

Chemical Concepts in the Era of Computational Chemistry

by

Tyler Mark Ahlstrom

B.A. Philosophy, University of Minnesota, 2011

Submitted to the Graduate Faculty of the
Dietrich School of Arts and Sciences in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

University of Pittsburgh

2019

UNIVERSITY OF PITTSBURGH
DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation is presented

by

Tyler Mark Ahlstrom

It was defended on

July 18, 2019

and approved by

James Woodward, Distinguished Professor, Department of History and Philosophy of Science

Robert Batterman, Professor, Department of Philosophy

Peng Liu, Assistant Professor, Department of Chemistry

Dissertation Director: John Norton, Distinguished Professor, Department of History and
Philosophy of Science

Copyright © by Tyler Mark Ahlstrom

2019

Chemical Concepts in the Era of Computational Chemistry

Tyler Mark Ahlstrom, PhD

University of Pittsburgh, 2019

Present work in the philosophy of chemistry has overlooked a foundational debate among chemists about the proper function of chemical concepts. The debate is fueled by a desire to connect computational models with traditional chemical concepts, and has divided chemists since the origins of quantum chemistry. By analyzing the history of the concepts of electronegativity and the atom in the molecule, I show that there are two camps with conflicting priorities. Theorists who favor rigor seek concepts that neatly summarize important elements of the underlying physical models. Theorists who favor understanding seek concepts that achieve a balance between simplicity and qualitative accuracy. The development of concepts for understanding is shown to involve the use of multiple quantification schemes in order to achieve consistency with other concepts. This practice might appear shortsighted if not for the diverse functionality of the resulting concepts. These concepts can i) help discover new reactions and structures, ii) allow comparison of different models in computational chemistry, and iii) guide chemists to develop more accurate and more interpretable computational models. Finally, it is shown that these conflicting modes of conceptual development have implications for the nature of chemical concepts. Chemists on each side of the debate adopt different positions, explicitly or tacitly, on reduction, pluralism, and the ontology of chemical concepts. Philosophers of chemistry who neglect this debate cannot responsibly interpret chemists' statements on these issues.

Table of Contents

1.0 Overview	1
2.0 The Problem of Interpreting the Computational Models of Chemistry	8
2.1 Preamble: chess concepts in the age of chess engines.....	8
2.2 Turning to chemistry	17
2.3 Brief comments on the concepts of valence bond theory	18
2.4 Brief comments on modern computational chemistry.....	23
2.5 The difficulty of quantifying traditional chemical concepts: two competing approaches	27
3.0 Two Case Studies in Chemistry: Electronegativity and the Atom in the Molecule.....	34
3.1 Introduction	34
3.2 Electronegativity	36
3.2.1 Pauling's ambiguous introduction.....	37
3.2.2 Dealing with the ambiguity: two approaches carve the boundaries	42
3.2.2.1 Concerns about the theoretical justification of electronegativity	45
3.2.2.2 Concerns about the variable electronegativity values of atoms in molecules	50
3.2.3 Electronegativity concluded	54
3.3 The Atom in the molecule	55
3.3.1 Three ways to identify the atom in the molecule.....	56
3.3.2 Debate about the correct way to identify the AIM	60
4.0 Two Approaches to the Development of Chemical Concepts	64

4.1 Introduction	64
4.2 Concepts for rigorous description	65
4.2.1 Practical motivations and hurdles to their payoff.....	67
4.2.2 Willingness to embrace conceptual modification	71
4.2.3 Philosophical motivations: definiteness & observability	73
4.3 Concepts for qualitative understanding	76
4.3.1 Practical motivations and their implications	79
4.3.2 Developing flexible conceptual networks.....	84
4.3.3 Willingness to embrace the strangeness of ‘fuzzy’ chemical concepts	92
5.0 Extending the Implications and the Future of Philosophy of Chemistry	96
5.1 Introduction	96
5.2 Misguided philosophical pursuits and the nature of the chemical bond	98
5.3 Limitations of concepts built for qualitative understanding: the strain case study.....	101
5.3.1 The barrier to rotation of ethane	103
5.3.2 Intractable questions: an implication of working with concepts developed for qualitative understanding.....	106
5.4 Auxiliary functions of concepts built for qualitative understanding.....	109
5.4.1 Concepts able to facilitate discourse about the many models of computational chemistry	110
5.4.2 Concepts able to guide the development of more accurate and more interpretable computational models	113
6.0 Conclusion	115
Bibliography	116

List of Figures

Figure 1: Positional analysis in chess	10
Figure 2: Chess engine vs. network of chess concepts	15
Figure 3: Two ways of identifying the atom in the molecule	59
Figure 4: Success and failures(?) of the topological AIM	73
Figure 5: Two ways of identifying the atom in the molecule (revisited)	87
Figure 6: Hirshfeld AIM charge test	88
Figure 7: Computational model vs. network of chemical concepts	91
Figure 8: Two conformations of the ethane molecule.	104

1.0 Overview

This dissertation is about a foundational debate in chemistry that has divided chemists into two communities. It is an enduring debate, initiated in the 1930s when quantum mechanics was applied to chemistry, and still shows no sign of resolution. The debate concerns how quantum mechanics should impact chemical theory.

That quantum mechanics should have an impact on chemical theory is motivated by its success in accurately modeling individual chemical systems. However, these models cannot answer many questions with broad and general scope. Questions about atoms and bonds, for instance, cannot be answered using these models because these models do not ‘speak’ in terms of atoms and bonds. Work must be done to bridge these models and chemical concepts, to show how atoms and bonds might be identified within the quantum mechanical models.

While pursuing this task, chemists have engaged in a divisive and longstanding debate. There are in fact many ways one might bridge these models and chemical concepts. How chemists choose among and prioritize these options depends entirely on what they are trying to accomplish. There are two fundamentally different things that different chemists seek to accomplish.

Some chemists seek to define chemical concepts in a way that is motivated by quantum mechanics. This is done by searching the physical equations underlying computational models for terms or groups of terms that behave in a manner that is similar to chemical concepts. Concepts are then defined by

these terms, and a clean connection between quantum mechanics and chemical concepts has thereby been made.

Other chemists seek to define chemical concepts in a way that maintains core elements of their traditional role in chemical theory. This includes maintaining the familiar idea that an atom's electronegativity is predictive of its charge in molecules, or the idea that the number of bonds an atom can form depends on its valence. In order to maintain concepts in this manner, proposals for bridging models and concepts are evaluated based on how well these conceptual behaviors are preserved. The motivation for this approach is that traditional chemical concepts are a simple means to qualitatively predict a wide array of molecular properties.

