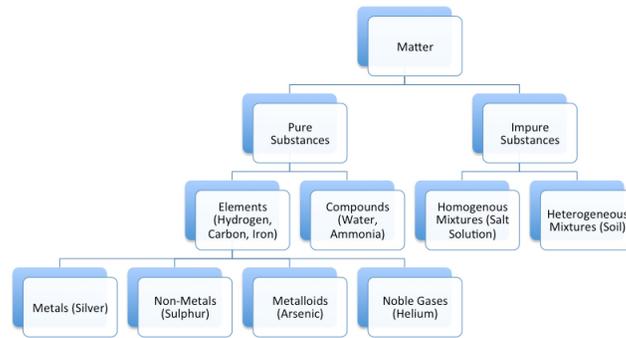


Figure 5: Example of a classification from chemical theory.

Silver, Sulphur, Arsenic, Helium, Salt, and a Salt solution. How might they classify these various substances?

On Duhem's view, the scientists aim to classify the substances in a way that the perfect theory would classify these substances. Nevertheless, we can imagine an infinite number of logical classifications. For example, they might classify silver and wood together because they are both solids and water with alcohol because they are both liquids. Or they might classify salt with salt water because they both taste similarly. All of these classifications will be based on the ways the scientists *experience* the substances rather than the nature of the substances themselves. Duhem suggests that if the scientists continue their pursuit long enough, the hope is that even though their classification is built on the observable properties of these substances, it will reflect the natural classification of the substances according to their nature.

One important disanalogy between my example and Duhem's is that I have described a classification of substances and he is interested in a classification of laws. For example, Lagrange determined that the laws of classical statics ought to be classified together with the laws of classical dynamics. Therefore, laws that govern an object remaining at rest can be derived from the same principle as the laws that govern an object in motion. On Duhem's picture, the success of Lagrange's theory suggests that it mirrors the natural classification

of laws. In other words, if we understood the true causal structure of the world, it would be true that the laws which govern classical statics share something in common with the laws that govern classical dynamics. Moreover, these laws, which are all instances of reversible modifications, lack something in common with the laws that govern irreversible modifications. This is what it means to claim that a theory converges on a natural classification — that it orders the laws in the same way a perfect theory of metaphysics would classify the laws. However, in order to do this, a theory *need not* reveal the essential nature of matter. A theory can classify laws independent of knowledge about the microstructure of matter.

5.2.2 Summary

Duhem's disavowals of explanation are not at odds with his claim that successful theories reflect a natural classification. An explanation requires the exposure of the constitutive nature of matter and a revelation of the metaphysical causal structure of reality. Duhem denies that physical theory always aims to provide this sort of information. Sometimes, scientists unify disparate laws by appeals to variational principles — principles which ignore the fundamental constitution of matter and focus instead on its phenomenal qualities. These variational principles are mathematical experiments which occur in an abstract configuration space that is not meant to mirror the metaphysical structure of the world.

Duhem is not a contemporary antirealist. He does not have worries about the observable and unobservable entities in theories. Instead, he is concerned with defending abstract, mathematized formulations (such as Lagrange's mechanics and thermodynamics) as candidates for true, fundamental theories *despite* the ignorance of microstructure that such theories convey. Duhem's main purpose is to create a model of a fundamental scientific theory that breaks the mold of the dynamic laws strategists. He shows us that science can classify phenomena, and provide insight into the natural world, without specifying a set of dynamic laws.

Even though Duhem thinks we should employ variational principles in the composition of a general scientific theory, and the use of such principles avoids peeking into the essence of matter, he still believes that the classification of laws the theory provides mirrors the

classification of laws provided by a perfect metaphysical theory. Thus, science does not merely economize the data collected in a lab, it also tells us something about the relations between laws and the metaphysical structure of the world.

6.0 CONCLUSIONS

This project began as an investigation into a particular tension in the work of Pierre Duhem: what is the difference between a scientific explanation and a natural classification? I have explicated several case-studies that were central to the development of Duhem's thought, and argued that we can only understand the distinction outside of the *dynamic laws* picture of scientific theory. Duhem's conception of scientific theory is quite different from ours; he suspects that science progresses by developing general principles from which we could derive specific equations of evolution for particular systems. On the dynamic laws picture, science progresses by specifying equations of evolution that hold for the most fundamental level of matter, and then explaining how observable phenomena is reducible to these laws. While it's true that on a dynamic laws picture of science natural classifications and explanations are one in the same, Duhem's ideal theory — which rests on variational principles and quasi-static processes — provides classifications without explanations.

Having resolved the historical tension, I would like to both summarize and expand on some of the general points of application for contemporary philosophy. Duhem emphasizes several morals which are still relevant to current philosophical trends, and our current understanding of scientific theory can greatly benefit from the lessons Duhem teaches.

6.1 CONSEQUENCES FOR CONTEMPORARY THEORY

6.1.1 An Alternative Model of Scientific Theory

Perhaps the most significant contribution Duhem makes to contemporary philosophy is the sketching of an alternative model of scientific theory. Most philosophers assume that science is *obviously* progressing towards some version of a dynamic laws model. Duhem challenges this assumption by providing historically accurate examples of science progressing *away* from such a model by grounding itself in variational principles rather than dynamic laws.