Thus, the two camps disagree on how chemical theory should be modified to adapt to the introduction of a quantum mechanical underpinning. This dissertation analyzes the foregoing modes of conceptual development, the nature of the resulting concepts, and the implications for philosophy of chemistry.

Chapter 1 lays the foundation for this analysis. Here, I introduce the topic of using high-performance computation models to inform preexisting conceptual practices.

I begin with an excursion into chess theory, and address the question of how modern chess engines might impact traditional chess concepts. The chess engines are powerful, but do not 'speak' in terms of traditional chess concepts. Many chess theorists have debated how these engines might still inform the development of these concepts. Some argue that we can identify our concepts with certain sets of terms used by the engines, and thereby add precision to them; others argue that the engines show the futility of traditional conceptual practices. I offer an alternative account which I call computationally

guided conceptual development. Here, the aim is to use the engine as a means of testing the efficacy of our concepts without regard for how the engine comes to its decisions. When our concepts are shown to fail, the engine may be used to seek patterns in these failures. These patterns can then be captured by heuristics about the use of our concepts, and our concepts have thereby been refined.

I proceed to more thoroughly introduce the same question for chemical theory: how have the modern computational models of chemistry impacted traditional chemical concepts? To get clear on the chemical concepts discussed throughout this dissertation, I provide a brief overview of Valence Bond Theory and several concepts rooted therein, including the chemical bond, electronegativity, strain, and the atom in the molecule. I likewise provide a brief overview of modern computational models used in chemistry, and gesture at how they construct approximate solutions to the Schrödinger equation. Together these overviews provide the means to appreciate the fact that modern computational models—accurate as they may be—do not ‘speak’ in terms of traditional chemical concepts. As with chess, many theoretical chemists have debated how these models might still inform the development of chemical concepts.

The first chapter concludes with a brief description of two approaches to bridging the gap between computational models and traditional chemical concepts. The rigor-driven approach seeks to identify concepts with important elements of the physical equations on which the computational models are based. The understanding-driven approach uses the practice of computationally guided conceptual development in order to develop concepts that achieve a balance between simplicity and qualitative accuracy.

Chapter 2 shows how the tension between these two approaches to conceptual development can account for otherwise strange issues in the history of chemical theory.

We look first at the concept of electronegativity, and how ambiguities in Pauling's seminal introduction led to debate about the quantum mechanical basis of the concept and about whether standard electronegativity values were to be assigned to free atoms or to atoms in molecules. In dealing with these ambiguities, some chemists pursued more rigorous ways of conceiving of electronegativity, while others lamented that these rigorous conceptions were more difficult to use and did not provide practical advantage. To make sense of these concerns about usability, we examine how the concept is used to qualitatively organize various chemical behavior and how Pauling's simple conception of electronegativity performs this task well.

We turn then to the concept of the atom in the molecule (AIM) and I discuss several proposals for how atoms might be identified within molecules: Mulliken's approach, the stockholder approach, and the topological approach. These approaches carve the molecule into different AIMs with different properties, raising the question of how these various approaches ought to be evaluated. Despite several multifaceted debates surrounding the concept of the AIM, in practice there are only two fundamentally different strategies used to evaluate these proposals. Some theorists evaluate conceptions of the AIM based on how neatly it fits with underlying physical models, and other theorists evaluate them based on their ability to qualitatively satisfy expectations derived from other chemical concepts, such as electronegativity. This latter evaluation procedure is shown to be a means of extending the functionality of a network of qualitative concepts, and is aligned with the practice of computationally guided conceptual development.

Chapter 3 elaborates on the two approaches to conceptual development by identifying their respective appeals and hurdles to their success.

The rigor-driven approach is motivated by the philosophical presupposition that chemical concepts ought to be physically observable, as well as by several practical considerations. The latter include developing concepts that can be applied with precision and that can be applied in a manner that is consistent across all chemical systems. I show the ability to quantify chemical concepts with precision is only valuable if there is proper alignment between the accuracy of the conceptual scheme and the accuracy of the model as a whole. Likewise, I show that the value of being able to apply a conceptual scheme universally across all chemical systems is limited by the fact that model accuracy is often a function of the class of chemistry being modeled. Finally, I show how this approach often results in concepts that behave quite differently than traditionally understood, which we observe using the concept of the chemical bond as an example. Advocates of this approach are willing to embrace these changes to conceptual behavior.

The understanding driven approach is motivated by what is often called ‘chemical understanding’. By looking at how concepts valued for this understanding are developed and used, I cash out this understanding as the ability to use concepts that achieve a balance between simplicity and qualitative accuracy across important areas of chemistry. Qualitative accuracy is shown to be assessed against two different kinds of data: observable chemical properties, such as bond dissociation energy or molecular polarity, and other chemical concepts. It is often not possible for concepts to be qualitatively accurate in both respects for all classes of chemical systems, and for this reason advocates of this approach sometimes adopt a pluralistic attitude about their concepts. Concepts can be redefined for different classes of chemical systems, and there is no singularly correct way to define concepts. Advocates of

this approach are also willing to make use of ad hoc reference structures to quantify their ‘unobservable’ concepts in a way that serves their aims.

Chapter 4 explores future directions for philosophy of chemistry in light of the debate among chemists about the proper function of chemical concepts.

I first critically evaluate an ongoing literature in the philosophy of chemistry that aims to establish the ontological status of the chemical bond. Chemists on each side of the debate adopt different positions, explicitly or tacitly, on philosophical issues including the ontology of chemical concepts. Because philosophers of chemistry neglect this debate, they cannot responsibly interpret chemists’ statements on these issues, and cannot hope to achieve an answer that is consistent with chemical practice.

More broadly, I suggest a general problem with the current state of philosophy of chemistry is that there is a presumption among philosophers that the pragmatic conceptual practices adopted by many chemists means that the philosophical questions are being passed over. Instead, I suggest these practices are rich with philosophical implications and deserve the attention of philosophers of chemistry.

In particular, I suggest there is important philosophical work to be done exploring what concepts built for understanding can and cannot do. I begin this exploration by broaching a longstanding debate among chemists about whether strain or hyperconjugation is more responsible for the barrier to rotation in ethane. I argue that concepts built for understanding are unable to provide a definitive answer to this question because this would require a reliable quantitative functionality of the concepts that is sacrificed during the course of their development.

On the other hand, concepts built for understanding are able to perform a variety of functions besides organize chemical data in way that makes simple qualitative patterns accessible. They provide a basis for evaluating model performance in a way that transcends the details of individual models, and in this way facilitate discussion about when a model will be more or less reliable than others. For the same reason, these ‘fuzzy’ concepts provide a useful guide for developing patchwork corrections to existing models. Finally, they are also able to play a key role in the development of computational models that seek to bypass the problem of model interpretability by building these ‘fuzzy’ concepts into the computational procedure at the very beginning.