Non-dynamic theories, such as analytic mechanics and thermodynamics, focus on the role of equilibrium in physical phenomena. Rather than seek out ordinary differential equations which describe the *actual motion* of a system over time, these theories are based in information concerning how systems come to rest. The processes at the core of such theories are *quasi-static* — processes which jump from one static state to the next. Although such theories appear limited and idealized, they are actually able to provide us with quite accurate, applicable information about a wide range of systems. Moreover, these theories continue to be embraced in many areas of physical research.

Although these theories are based on general principles rather than dynamic laws, Duhem does not take the anti-realist position of philosophers such as Nancy Cartwright, who argue that this idealization implies the laws of physics are *not* real laws. Instead, Duhem thinks these abstract principles are *laws which are never perfectly instantiated* because of the messy complications which occur from friction.

Philosophers such as Cartwright think that true physical laws must describe the actual motion of, for example, a pendulum through space. Because the laws we use to describe the pendulum usually describe an “ideal” pendulum which is isolated in space, rather than an ordinary pendulum in the world and the frictional forces which influence its path, they are not actually *laws*. By contrast, Duhem suggests that these idealizations *are* the laws which govern systems, but because the laws are instantiated in a messy world with much complexity (such as friction), we never see systems perfectly obey them. Still, this does not imply that they are not the laws which govern the system.

Perhaps Duhem's point is best illustrated by an analogy from football¹. The game of football is played, in part, by the execution of a series of "plays." Plays are detailed strategies a team intends to deploy in order to move the ball towards the end zone. The execution of such plays is influenced by all sorts of factors: the ability and training of the competitors, the conditions of the field, the coach's relationship with his players, the mindset of the athletes during the game, the weather, etc. To provide a detailed causal description of the sequence of events which occur in any actual play would include information about how all such factors directly effect the play. This type of description would be extremely difficult, if not impossible, to provide. More importantly, the possession of such an elaborate story would fail to give us the formulation of a team's intended play as it is written in their playbook. Actual plays always deviate from intended plays, even if the deviation is very slight. A dynamic laws picture provides a description of the deviation, but it cant tell us anything about the the form of the team's strategy.

But surely the form of the play has an important, reliable role in explaining the execution of each instance of the play, even if such instances never go quite according to plan (due to the interference of other players, slips on the field, etc.). Knowledge about the structure of the intended play is, in many ways, more insightful than knowledge about the details of its many executions.

Duhem provides several cases where, historically, science did not progress in a manner consistent with the dynamic laws approach. In fact, Duhem at times seems to argue quite dogmatically that science will not ever reveal true dynamic laws. I think this aspect of Duhem's philosophy is perhaps too strong. Nevertheless, I think he's right to *question* the assumption of many philosophers that scientific progress always follows a particular route.

We can't be sure how scientific theory is going to evolve — it is an empirical question how scientific theory will progress and what it will or will not be able to tell us. However, it is a mistake to suppose that science "obviously" progresses towards mechanical, fundamental, dynamic laws which all phenomena can be reduced to. One of the most important lessons in Duhem, is that many of these paradigmatic cases of reduction, are not actually reductions (as in the case of thermodynamics and statistical mechanics) or cannot be formulated purely

¹Thanks to Sean Kelsey for first drawing this analogy.

by dynamic laws (as is the case of classical mechanics).

This being said, it's both an open question how scientific theory will progress, and so we must take caution when exploring the metaphysical implications of physics.

6.1.2 Theories of Causation

In Chapter 2, I discussed several theories of causation in contemporary philosophy of science which suggest scientific theories depict “causal processes.” Such views assume that our most fundamental scientific theory will contain processes described by dynamic laws, and it is the philosopher's job to determine why some of these processes are considered causal. For example, Salmon and Dowe have argued that certain processes count as causal because we can track the conservation and exchange of critical quantities.

However, Duhem's alternative picture of science presents several worries for such views. First of all, continuous processes cannot be depicted within thermodynamic state space. While thermodynamics admits state variables, and is built on the concept of the conservation of energy, we cannot track this energy through any kind of dynamic process. Still, thermodynamics is able to inform us about entropy, which plays a significant causal role in the behavior of thermal systems. How can we describe this notion of causation when our theory does not permit us access to the kind of processes Salmon and Dowe describe?

Additionally, consider the virtual modifications that lie at the heart of Lagrange's mechanics. Again we do not see causal processes, but systems which aim for an equilibrium that is defined by virtual modifications. Because these modifications necessarily occur “outside of time,” it doesn't make sense to try and describe them with dynamic laws.

Although causal process theorists aim to accurately represent what science tells us about causation, rather than understand the notion of cause through a type of armchair analysis, these views fail to take into account a wide range of physical theory that is central to scientific practice. Not only do Duhem's observations call into question the basic assumptions of these causal process theories, they also seem to point towards a second conception of causation in science.

Not all accounts of causation are causal process theories. Since the mid twentieth cen-

ture, there has been a rather unorthodox analysis of causation by manipulability theorists². Manipulability theories, sometimes called interventionist approaches, focus on the practical connection between the notion of cause and an agent’s ability to manipulate circumstances. Early arguments in this vein suggested a connection between causation and agency. For example, von Wright claims “...to think of a relation between events as causal is to think of it under the aspect of action” [von Wright1971, 74].

Many objections have been raised against manipulability theories because they appear too *anthropomorphic*. Our conception of causation appears to stipulate that causal relations are “out in the world” and not entirely dependent on us or our abilities. Plenty of relationships we view as causal, such as the moon and the tides, are certainly not anything an agent would be capable of manipulating.

An interesting and effective response to this objection is to define the manipulability relationship in a non-agentive way. One popular champion of this approach is Judea Pearl, a computer scientist turned philosopher, who associates the causal relationship with properties of directed graphs [Pearl2000]. James Woodward follows this approach, providing a counterfactual analysis of causation which is also deeply connected to models of directed graphs.

This more objective interventionist approach suggests that causes are related to potential interventions on the nodes of directed graphs. Causal relationships can be defined based on how a change in one node of the graph would have an effect on the other nodes. Loosely speaking, Woodward claims that X causes Y if and only if an intervention on X causes a change in Y (where these interventions are defined by counterfactual relationships between quantities)³⁴.

Interestingly, the virtual modifications that Duhem defines strongly support these interventionist approaches to causation in physical theory. Causal process argue that their view of causation is grounded in science, while interventionist theories focus too strongly on a

²For an excellent treatment of manipulability theories, see [Woodward2013].

³Interventions need not be humanly possible, a mere theoretical intervention does the trick [Woodward2003, 91].

⁴It should also be noted that Woodward very carefully articulates his conception of intervention and spends much time explaining how his theory effectively distinguishes causes from correlations, allows for partial and direct causes, etc., all of which I here gloss over [Woodward2003, 47-61].

subjective approach to the notion of cause. However, virtual modifications, such as those at the core of analytic mechanics and thermodynamics, fit much better with a manipulationist account of causation.

Recall how Duhem defines a virtual modification:

To impress upon the variable quantities that characterise the state of the system some infinitesimal alterations allowed by the constraints is to impose upon the material system a virtual modification [Duhem2011].

Virtual modifications are infinitesimal “wiggles” that theoreticians can apply to the mathematical representation of a system in order to determine how a change in one quantity effects the other. In analytic mechanics, these virtual modifications are used to determine equilibrium states, which allows for a unified treatment of statics. The mathematical experiments which serve as a foundation for theories based on variational principles exemplify the type of counterfactual dependencies that Woodward and Pearl describe.

To summarize, not only do Duhem’s insights raise serious worries for causal process theories, they also point towards a scientific underpinning for interventionist accounts of causation.

6.2 MIXED LEVEL THEORIES

Another interesting conclusion we can draw from Duhem, which is quite relevant to contemporary philosophical debate, is the dependence of microlevel theories on macrolevel theories. Dynamic laws strategists suspect that we can reduce all phenomena to information about fundamental pieces of matter and the dynamic laws which describe their motion. But, at least in the case of the reduction of thermodynamics to statistical mechanics, it does not appear that these dynamic laws are wholly independent of constraints provided by macro level information such as temperature and entropy.

Robert Batterman has argued on several occasions that this standard, reductionist picture of philosophy fails to account for the way microlevel theories interact with macro concepts. In his book *The Devil in the Details: Asymptotic Reasoning in Explanation, Re-*

duction, and Emergence, Batterman argues that seemingly fundamental theories often rely on concepts which are defined by higher level theories. For example, he argues that the explanation of rainbows based on the wave theory of light still relies on certain geometrical concepts from the classical theory of light in order to explain why we are interested in certain asymptotic behavior [Batterman2002].

This type of dependence can be seen in the relationship of thermodynamics and statistical mechanics. Recall Helmholtz’s work on monocyclic systems. It was Helmholtz’s goal to pick out a class of models which could, on a fundamental level, represent thermodynamic behavior. Although Helmholtz succeeds in finding this class of models, it is unclear what makes this class *significant* from a dynamic laws perspective. Were we to view the world from a strictly dynamic laws perspective, without any knowledge of thermodynamic phenomena or the concept of entropy, we would have no reason to define this class of models as a class at all. It is only by reflecting this higher level information back into our lower level theory that we are able to describe thermal phenomena from a dynamic laws perspective.

Moreover, as Duhem points out, Helmholtz’s “reduction” is based in concepts from analytic mechanics — such as the energy — which do not appear in a strictly dynamic laws approach to vectorial systems. Additionally, Helmholtz’s theorem takes into account macro properties — such as the volume — treating them as *constraints* on the system of interest. In fact, the notion of a constraint in classical mechanics is quite helpful in illuminating just what sort of dependency exists between thermodynamics and classical mechanics.

On Duhem’s picture, which he illustrates through Helmholtz’s models, higher level theories factor into lower level theories as “constraints.” We use information about what happens higher up to constrain which models in a theory that we are interested in. It is not the case that these higher level theories are straightforwardly reduced — their results still play a role in determining which solutions of the classical equations are relevant. As Duhem says, these macro concepts are “illustrated” by our lower level theories — we incorporate higher level results into our lower level theories, *in the same way we might, at times, factor lower level information into our higher level theories*⁵.

To summarize, Duhem challenges the dogmatic assumption that fundamental theories are

⁵For further treatment of the role played by scales see [Wilson2011].

theories which describe the smallest bits of matter and the laws which govern their motion. On his picture, macro level theories like thermodynamics are a type of fundamental theory because they provide information that must be incorporated into our lower level theories in order to appropriately model experimental results.

6.3 PHYSICAL INSIGHT

Classifying Duhem as an anti-realist unfortunately glosses over his insightful comments about the way in which science reveals deep physical insight. While it's true that Duhem denies the dynamic laws strategy, claiming that our human limitations prevent us from getting down to the bottom of things, and asserts that scientific theories are not metaphysical theories, he also insists that our best physical theories converge on natural classifications; such theories are reflections of the classificatory order that would be provided by a "perfect theory" or explanation. What's especially novel in Duhem is his characterization of *how* science latches on to the world.