2.0 The Problem of Interpreting the Computational Models of Chemistry

This chapter introduces the problem of quantifying chemical concepts using computational models and serves as a broad introduction to the dissertation as a whole. I begin with an analogy from the realm of chess, where high-performance models are also widely used and pose similar problems for those seeking to use these models to quantify their traditional chess concepts. Upon returning to the realm of chemistry, I briefly introduce the chemical concepts that will be of interest later, namely electronegativity, the atom in the molecule, and strain. I then briefly discuss some complexities surrounding modern computational models in chemistry, which will help to illuminate why chemists make use of many different modeling approaches that are based on different approximations. Finally, I discuss the difficulty of connecting the chemical concepts and the computational models—that is, the difficulty of using the models to quantify or otherwise shed light on the chemical concepts. I outline two strategies for making this connection between computational models and chemical concepts, two strategies which will be engaged more thoroughly in the chapters to follow.

2.1 Preamble: chess concepts in the age of chess engines

This work focuses on chemical concepts, and their contentious survival in the era of computational chemistry. However, I would like to begin with a brief analogy that takes us into the realm of chess theory. The reason for this preamble is that the increased use of computers brought about similar problems for chess theorists as it did for chemical theorists—it is, I argue, a general phenomenon that is due to shared computational and modeling strategies in a variety of fields. Briefly examining the contours of these problems in the context of chess will prepare us for their more thorough

examination in the context of chemistry, and I hope, keep us from losing the forest for the trees. It will also prepare us for the broader discussion of conceptual thinking that will follow.

Like chemists, human chess players face complex problems of analysis. While sometimes these problems can be exactly solved—when a forced checkmate is only a few moves away, or during the endgame, when few pieces remain on the board—for the most part, the chess player must rely on something that falls short of exact solution. In such situations, players must evaluate the position of their pieces relative to that of their opponent, and to do so without necessarily having a detailed plan for how that position will unfold in the moves to come. How do chess players make these positional evaluations?¹

Consider these two reasonable moves by white (Dvoretsky and Neat 2002, p. 62):

¹ A distinction is often made between positional considerations and tactical considerations. Tactical considerations involve foreseeing the outcome of some exchange of pieces and are typically resolved by determining if the exchange is beneficial or not (which may or may not involve positional considerations). Positional considerations, the focus of the present discussion, do not involve a calculation of exchanges, and must be resolved by considering merely the position of the pieces, and often with a very tentative grasp on how the game will henceforth develop.




Figure 1: Positional analysis in chess

Two reasonable chess moves (top) and their most reasonable follow-up moves (bottom). How do expert chess players evaluate the resulting positions, and determine which position is stronger?

The rook to D2 would naturally be followed up by doubling the rooks on the D-file, whereas the knight advance to H4 would likely be followed with its movement to F5, in position to attack kingside. How do experienced chess players evaluate the relative merits of these two positions?

To find the answer to this, we must look to chess theory, the foundation of which is conceptual in a manner that will be familiar when we turn to chemistry. When facing complex positions such as these, where one cannot see a concrete route to victory, chess students and professionals alike are advised

to base their decisions on longstanding, tested concepts; in this case, a concept of key interest in *center control*. The importance of this concept was laid out in Aron Nimzowitsch's *My System* (1925), which is generally considered to have brought about the hypermodern school of chess (Averbakh 2012) and remains of central consideration to modern chess theory. In short, the idea is that a good position will, *ceteris paribus*, exert control over the central 4 squares of the board.² The importance of controlling these squares is primarily due to the opportunities it affords: one can simultaneously exert pressure on both wings, and deploy when advantageous. "Without healthy conditions in the center, a healthy position is definitely unthinkable" (Nimzowitsch and Hays 1991, p.118; translation from Nimzowitsch 1925).

In the analysis of the two moves offered above, the concept of center control is employed to favor the doubling of the rooks on the d-file (Dvoretsky and Neat 2002, p. 63). This position, more so than the knight flank, builds white's control over the center, and this control has the potential to grow significantly if white decides to open the d-file with a pawn exchange. Because of the work that has been put into confirming the importance of center control—namely, having shown it to be predictively significant for a large number of games throughout history—the player can be reasonably confident that this position is stronger than the alternative.

However, the *ceteris paribus* clause is important here; an overzealous pursuit of center control can lead to weak positions. To reconcile this fact with the concept's strong reputation, one must appreciate

² Control here is importantly distinct from occupation; the former can be measured by the number of pieces in position to attack the central squares.

that concepts come in sets. Center control is important, but so are other things, such as *material advantage* (roughly, having more pieces than an opponent) and *piece mobility* (having pieces where they are not blocked by others or restricted by the edge of the board). It is the shared purpose of the set of concepts to break down something vague but important—positional advantage—into comprehensible components that one may employ ‘on the fly’ while playing a game, at least after training. The analysis of positional advantage using chess concepts would ideally result in the correct assessment regarding who has the advantage, and perhaps even how strong that advantage is. Put more abstractly, we would like this set of concepts to be exhaustive in its assessment of positional advantage; if it leaves some ‘advantage’ out of the analysis, we would not know whether the remaining advantage favors us or our opponent, and therefore we would have some uncertainty regarding which side possesses the overall advantage. It is the job of chess theorists to continue to develop concepts in such a way that enables more exhaustive analyses in the foregoing sense.

Let this, then, serve as a brief example of the kind of conceptual thinking on which chess theory has been built. For our purposes, the story continues with the rise of chess engines, capable today of outperforming the top human players by a significant margin. The question that has since been asked, on which the core of this dissertation hinges: how might we use the high-performance computational engines to update our conceptual tools? That is, can we use the engines to determine whether concepts ought to be kept or rejected; can we use the engines to employ the concepts with greater precision, or to modify their application in some way? Or, put yet differently, can we use the engines to create a more general and more exhaustive set of concepts that we, as humans, are able to employ in our reasoning?

There is some reason for thinking this ought to be possible in the case of chess. Chess engines are not merely a feat of brute force. They do not search through all the possible moves, stopping only when a winning position is found. While they indeed search deep—often several dozens of moves ahead—due still to technical limitations and the complexity of the game, the engine must stop at some point and evaluate the position (see Ferreira 2013 for discussion). In this evaluation process, the engines employ parameters that, at the very least, may resemble traditional chess concepts, such as material advantage, center control, mobility, and passed pawns (Levy 2009). Coefficients that represent the value of these parameters are determined by training the engine over a large number of games. Can we not, at the very least, gain from a study of the chess engine an appreciation for the relative value of these different positional elements? This question has become especially pertinent as some recent chess engines, which though only performing at the level of a FIDE International Master thus far, have eschewed all deep searches and learned to rely almost exclusively on positional analysis alone (Lai 2015).