In contemporary philosophy, the notion of scientific progress is heavily influenced by strong commitments to dynamic laws views. In the first Chapter I discussed Railton, who argues that scientific explanations approximate an ideal DN text. This suggests that science only progresses when it approximates such a text: theories are more or less fundamental depending on how closely they emulate the ideal DN text. On this view, science advances linearly, by digging deeper and deeper into the material nature of things; we learn more about the world as we learn about how its smallest pieces work.

Yet one of the most interesting and influential discoveries in scientific theory — entropy — breaks the mold of what scientific progress looks like. Entropy was not discovered through an understanding of the nature of heat. Thermodynamic theory does not probe the essential nature of heat, telling us what it is composed of or what sort of metaphysical status it has. Instead, Carnot formulated a highly abstract model of an ideal heat engine which models the behavior of heat in actual systems. Through a study of this model, Clausius formulated the quantity known as entropy, in order to understand why the machine was so efficient.

Of course, in the vein of Reichenbach, one might argue that this insight conflates context of discovery with context of justification. While it's true that we discovered entropy by means of an abstract model, insofar as we discovered some interesting fact about the world, what we discovered was a fact about dynamic laws and the behavior of lower level systems.

But based on our current knowledge, this is just false. Entropy, like temperature, is a property that can only apply to complex systems rather than individual atoms. Although it is sometimes possible to reflect our macro level understanding of entropy back down to constrain the behavior of lower level particles, this by no means suggests that entropy can be a property of an atom. In this sense, physical theory points towards the importance of macro properties of systems. Additionally, the fact that entropy always increases is *multiply realizable*. Systems obey entropic constraints independent of their microscopic constitution. If this is the case, it seems strange to say that entropy is a fundamental property of atoms.

Moreover, entropy is impossible to observe from a dynamic laws perspective. Though entropy is a difficult quantity to understand, it is roughly considered a measure of thermodynamic work that a system is capable of doing. For example, if a glass of water is at thermodynamic equilibrium, it cannot engage in any thermodynamic processes (unless it is disturbed). If we drop ice in the water, however, the system is capable of performing thermodynamic work (i.e. the melting of the water).

Entropy is a reflection of the energy that cannot be used to perform work. When Clausius originally introduce the quantity, he understood it as the “waste heat” of the system, or the heat of the system that was not available to be used as energy. More specifically, he took it to be a measure of *heat loss*. Entropy was originally formulated as the heat lost by a Carnot engine — an idealized engine which moves through a series of equilibrium states. Were it the case that we described, according to dynamic laws, the behavior of an actual heat engine, this engine would surely lose a great deal of heat and energy to frictional forces. All engines lose a great deal of energy to friction. As machines become more and more efficient, they lose less and less energy. This is why Carnot suspected that his engine would not lose any heat whatsoever — it was fully isolated from the effects of friction. However, as Clausius articulates with his formulation of entropy, it turns out that even without friction there is a certain amount of energy that cannot be converted into work — and that “unusable work”

is the entropy. But from a dynamic laws point of view, it is impossible to distinguish the energy lost due to entropy from the energy lost due to friction — another reason why the thermodynamic perspective is indispensable.

Of course, it is possible that at some later time we will discover regular behaviors at an even lower level of matter than is currently accessible to us. And, of course it's possible that this lower level of matter will obey laws which explain the nature of entropy. Perhaps this lower level of matter will not experience frictional effects, and it will be quite clear how and why entropy comes in to play. However, this hypothesis, albeit theoretically possible, seems more like an unsupported conjecture than an accurate prediction based on the way science has evolved so far. While few philosophers presume that the micro mechanical laws we currently possess (e.g. quantum mechanics) are the final story, many suppose that we are “on the right track,” and that when we find a lower level theory, we will be able to cleanly build up our macro level physical theories. This belief is based on the observation that science always progresses by discovering smaller bits of the universe and general laws which describe them.

But Duhem's examples radically challenge this view of scientific progress. He points towards theories which are able to become more unified and more general by *ignoring* microscopic details. He even shows us how, in the case of thermodynamics, we gain important physical insights such as entropy by ignoring the micro level models of our system. Additionally, he shows us that scientific theory often engages in abstraction in order to get around the muddling effects of frictional forces. It is only by looking at these abstract, general regularities that we can see “beyond the veil” to how systems would behave if they weren't consistently attacked by the difficulties of frictional forces. Considering what Duhem emphasizes in the history of physical theory, it is unclear what basis we have to expect that scientific theory will ultimately reveal to us the basic building blocks of the universe and the simple laws which govern them, and then go on to explain how our other theories can be reduced to this micro foundation.

Of course, one can reasonably object by pointing out that science often does manage to successfully uncover dynamic laws, and I have here ignored all such case studies. I certainly don't want to deny that in certain branches of science, practitioners do aim and at times

successfully uncover the dynamic laws which describe a particular phenomena. However, I want to emphasize that this is not *always* the case, and there are plenty of other methods which scientists employ in their development of scientific theory.

As I have previously noted, even Duhem describes the wave theory of sound as an example of a theory that provides an “explanation.” What Duhem denies, I think, is not that theories can ever provide apparent explanations, but that one single theory will provide a straightforward and unified explanation of everything. Now, I think even this might be too strong — given that we really have no idea what nature is going to tell us or in what direction science will proceed. Instead, I want to argue that we have no idea what structure scientific theory will take next, and the theories that we have now are not clear, straightforward dynamic laws models. Instead, many of the most central pieces of physical theory involve variational principles and quasi-static processes.