Attempts to answer this question have led to some dispute. There are those who say computational engines have made much of past chess theory irrelevant (Campitelli 2013), or less strongly, that the style of modern grandmasters resembles a computer more so than any historical school of chess theory (Sardesai and Anand 2013). This general point of view is driven by the fact that computer engines have shown traditional chess concepts to fail in so many cases, arguably diminishing them to nothing better than pedagogical tools for beginners, and not something appropriate for professional understanding or guidance. About the potential for identifying chess concepts with terms in the positional evaluation functions used by the engines, it is said that this would be highly misleading for two reasons. First, the apparent use of these concepts in conjunction with an advanced search algorithm morphs the would-be-familiar concepts into something foreign; and second, because of the

dependence among the different terms in positional evaluation functions and the fact that several—or sometimes even several hundred—unfamiliar terms are included, we cannot trust that seemingly familiar terms truly capture the concept they might resemble. We are thus in a position of being shown how to perform at a much higher level and having to eschew our traditional understanding to achieve this, if indeed we are capable of doing so.

Conversely, there are those that say human players ought to hang on to their manner of conceptual thinking, and be extremely cautious when using an engine to inform positional considerations, if not avoid them altogether (Silman 2014). The primary argument here is that, over many decades of theoretical development, chess concepts have been forged into something that works for us humans, and we will never be able to effectively think in the terms employed by the chess engines. As such, when a chess engine tells us our concepts have led us astray in a certain position, the correct attitude to have is one of determination: work must be done to refine our concepts. And it is this practice in which we are primarily interested, a practice I call *computationally guided conceptual development*.

The goal of this practice is to find a way of using our concepts to make sense of what the computer is telling us. So long as there is good reason to trust the overall assessment of the engine (in this case, that one position is better than another), we can carry out this practice without regard for how the computer came to that assessment. This is because, again, the goal is not to understand how the computer came to some conclusion, but rather to develop a way of thinking about the results that works for us, and ultimately allows us to play better chess.

So, the computer has one way of thinking about chess positions, involving many inscrutable parameters trained over many games; and we have another way of thinking about chess positions,

involving longstanding and familiar concepts like center control, developed over centuries of theorizing. These two frameworks, different as they are, share a *point of conversion*, where they are talking about precisely the same thing: positional evaluation. The practice of computationally guided conceptual development aims to make this point of conversion a boon to our preferred conceptual framework. In this way, the engine becomes a tool for the chess theorist to find work, but where the details of how the engine works have little or no bearing on how the theorist does this work.

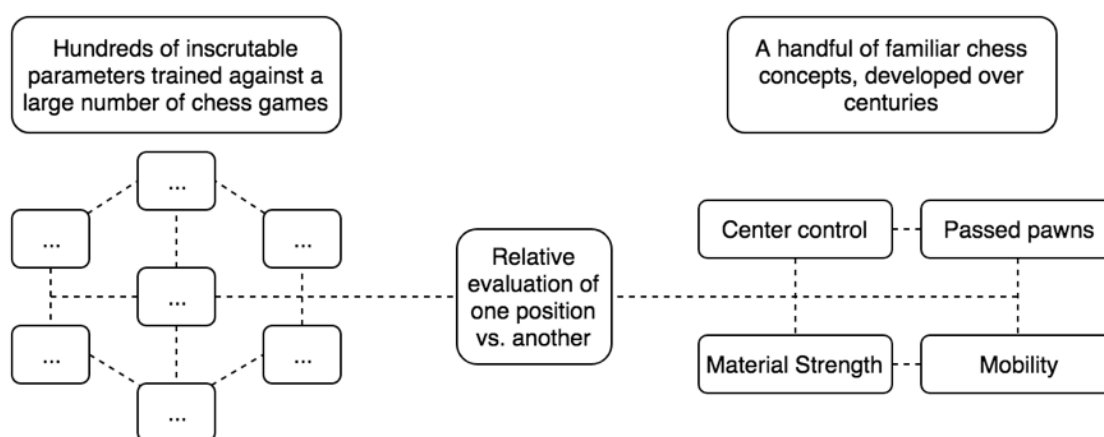


Figure 2: Chess engine vs. network of chess concepts

The computational framework of a chess engine (left) and the conceptual framework of chess theory (right) share point of conversion (center) where they are talking about the same thing.

The practice of computationally guided conceptual development can be understood as follows. Concepts, such as those that make up a large part of theoretical chess, identify rough or imperfect trends and patterns. They might be considered highly efficient, though lossy, compressions of large swaths of information—in this case, about evaluations of chess positions. These concepts are useful to us because they straddle the line between predictive accuracy and simplicity, and by doing so they push our human computational limits. When these concepts lead us astray, as chess concepts are

shown to sometimes do when confronted with results from chess engines, they are not abandoned, but rather subject to a sort of conceptual calibration. Theorists will seek patterns in these errors, patterns which might be articulated in the form of a heuristic about the application of the concept. Often these heuristics are articulated in terms of other concepts; for example, *center control* is less important in ‘closed’ games (locked pawn structure) where pieces have lower *mobility*. If these heuristics are found, theory can be said to have advanced. The computational engine provides the data, but it is left to the theorist to organize it using domain specific concepts. Note that when this task is carried out for well-established concepts, the practice may resemble ad-hoc curve fitting: the theorist pursues their task under the assumption that these concepts are worth saving, and it is difficult to imagine evidence that could convince them otherwise. As such, they can be said to be working within a conceptual paradigm in the Kuhnian sense.

So, what is the relationship between our conceptual thinking and high-performance computational models? My answer is this practice of computationally guided conceptual development. And the same practice characterizes the relationship between chemical concepts and physical models used in computational chemistry. As with chess theory, much of modern chemical theory can be traced to a several major figures in the first half of the 20th century (Mulliken 1935; Pauling 1939), who promoted conceptual thinking that would be challenged by computational models developed in the latter part of the 20th century. There too is debate about how traditional chemical concepts ought to fair in the context of highly successful computational models—should they be eschewed, connected in some explicit way to the terms of the computational models, or some other alternative? And lastly, there too we will see the controversial but rather successful practice of computationally guided conceptual development.

2.2 Turning to chemistry

Beginning in the 1920s and taking off in the 1930s, there was a rise in the amount of theory-generated data available about chemical systems. Early quantum chemistry was beginning to generate descriptions of chemical systems in terms of electronic orbitals, and at least for many simple systems, these descriptions were reasonably consistent with associated experimental data. Importantly, these descriptions also extended to chemistry that was not able to be observed in the laboratory, and even to chemistry that we might never see in the laboratory.