Nevertheless, our physical theories do provide us with genuine insight, even if they do not always provide us with straightforward dynamic laws.

6.3.1 An Objection Based on Equivalent Formulations

Another possible objection for this line of argumentation is as follows: The vectorial description of classical mechanics and the Lagrangian formulation of classical mechanics are *equivalent* theories. Therefore, it’s wrong to think Lagrange’s theory is “superior” to, or an advancement from, the point-mass formulation. Why not just suppose that Lagrange’s mechanics is true *because* the point-mass formulation is true? Sure, Lagrange provides us with methods to help make calculations simpler, but we can show that it’s mathematically equivalent to the vectorial approach. If this is the case, the vectorial approach is clearly the more “metaphysically accurate” theory, and Lagrange’s theory is only superior with respect to its usefulness, but not in terms of its physical insight.

One way to respond to this objection is to argue that Lagrange’s formulation is in fact superior to the vectorial approach. I have already argued something like this — or, I have at least argued that Duhem had good reasons for believing something like this — in Chapter 3. Here, I respond to this objection in a different way. While I certainly think it’s a mistake

to take the mathematical equivalence of the vectorial approach and the analytic approach as full physical equivalence, let's suppose, for the sake of argument, that these two theories are completely equivalent.

Now, the vectorial approach to mechanics suggests that classical behavior is derived from the basic laws of motion which govern the discrete, rigid, extensionless points we call atoms. By contrast, Lagrange's theory suggests that a particle will always travel the path of minimum energy: that is, a particle's path is determined not by some dynamic laws which govern it but by nature's insistence to conserve energy. The point then is, these two equivalent theories offer radically different metaphysical perspectives. If the theories are genuinely equivalent, what reason do we have to assume the metaphysical content of one is more significant than the metaphysical content of another.

For example, why should we assume that particles always take the path of least energy *because* that is what is stipulated by the dynamic laws, rather than assuming that the dynamic laws are what they are because they model nature's strict rule that particles must always travel the path of least energy? On what basis can we argue that the vectorial approach is "more fundamental," or that the analytic approach is grounded in the point-mass formulation? Given that the theories are equivalent, and each one offers a "metaphysical reason" why particles behave as they do, how do we bestow one or the other with metaphysical priority?

One way to reconcile this metaphysical incongruity is by suggesting that the different formulations offer different perspectives of the same metaphysical phenomenon. While analytic mechanics provides us with a global approach to the laws of classical motion, which is based on the overall energy used by the system rather than dynamic laws which govern its parts, the vectorial approach is a local approach, describing the behavior of the individual atoms. But insofar as we use this sort of reasoning to ease the philosophical tension between the two views, we must concede that neither theory provides a full metaphysical explanation of classical motion. Both might provide physical insight, and both might tell us different ways of conceiving of the laws of classical motion, but neither gives an outright *explanation* of what *causes* physical motion. Even if we assume these theories are entirely physically equivalent, Duhem's worry still stands.

6.4 FINAL THOUGHTS

In this dissertation, I have considered the distinction Duhem makes between natural classifications and explanations. I have argued that much of the exegetical confusion surrounding this distinction arises from our current conceptions of how scientific theories are structured. A large number of contemporary philosophers assume that physical science adheres to a dynamic laws model: it provides us with laws that describe how the basic bits of matter evolve over time. However, many important pieces of physical theory do not follow the dynamic laws model. Instead, they invoke what I have called “quasi-static reasoning,” a type of abstract modeling based on information about various equilibrium states of systems. I have shown that Duhem was quite familiar with this alternative model of science, and that such models provide the undergirding for the seemingly contradictory claims he makes about scientific theory.

I have further argued that Duhem’s project is not antiquated: the issues at stake are not of merely historical interest. Duhem provides an alternative prototype for how scientific theory latches on to the natural world and provides us with genuine physical insight. I began by suggesting that different philosophers have placed Duhem all over the realism/anti-realism continuum. Perhaps this is because our understanding of scientific realism is difficult to interpret outside of the dynamic laws model of physical theory. The realism/antirealism debate, along with much of the discussion of grounding in metaphysics, relies on a dynamic laws conception of physical theory.

But, in truth, we have no idea how physical theory will shape up. Thermodynamics and Analytic Mechanics are at least two examples of how science can progress by moving *away* from a dynamic laws strategy or ideal DN text. Of course, there have been plenty of important discoveries in the history of science which came in the form of dynamic laws. Given that science moves in such mysterious ways, it’s an empirical question how scientific theory will evolve and what, if anything, our theories will reveal about a “fundamental level” behind the veil of appearances. Nevertheless, it is important for philosophers to look at these cases in detail, in order to develop an understanding of the different ways scientific theories articulate natural insights. Only by a serious, careful examination of the history of science

and the technical intricacies of physical theory — the sort of study modeled for us by Pierre Duhem — can we begin to sketch a portrait of how nature reveals herself.

BIBLIOGRAPHY

- [Albert2003] Albert, David Z. 2003. *Time and Chance*. Harvard University Press.
- [Ariew2008] Ariew, Roger. 2008. Pierre Duhem. In *Stanford Encyclopedia of Philosophy*.
- [Batterman2013] Batterman, Robert. 2013. *The Oxford Handbook of Philosophy of Physics*.
- [Batterman1998] Batterman, Robert W. 1998. Why Equilibrium Statistical Mechanics Works: Universality and the Renormalization Group. *Philosophy of Science* 65 (2): 183–208.
- [Batterman2002] Batterman, Robert W.. 2002. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press.
- [Batterman2009] Batterman, Robert W.. 2009. Idealization and Modeling. *Synthese* 169 (3): 427–446.
- [Batterman2011] Batterman, Robert W.. 2011, June. The Tyranny of Scales.
- [Belot2005] Belot, Gordon. 2005. Whose Devil? Which Details? *Philosophy of Science* 72 (1): 128–153.
- [Bisgaard2014] Bisgaard, Morten. 2014. pith ball electroscope.
- [Boltzmann1962] Boltzmann, L. 1962. Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. In *Wissenschaftliche Abhandlungen*, edited by F. Hasenohrl, 498–500. New York: Chelsea Publishing Company.
- [Brenner et al.2011] Brenner, Anastasios, Paul Needham, David Stump, and Robert Deltete. 2011. New Perspectives on Pierre Duhem’s *The Aim and Structure of Physical Theory*. *Metascience* 20 (1): 1–25.
- [Callen1985] Callen, Herbert B. 1985. *Thermodynamics and an Introduction to Thermostatistics*. Wiley.
- [Callender2001] Callender, Craig. 2001. Taking Thermodynamics Too Seriously. *Studies in History and Philosophy of Science Part B* 32 (4): 539–553.

- [Campisi2005] Campisi, Michele. 2005. On the mechanical foundations of thermodynamics: The generalized Helmholtz theorem. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 36 (2): 275 – 290.
- [Capecchi2012] Capecchi, D. 2012. *History of Virtual Work Laws: A History of Mechanics Prospective*. Science networks historical studies. Springer Milan.
- [Caratheodory1909] Caratheodory, C. 1909. Untersuchungen über die Grundlagen der Thermodynamik. *Mathematische Annalen* 67:355–386.
- [Caratheodory1925] Caratheodory, C.. 1925. Über die Bestimmung der Energie und der absoluten Temperatur mit Hilfe von reversiblen Prozessen. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, pp. 39–47.
- [Cartwright1983] Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford University Press.
- [Cartwright1989] Cartwright, Nancy. 1989. *Nature’s Capacities and Their Measurement*. Oxford University Press.
- [Clausius1864] Clausius, Rudolf. 1864. *Abhandlungen über die mechanische Wärmetheorie*, Vol. 1. Braunschweig: F. Vieweg.
- [Clausius1867a] Clausius, Rudolf. 1867a. *Abhandlungen über die mechanische Wärmetheorie*, Vol. 2. Braunschweig: F. Vieweg.
- [Clausius1867b] Clausius, Rudolf. 1867b. *The Mechanical Theory of Heat: With Its Applications to the Steam-engine and to the Physical Properties of Bodies*. J. Van Voorst.
- [Clausius1876] Clausius, Rudolf. 1876. *Die mechanische Wärmetheorie*. Braunschweig: F. Vieweg.
- [Clifford Truesdell1977] Clifford Truesdell, Subramanyam Bharatha. 1977. *The concepts and logic of classical thermodynamics as a theory of heat engines, rigorously constructed upon the foundation laid by S. Carnot and F. Reech*. Springer-Verlag.
- [Cohen1985] Cohen, I. Bernard. 1985. The Evolutions of Mechanics. *Physics Today* 35 (5): 86.
- [Colyvan2003] Colyvan, Mark. 2003. *The Indispensability of Mathematics*. OUP USA.
- [Darling2003] Darling, Karen Merikangas. 2003. Motivational Realism: The Natural Classification for Pierre Duhem. *Philosophy of Science* 70 (5): 1125–1136.
- [Deltete2008] Deltete, Robert J. 2008. Man of Science, Man of Faith: Pierre Duhem’s ”Physique de Croyant”. *Zygon* 43 (3): 627–637.
- [Derkleinebauer2008] Derkleinebauer. 2008. p-V diagram.

- [Dijksterhuis1956] Dijksterhuis, E.J. 1956. Archimedes. *Acta historica scientiarum naturalium et medicinalium*. E. Munksgaard.
- [Duhem1980] Duhem, Pierre. 1980. *Evolution of Mechanics*. Springer.
- [Duhem1987] Duhem, Pierre. 1987. *Medieval Cosmology: Theories of Infinity, Place, Time, Void, and the Plurality of Worlds*. Translated by Roger Ariew. University of Chicago Press.
- [Duhem1990a] Duhem, Pierre. 1990a. Logical Examination of Physical Theory. *Synthese* 83 (2): 183–188.
- [Duhem1990b] Duhem, Pierre. 1990b. Research on the History of Physical Theories. *Synthese* 83 (2): 189–200.
- [Duhem1991] Duhem, Pierre. 1991. *The origins of statics: the sources of physical theory*. Kluwer Academic.
- [Duhem2006] Duhem, Pierre. 2006. A New Theory of the Inorganic World. *Philosophical Forum* 37 (1): 3–28.
- [Duhem2009] Duhem, Pierre. 2009. *Thermodynamics and Chemistry - A Non-Mathematical Treatise for Chemists and Students of Chemistry*. Lightning Source Inc.
- [Duhem2011] Duhem, Pierre. 2011. *Commentary on the Principles of Thermodynamics*. Springer.
- [Duhem1954] Duhem, Pierre Maurice Marie. 1954. *The Aim and Structure of Physical Theory*. Princeton, Princeton University Press.
- [Duhem, Ariew, and Barker1996] Duhem, P.M.M., R. Ariew, and P. Barker. 1996. *Essays In The History And Philosophy of Science*. Hackett Publishing Company.
- [Dupré1983] Dupré, John. 1983. The Disunity of Science. *Mind* 92 (367): 321–346.
- [Dupré1998] Dupré, John. 1998. Against Reductionist Explanations of Human Behaviour: John Dupré. *Aristotelian Society Supplementary Volume* 72 (1): 153–211.
- [Einstein1905] Einstein, A. 1905. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik* 322 (8): 549–560.
- [FrostArnold2010] FrostArnold, Greg. 2010. The NoMiracles Argument for Realism: Inference to an Unacceptable Explanation. *Philosophy of Science* 77 (1): 35–58.
- [Gallavotti1999] Gallavotti, G. 1999. *Statistical Mechanics: A Short Treatise*. Texts and monographs in physics. Springer.