These computational models were more than just a new source of raw data, more than just a new instrument to the experimentalist. For not only were these models generating data that needed to be accounted for, but they were also, in virtue of their theoretical underpinning, candidate explanations for the data they generated. Why is the barrier to rotation in ethane 3.2 kcal/mol? Why is fluorine the most electronegative element? Why is cyclopropane unusually stable? Because, so those most enamored by the new quantum chemistry would claim: the theory says it is so. Demonstrations offered using clearly defined terms, and despite the approximate nature of the endeavor, the numbers that backed these assertions would only become more accurate. In this context, it is easy to appreciate the conviction with which some chemists promoted the new quantum chemistry; the physical basis for chemistry has been found, and it was generating results.

For the majority of chemists at the time, however, the descriptions generated by quantum chemistry were problematic for the reason that they were, of course, couched in terms of quantum mechanics. It was as if the models were generating answers to longstanding questions with the rider that *you must think like this to understand*. One might critically assume any problems these chemists had with the new

quantum chemistry was due to the fact that they were unfamiliar with the language and physical developments underpinning the theory, and while this was certainly true of many chemists at the time, I believe this assumption conceals a deeper and more interesting problem with these descriptions being couched in the terms of quantum mechanics. There is something to be said, I argue, about the terms and equations found in quantum chemistry not doing the same kind of explanatory work that traditional chemical concepts are capable of doing.

This idea is something several prominent chemists have raised (esp. Pauling 1939; Coulson 1961; Hoffman 1998). Typically, the argument is that traditional chemical concepts facilitate *chemical understanding*, and conversely that concepts taken straight from quantum physics—which along with their equations, may provide very accurate quantifications of chemically relevant properties—fail to bestow this understanding. Before pursuing this line of inquiry, however, we will first briefly explore several of the concepts that are said to bestow understanding, as well as the high-performance models that do not, in any straightforward way, speak in terms of these concepts.

2.3 Brief comments on the concepts of valence bond theory

This section is intended to provide a brief introduction to the chemical concepts that will be investigated more thoroughly in the following sections, and it will do so by offering a very brief introduction to valence bond theory. A semi-historical introduction to these concepts will serve as a basis for Chapter 2, which will investigate their history in more detail.

The concepts employed in Valence Bond Theory, laid out in the first half of the 20th century, remain at the heart of much theoretical dialogue in chemistry. While they have roots that reach back to the 19th century, it is Pauling's 1939 discussion of the concepts in his textbook *The Nature of the Chemical Bond* that cast them into their modern form.³ Superficially, they should be familiar to those with even a basic education in chemistry; their more technical forms will be discussed later.

Valence Bond (VB) Theory is a direct development of Gilbert Lewis' (1916a, 1916b) work on chemical valence and bonding. Prior to Lewis, it was already known that atoms of a given element typically—though not always—engage in a characteristic number of bonds, which can be inferred by their group (column) in the periodic table;⁴ for example, hydrogen, in the first group, typically forms a single bond, while carbon, in the fourth (or 14th—⁵), often forms four bonds. Lewis posited that this behavior is due to the number of 'valence' electrons an atom has, and that it is these electrons that are able to participate in bonding interactions. Bonds, for Lewis, were formed when electrons pair up; typically each atom engaged in the bonding interaction will contribute one electron to the bond. To account for the fact that elements like fluorine, which would have seven valence electrons, typically only form one bond, Lewis posted his 'rule of eight': stable molecules contain atoms with eight total valence

³ In some cases, the concepts can be traced through to the 18th century or earlier, though answering questions of origin is not the aim of this work. See Jensen (1996, 2003) for nuanced account of the origin of the concept of electronegativity.

⁴ It was largely these bonding patterns that led to Mendeley's articulation of the periodic table (Brock 1993).

⁵ 14th rather than 4th due to the transition metals, which take up ten columns but are not present in the first three rows.

electrons, where this number is the sum of their native valence and any electrons shared with them through bonding; for example, carbon, fluorine, and other main group elements, prefer a valence of eight.⁶ The foregoing principles suffice to determine reasonable chemical structures for many molecules.

These ideas, and in particular the diagrams Lewis developed to represent them, had already found widespread use among chemists when, a decade later, they were placed on more sound basis in terms of quantum physics (Heitler and London 1927; Pauli 1927; London 1928). Pauling believed this new, more rigorous basis was, “in the main equivalent to the rules which [Lewis] postulated”⁷, and presented these developments first in a series of papers titled “The Nature of the Chemical Bond” (Pauling 1931a, 1931b, 1932a, 1932b) and then in a popular textbook of the same title (Pauling 1939). He put the developments in a chemical language that would stick in the minds of chemists through to today, and which constitutes the basis for VB theory.

Among the key ideas of VB Theory are the three concepts I will examine in this dissertation: the concept of the *atom in the molecule*, *electronegativity*, and *strain*. Briefly put, Pauling showed how it makes sense to talk about the properties of the atom as it exists within the molecule, such as its charge. The

⁶ This ‘rule of eight’, later coined the octet rule (Langmuir 1919), applies only to main group atoms. Hydrogen and Helium are satisfied with a valence of 2, while lower period elements can fit 18 electrons in their valence.

⁷ L. Pauling, March 7, 1928. Lewis, G. N.: Correspondence, 1926-1945. Courtesy Ava Helen and Linus Pauling Papers, Special Collections & Archives Research Center, Oregon State University Libraries, Correspondence Box 216, Folder 216.1.

charge of an *atom in a molecule* can be used, for instance, to predict the molecule's polarity, and therefore how the molecule will orient itself in the presence of other charged entities.⁸ To determine the charge of an atom in a molecule, VB theory offers another concept, *electronegativity*, described by Pauling as “the power of an atom in a molecule to attract electrons to itself” (1939, p. 58). Pauling assigned a single electronegativity value to each element, which one could use to semi-quantitatively calculate its charge in the presence of other atoms; for example, since oxygen is more electronegative than carbon, when atoms of these respective elements are bonded, one can be sure the carbon possess a partial positive charge, and the oxygen a partial negative charge. As in Lewis' work, bonds were something that existed locally between atoms; as such, it made sense to talk about the properties of bonds, such as their length or the angles between them. Already it was recognized that lone pairs were ‘bulkier’ than bonding pairs, which could be used to qualitatively predict preferred bond angles (e.g., the O—H bond in H₂O will be more acute than the 109.5° tetrahedral). In this theoretical context, it was

⁸ The most thorough discussion of the charge of atoms in molecules can be found in the 1st and 2nd editions of his textbook (pp. 65-66), and less so in the 3rd edition (Pauling 1939, 1940, 1960). Personal archival research suggests the change in the 3rd edition is due to increased criticism of Pauling's discussion of electronegativity, which will be discussed in Chapter 2. In any case, theorizing by other chemists about the charge of the atom in a molecule (esp. Mulliken 1935), and other properties thereof, certainly does not stop due to Pauling's reservations.

straightforward to introduce the concept of *strain* to characterize molecules that—for one reason or another—are unable to reach their preferred bond angles or lengths.⁹

These concepts all had a semi-quantitative character, often being employed only to a single decimal point. The degree to which they were related to the new quantum chemistry varied, and as we will see, there will be many attempts to make them more rigorous and more precise. At the same time, they were more than a first approximation; like the chess concepts examined earlier, they would serve as a conceptual basis that future chemists would *work with* and *work to maintain*, often despite some difficulty. In this regard, the particular difficulty we will examine is their apparent lack of straightforward relevance to modern computational schemes, and the subsequent difficulty of using computational models to inform these concepts. This kind of problem is not, generally speaking, new to chemical concepts. During the latter half of the 20th century there was a dispute—now well documented¹⁰—between advocates of VB Theory and those of Molecular Orbital theory. This

⁹ For discussion of steric strain, see Pauling (1939) section 14 and 27. Note: I will not be focusing on the concept of the chemical bond, despite the fact that it has received more focus than any other concept from philosophers of chemistry. This is intentional, as later it will allow us to return to these discussions with a fresh perspective and see how longstanding debates can be shown to be misguided, without clouding our thinking beforehand.

¹⁰ See especially (Brush 1999a, 1999b; Hoffmann, Shaik, and Hiberty 2003). Related, it is interesting to note that some historians have argued Pauling became somewhat estranged from the quantum chemical community because his ideas could not be readily integrated with the culture of increasing computational precision (Gavroglou and Simões 2012). I find this narrative at odds with the fact that many theoreticians put so much time into attempts to integrate Pauling's ideas with quantum chemistry despite the difficulty of doing so.

episode in the history of chemistry displayed many of the core features we will see in our own case studies, and for many of the same reasons. For now, let me only promise that the reasons for the ‘eclipse’ and eventual return of valence bond theory become clear on my account of chemical concepts.

2.4 Brief comments on modern computational chemistry

The underlying physical laws for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.

-Paul Dirac 1929

This quote conjures mixed feelings for many chemists. On the one hand, it nicely encapsulates many of the subsequent developments in computational chemistry, which aim to meet “Dirac’s Challenge” (Marvin Cohen 2003) by devising increasingly accurate methods of approximating the aforementioned physical equations; it is for this reason that the quote has become a staple to the introductory section in textbooks on computational chemistry. On the other hand, many chemists consider the quote a “backhanded swipe at chemistry” (Bachrach 2007) for insinuating that there is little interesting theoretical work to be done in chemistry.

While it is the direct counter to this suggestion that I will ultimately pursue—that indeed, there remains plenty of interesting theoretical work to be done in chemistry—we cannot avoid some discussion of the computational details. That is, we must briefly discuss the approximate methods that have been developed following Dirac’s pronouncement. The reason we cannot avoid these details is the same

reason that theoretical chemists cannot avoid them: methods to approximate the exact solution to insoluble physical equations have become the source material for chemical theory. It is using these methods—though sometimes just the numbers they output—that theoretical chemists articulate their concepts, whether quantitatively or qualitatively. The following does not amount to a primer on computational chemistry, and those who seek such should look elsewhere (Cramer 2004; Bachrach 2007).

The insolubility of the physical equations underlying chemistry can be viewed as a twofold problem. First is the computational intractability resulting from the correlation of the particles. The form of the Hamiltonian (H), which when applied to the wavefunction of a chemical system results in its energy, is typically expressed in the following form:

$$H\Psi = E\Psi$$

$$H = -\sum_i \frac{\hbar^2}{2m_e} \nabla_i^2 - \sum_k \frac{\hbar^2}{2m_k} \nabla_k^2 - \sum_i \sum_k \frac{e^2 Z_k}{r_{ik}} + \sum_{i<j} \frac{e^2}{r_{ij}} + \sum_{k<l} \frac{e^2 Z_k Z_l}{r_{kl}}$$

where the five terms on the right quantify, in order from left to right: the kinetic energy of the electrons, the kinetic energy of the nuclei, electron-nuclei attraction, electron-electron repulsion, and nuclei-nuclei repulsion. This is an instance of the many body problem; even in a single carbon atom there are six electrons in addition to the nuclei.¹¹ Various approximations are employed to get around

¹¹ Treating the nucleus as a single particle is a practice shared throughout computational chemistry.

this computational intractability. For example, the widely adopted Born-Oppenheimer approximation assumes the nuclei are stationary relative to the electrons, significantly simplifying the second and fifth terms above. The manner of simplifying the remaining electronic terms is, for the most part, what distinguishes one computational approach from another. The Hartree-Fock method, for example, manages the correlation between the electrons by iteratively treating each electron as if it were subjected only to the mean field of all other electrons, rather than calculating all of their respective interactions. While many other methods have been devised to better capture electron-electron correlation, the approach that has gained widespread support in the past couple decades is rooted in density functional theory (DFT), where all electron correlation is handed in an approximate and computationally less expensive functional.

A second class of computational difficulties results not from the complexity of the of the Hamiltonian, but from the fact that the form of the wavefunction is unknown. A rough but intuitive way of restating this: we do not know where the electrons are in the first place, and we must have at least a guess before we can apply the Hamiltonian. The general strategy for dealing with this problem is to construct the wavefunction from a linear combination of some basis functions, and to use the ‘variational principle’ to iteratively approach a more accurate wavefunction (Cramer 2004, p.108). In short, the principle states that the actual energy of any chemical system will be lower than that of any approximate wavefunction, which affords us the strategy of proposing a wavefunction formed from

the basis functions, and checking to see whether it is lower in energy than our previous best guess.¹²

But what should be included in this basis set? Here again, matters get complicated, and vary from application to application.¹³

Thus, due to the intractability of the Schrödinger equation as well as our uncertainty regarding the form of the wave function, the actual calculations that are carried out in modern computational models in chemistry are quite messy, and stray from the models motivated by our physical understanding of these systems.