- [Gibbs1906] Gibbs, J. Willard. 1906. *The Scientific Papers of J. Willard Gibbs. Vol. 1: Thermodynamics*. Ox Bow Press, 1993.
- [Goddu1990] Goddu, André. 1990. The Realism That Duhem Rejected in Copernicus. *Synthese* 83 (2): 301–315.
- [Godfrey-Smith1996] Godfrey-Smith, Peter. 1996. *Complexity and the Function of Mind in Nature*. Cambridge University Press.
- [Goldstein, Poole, and Safko2002] Goldstein, H., C.P. Poole, and J.L. Safko. 2002. *Classical Mechanics*. Addison-Wesley Longman, Incorporated.
- [Hempel1965] Hempel, C.G. 1965. *Aspects of scientific explanation: and other essays in the philosophy of science*. Free Press.
- [Kutach2013] Kutach, Douglas. 2013. *Causation and Its Basis in Fundamental Physics*. Oxford University Press.
- [Lagrange, Boissonnade, and Vagliente2010] Lagrange, J.L., A. Boissonnade, and V.N. Vagliente. 2010. *Analytical Mechanics*. Boston Studies in the Philosophy of Science. Springer.
- [Lanczos1964] Lanczos, C. 1964. *The variational principles of mechanics*. Mathematical expositions. University of Toronto Press.
- [Lieb and Yngvasson1999] Lieb, and Yngvasson. 1999. The Physics and Mathematics of the Second Law of Thermodynamics. *Science Reports* 310 (314): 1–96.
- [Lugg1990] Lugg, Andrew. 1990. Pierre Duhem’s Conception of Natural Classification. *Synthese* 83 (3): 409–420.
- [Maddy2008] Maddy, Penelope. 2008. How Applied Mathematics Became Pure. *Review of Symbolic Logic* 1 (1): 16–41.
- [Martin1991] Martin, R. N. D. 1991. *Pierre Duhem: Philosophy and History in the Work of a Believing Physicist*. Open Court Publishing.
- [Matheson1998] Matheson, Carl. 1998. Why the No-Miracles Argument Fails. *International Studies in the Philosophy of Science* 12 (3): 263–279.
- [Maxwell1890a] Maxwell, J. 1890a. Diffusion. In *Scientific Papers, Vol. 2*, edited by W.D. Niven, 625–646. Cambridge: Cambridge University Press.
- [Maxwell1890b] Maxwell, J.. 1890b. On the Equilibrium of Heterogenous Substances. In *Scientific Papers, Vol. 2*, edited by W.D. Niven, 498–500. Cambridge: Cambridge University Press.

- [Maxwell1890c] Maxwell, J.. 1890c. Tait's Thermodynamics. In *Scientific Papers*, Vol. 2, edited by W.D. Niven, 660–671. Cambridge: Cambridge University Press.
- [Mitchell2008] Mitchell, Sandra D. 2008. Exporting Causal Knowledge in Evolutionary and Developmental Biology. *Philosophy of Science* 75 (5): 697–706.
- [Mitchell2009] Mitchell, Sandra D.. 2009. *Unsimple Truths: Science, Complexity, and Policy*. The University of Chicago Press Chicago and London.
- [Müller2007] Müller, I. 2007. *A History of Thermodynamics: The Doctrine of Energy and Entropy*. Springer.
- [Needham1991] Needham, Paul. 1991. Duhem and Cartwright on the Truth of Laws. *Synthese* 89 (1): 89–109.
- [Needham2008] Needham, Paul. 2008. Resisting Chemical Atomism: Duhem's Argument. *Philosophy of Science* 75 (5): 921–931.
- [Nortona] Norton, John D. *Approximation and Idealization: Why the Difference Matters*.
- [Nortonb] Norton, John D.. *Infinite Idealizations*.
- [Pearl2000] Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- [Pierre Duhem2002] Pierre Duhem, Paul Needham. 2002. *Mixture and chemical combination: and related essays*, Volume 223 *Boston studies in the philosophy of science*. Translated by Paul Needham. Springer.
- [Pierre Duhem1996] Pierre Duhem, Roger Ariew, Peter Barker. 1996. *Essays In The History And Philosophy of Science*. Translated by Peter Barker Roger Ariew. Hackett Publishing.
- [Pincock2004] Pincock, Christopher. 2004. A Revealing Flaw in Colyvan's Indispensability Argument. *Philosophy of Science* 71 (1): 61–79.
- [Pincock2011] Pincock, Christopher. 2011. On Batterman's 'On the Explanatory Role of Mathematics in Empirical Science'. *British Journal for the Philosophy of Science* 62 (1): 211–217.
- [Planck1897] Planck, M. 1897. *Vorlesungen über Thermodynamik*. Leipzig: Veit & Comp.
- [Planck1926] Planck, M.. 1926. Über die Begründung des zweiten Hauptsatzes der Thermodynamik. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, pp. 453–463.
- [Planck1945] Planck, M.. 1945. *Treatise on Thermodynamics*. New York: Dover.
- [Planck1948] Planck, M.. 1948. *Wissenschaftliche Selbstbiographie*. Leipzig: J. A. Barth.