As we begin to turn to the issue of conceptual understanding in chemistry, it is important to note that the manner in which these models stray from the physically motivated equations are often bracketed for the purpose of extracting conceptual understanding. That is, due to the wide variety and frequently changing approximations made by these highly successful models, many seek to quantify traditional

¹² The variational principle does not hold true for all computational approaches, including DFT. Despite this, the iterative process seeks to minimize the energy, and the possibility of predicting a lower-than-actual energy must be considered.

¹³ Following in the spirit of Molecular Orbital Theory, the atomic orbitals of the atoms in the molecule might be used as the basis functions. For example, the wave function of a water molecule might be constructed using a linear combination of the s orbital of the hydrogen atoms, and the s and p orbitals of the oxygen atom. The constructed wave functions only resemble molecular orbitals within a certain energy range, outside of which they take on strange and unfamiliar shapes. For reasons of computational efficiency, most computational approaches make use of modified atomic orbitals with different mathematical properties, be it Gaussian type orbitals or the Kohn Sham orbitals of DFT.

chemical concepts by identifying the concept with some aspect of the ‘pure’ model, as formulated without the approximations that are made for the sake of computational tractability. For example, this is strategy of those pursuing what has been called ‘conceptual DFT’ (P. Geerlings, De Proft, and Langenaeker 2003), where all approximations are bracketed and conceptual definition begins from the true electron density (which itself can only be approximated). Despite this bracketing of approximations, it is important to understand the mess to appreciate why the conceptual definitions take the form they do; most notably, to understand why so much of the conceptual talk centers around density rather than the more familiar orbitals.

Once the ‘pure model’ has been selected, the exercise is to use its features to define the concepts that chemist want to maintain and make more precise. Sure, we know oxygen is more electronegative than carbon, but can we use the models to tell us exactly how much more? What, exactly, is the charge of the oxygen atom in cyclopropanone? How much strain is there in the molecule? More generally, as in the case of chess engines, many theorists are motivated to improve on our traditional methods of understanding these complex systems using modern high-performance models. And also again, how this ought to be done is a matter of great contention.

2.5 The difficulty of quantifying traditional chemical concepts: two competing approaches

‘The underlying physical laws [...] are thus completely known and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.’ We amend this well-known statement of Dirac to read that ‘...these laws lead to equations much too complicated to be chemically useful even if they were soluble.’

- Nicolaos Epiotis (1996) *commenting on Paul Dirac (1929)*

With the Schrödinger equation, chemistry becomes structureless; each chemical system is described by its own wavefunction

- Jansen and Wedig (2008)

Despite the significant progress made towards approximating the exact solutions of these physical equations, there remains the task of showing how these high-performance models can be used to shed light on problems that are posed in terms familiar to chemists. Continuing the example above, modern computational models have no difficulty providing a very accurate value for the total energy of cyclopropanone, but do not in any straightforward way answer questions such as: ‘what is the charge of the oxygen atom?’, or ‘what is the steric strain of the molecule’. The simple reason for this is that these questions are formulated in terms of *parts* of the whole—the oxygen atom is part of the molecule, and the strain energy is part of the total energy—and the model is not designed to be straightforwardly decomposed into these parts. This is sometimes partially due to the various approximations that are made in order to make the calculations soluble.¹⁴ However, as Epiotis suggests, even when these concerns are bracketed, it is often not straightforward how we ought to make the connection between these physical models and talk of traditional chemical concepts. How should the molecular electron

¹⁴ For example, the local-density approximation made in many applications of DFT typically overestimates the correlation energy and underestimates the exchange energy. While these offsetting errors are usually tolerated due to the good overall performance of the models, these ‘parts’ of the calculation cannot be trusted (Becke 2014).

density be decomposed into its constituent atoms? Which of the five terms in the Hamiltonian contribute to steric strain, and how much in a given instance?

In the absence of a single straightforward way to make this connection, many proposals are made. There are dozens of proposals for how one might quantify the electronegativity of the atoms using physical equations, and the same can be said of atomic charge. As above, these concepts are employed to identify and describe ‘parts’ of the molecular system, and as such, the competing proposals would lead users to identify the parts in alternative ways, quantitatively and sometimes even qualitatively. The following chapter discusses two case studies—electronegativity and atomic charge—in greater detail to show how these competing proposals emerge, and why the plurality often proves to be resilient. For now, I would like to briefly consider the general question of how these proposals ought to be evaluated. That is to say, given several two or more proposals for how to quantify the idea of electronegativity (or some other chemical concept), how should one determine which among them are better, or possibly best? What are the criteria by which they are judged? This is the central question explored in this dissertation.

I tackle this question by first analyzing the debates among chemists surrounding this issue (Chapter 2), and subsequently offering an account of conceptual development in the context of modern computational chemistry (chapter 3, and in outline here). The analysis of existing debates reveals the presence of two distinct camps among chemists, who approach the evaluations of competing proposals in categorically different manners.

The rigor-driven camp. On the one hand are chemists who aim to develop concepts that are motivated by the physical theory on which successful computational models are based. According to these

theorists, chemical concepts ought to be clearly definable within the physical equations used to accurately calculate the molecular structure and energy. The primary advantages of defining concepts in this way is that they can be assigned quantitative precision, and in a manner that is consistent across all chemical systems. During the past two decades, this approach to the development of chemical concepts has largely taken up camp in the context of DFT approaches to chemical computation, and posited the identification of features of the molecular electron density and chemical concepts. For example, such theorists commonly assert that electronegativity is nothing but the negative of the chemical potential, that the atom in the molecule is identical to basins in the electron density, and that bonds are just the path in the electron density that connects nuclei and passes through the intervening saddle point in the electron density (see Bader 2008 for a prominent example). Each of these features of the electron density—the chemical potential, the basins, and the ‘bond-paths’—is unambiguously defined in the physical equations, and behaves in ways that are similar to how the corresponding chemical concept was previously employed. Similar but not identical, and thus changes to the chemical concepts are being proposed with the goal of establishing a neater connection between the concepts and the physical models. Philosophical positions sometimes supply additional motivation for this way of developing chemical concepts.

The understanding-driven camp. Other theoretical chemists place less value in establishing a neat connection between physical theory and chemical concepts, and instead prioritize what is often referred to as ‘chemical understanding’. That is, chemical concepts ought to be defined and developed in ways that allow chemists to better understand molecular phenomena and the computational data that describes it (see Hoffmann 1998 for a prominent example). However, what this understanding consists in has never been spelled out clearly or in detail. A large part of my dissertation will be the elaboration of the precise nature of this understanding and how it guides the development and

quantification of chemical concepts. In brief, I argue that chemical concepts afford understanding when they optimize a tradeoff between simplicity and qualitative accuracy across important domains of chemistry.¹⁵ That is, chemical understanding is possessed when concepts can be employed in operations that are simple enough to be easily performed without the use of a computer calculation, and yet are qualitatively accurate across large swaths of pertinent chemistry. The following chapter contain several case studies that aim to show how these criteria interact to guide theorists of this camp in the development of chemical concepts. We will see how, in the pursuit of this chemical understanding, theoretical chemists of this second camp are willing to propose ways of defining chemical concepts that break from the physical models in ways that otherwise appear stubborn and shortsighted.