- [Psillos1999] Psillos, S. 1999. *Scientific Realism: How Science Tracks Truth*. Philosophical Issues in Science. Routledge.
- [Psillos2013] Psillos, Stathis. 2013. *The Scope and Limits of the No-Miracles Argument*.
- [Putnam1975] Putnam, Hilary. 1975. What is Mathematical Truth? In *Mathematics, Matter and Method*, 60–78. Cambridge University Press.
- [Quine2010] Quine, Willard Van Orman. 2010. Two Dogmas of Empiricism. In *Arguing About Language*, Volume 60, 20–43. Routledge.
- [Railton1978] Railton, Peter. 1978. A Deductive-Nomological Model of Probabilistic Explanation. *Philosophy of Science* 45 (2): 206–226.
- [Railton1981] Railton, Peter. 1981. Probability, Explanation, and Information. *Synthese* 48 (2): 233–256.
- [Rankine1852] Rankine, W. 1852. On the Reconcentration of the Mechanical Energy of the Universe. *Philosophical Magazine* 4:358.
- [Renn2005] Renn, J. 2005. Einstein’s invention of Brownian motion. *Annalen der Physik* 14 (S1): 23–37.
- [Roger Ariew1990] Roger Ariew, Peter Barker. 1990. *Pierre Duhem: Historian and Philosopher of Science*. Kluwer Academic Publishers.
- [Russell2009] Russell, B. 2009. *Human Knowledge: Its Scope and Limits*. Routledge classics. Taylor & Francis.
- [Russell1912] Russell, Bertrand. 1912. “On the Notion of Cause”. *Proceedings of the Aristotelian Society* 7:1–26.
- [Russell1948] Russell, Bertrand. 1948. Volume 24.
- [Sadi Carnot1977] Sadi Carnot, E. Mendoza. 1977. *Reflections on the motive power of fire: and other papers on the second law of thermodynamics by E. Clapeyron and R. Clausius*; ed. with an introduction by E. Mendoza. Peter Smith.
- [Salmon1994] Salmon, Wesley C. 1994. Causality Without Counterfactuals. *Philosophy of Science* 61 (2): 297–312.
- [Sklar1993] Sklar, Lawrence. 1993. *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge University Press.
- [Stump1989] Stump, David. 1989. Henri Poincaré’s Philosophy of Science. *Studies in History and Philosophy of Science Part A* 20 (3): 335–363.
- [Truesdell1968] Truesdell, Clifford. 1968. *Essays in the History of Mechanics*. New York: Springer.

- [Truesdell1980] Truesdell, Clifford. 1980. The Tragicomical History of Thermodynamics 1822-1854. New York: Springer.
- [Truesdell1986] Truesdell, Clifford. 1986. What Did Gibbs and Caratheodory Leave Us About Thermo- dynamics? In New Perspectives in Thermodynamics, edited by J. Serrin, 101–123. Berlin: Springer.
- [Uffink2001] Uffink, Jos. 2001. Bluff Your Way in the Second Law of Thermodynamics. Studies in History and Philosophy of Science Part B 32 (3): 305–394.
- [von Helmholtz] von Helmholtz, H. Principles of Static Monocyclic Systems, year = 1884, url = http://neo-classical-physics.info/uploads/3/0/6/5/3065888/helmholtz_static_monocyclic_systems.pdf.
- [von Helmholtz1847] von Helmholtz, H. 1847. Über die Erhaltung der Kraft. Berlin: G. Reimer.
- [von Wright1971] von Wright, G. H. 1971. Explanation and Understanding. Volume 23. Cornell University Press.
- [Wallace2010] Wallace, David. 2010. Gravity, Entropy, and Cosmology: In Search of Clarity. British Journal for the Philosophy of Science 61 (3): 513–540.
- [Wilson2011] Wilson, Mark. 2011. Physics Avoidance. Unpublished, pp. 305–394.
- [Wilson2012] Wilson, Mark. 2012. What is Classical Mechanics Anyway? In The Oxford Handbook of Physics, edited by Robert Batterman, ? Oxford.
- [Wilson2013] Wilson, Mark. 2013. Two Cheers for Anti-Atomism.
- [Woodward2003] Woodward, James. 2003. Making Things Happen: A Theory of Causal Explanation. Oxford University Press.
- [Woodward2013] Woodward, James. 2013. Causation and Manipulability. In The Stanford Encyclopedia of Philosophy, edited by Edward N. Zalta, Winter 2013.
- [Woodward1997] Woodward, Jim. 1997. Explanation, Invariance, and Intervention. Philosophy of Science 64 (4): 41.
- [Woodward2001] Woodward, Jim. 2001. Law and Explanation in Biology: Invariance is the Kind of Stability That Matters. Philosophy of Science 68 (1): 1–20.
- [Worrall1989] Worrall, John. 1989. Structural Realism: The Best of Both Worlds? Dialectica 43 (1-2): 99–124.