To better show what the understanding-driven camp is doing, I describe the function of chemical concepts and articulate the practice of computationally guided conceptual development. Although the evidence for these claims will be found in the following chapter, their outline can be stated as follows.¹⁶

Concepts, such as those that make up a large part of traditional chemical theory, identify rough or imperfect trends and patterns. They might be thought of as highly efficient, though lossy,

¹⁵ Here I draw on de Regt's (2017) account of scientific understanding, which also focuses on qualitative prediction. See also Hoffman's (1998) brief remarks on qualitative prediction. In regards to the tradeoff, c.f. Lewis (1973).

¹⁶ Note how well this description fits with the chess analogy, or specifically traditional chess concepts now confronted with higher performance chess engines.

compressions of large swaths of information—in this case, about energy and other properties of molecules. These concepts are useful to us because they straddle the line between predictive accuracy and simplicity, and by doing such they push our human computational limits. When these concepts lead us astray, as traditional chemical concepts are shown to sometimes do when confronted with results from modern computational models, they are not abandoned, but rather subject to a sort of conceptual calibration. Theorists will seek patterns in these errors, patterns which could be articulated in the form of a heuristic about the application of the concept. Often these heuristics are articulated in terms of other concepts; for example, Pauling’s electronegativities are effective in the context of main group chemistry, but less reliable when transition metals are involved. If these heuristics are found, theory can be said to have made advances. The computational model provides the data, but it is left to the theorist to organize it using domain specific concepts. Note that when this task is carried out for well-established concepts, the practice may resemble ad-hoc curve fitting: the theorist pursues their task under the assumption that these concepts are worth saving, and it is difficult to imagine evidence that could convince them otherwise. As such, they can be said to be working within a conceptual paradigm in the Kuhnian sense.

It is this practice of computationally guided conceptual development that defines the relationship between traditional chemical concepts and modern computational models in chemistry. The same practice relates chess concepts and high-performance chess engines, and yet additional cases where

theorists seek to use high-performance computational models to inform traditional conceptual thinking.¹⁷

¹⁷ I talk about concepts, but this research might be pursued in the context of several different topics in the philosophy of science. It touches on scientific explanation, models, concepts, and more recently, understanding. The proceeding discussion—and perhaps this entire dissertation— might be tailored to any of these literatures. I have chosen to employ the language of conceptual practices in part because it fits most closely with the pertinent discussions within the chemical literature, but also, as I hope to show in what follows, because I believe it sheds more light on these other topics in philosophy of science than vice-versa.

3.0 Two Case Studies in Chemistry: Electronegativity and the Atom in the Molecule

The central aim of this chapter is to provide a background on some chemical concepts that will be used as examples in the following chapters. I also lay the groundwork for showing how their contested development can be organized by the framework I offer, namely the division of theorists into the rigor- and understanding-driven camps. If this chapter contains a thesis, it is merely that the history of these concepts is well suited for this organization.

As such, this is primarily a standalone chapter, and readers not interested in the particulars of the history of the concepts would do fine to proceed to the following chapter, which is more philosophically oriented.

3.1 Introduction

As previously stated, many chemical concepts are attributed to parts of the molecule, such as atoms and bonds, and these parts are not explicitly recognized in the physical equations used in modern computational models. A primary job of the theoretical chemist is to generate strategies for interpreting the results of computational models using these familiar chemical concepts, and to show how the results make sense for reasons stemming from the expected behavior of the concepts. This is to say, there is a gap between our ability to precisely calculate the properties of molecules and our ability to explain the results of these calculations using chemical concepts, and it is the task of the theoretical chemist to reduce this gap.

As we will see, a common feature of the development of chemical concepts is a plethora of ways they might be quantified, and concomitant disagreements about which quantification scheme ought to be preferred. While the details vary from concept to concept, I argue that there is a general story to be told, a story that captures the common elements in the disagreements, and that reveals the conflicting motivations held by the interlocutors. In short, I argue there is a disagreement among theoretical chemists regarding the proper function of chemical concepts, and that in this regard, theoretical chemists fall into one of two camps. Here I show how these camps can be identified in the history of the concepts.

Note that here and elsewhere I choose not to focus on the concept of the chemical bond, which is the chemical concept most discussed by all parties involved, especially philosophers of chemistry. I avoid the concept of the bond in order to emphasize the generality of the trends I identify, and not because the concept poses problems for my account. In fact, debate about the nature of the chemical bond has spurred precisely the same philosophical and intransigent debates; for an example in this context, see a recent defense of the so-called unreal: “There is a future for unicorns in chemistry even after quantum theory has been introduced” (Frenking and Krapp 2007). It is my hope that focusing on general trends common to less explored chemical concepts will provide a fresh perspective that we can later use to return to discussions about the chemical bond, and shed light on why the philosophical literature dedicated to it has found little traction (Chapter 4).

3.2 Electronegativity

The concept of electronegativity—often characterized as the power of an atom to attract electrons to itself—has become a cornerstone in chemical theory, being used to rationalize the polarity of molecules, different ‘kinds’ of bonding interaction, and the larger part of acid-base chemistry, among other things. But despite the widespread use of the concept, there remains debate about the proper way to employ and understand it, and the debate spans several questions. Two questions in particular have survived nearly a century now: i) should a single electronegativity scale be privileged over others, and ii) should electronegativity be viewed as an *intrinsic* property of free atoms and ions, or rather as an *in-situ* property, defined by an atoms interaction with other atoms and changing along with the chemical environment. Following a brief historical introduction to the concept, these questions will be laid out, and I will gesture at how my philosophical framework can make sense of these longstanding debates.¹⁸

¹⁸ This does not amount to a survey, historical or otherwise. For a recent and accessible survey of many of these developments, I recommend Haghi (2012). There are many more technical takes on the same developments, but see especially Batsanov and Batsanov (2012, section 2.4) and Bergmann and Hinze (1996).

3.2.1 Pauling's ambiguous introduction

The groundwork for Pauling's seminal work on electronegativity was laid in Gilbert Lewis' (1916) paper, "The Atom and the Molecule," in which the difference between covalent and polar bonds is presented as a difference in degree. Lewis plotted the relative polarity of different classes of molecules along a smooth quadratic "without attempting to give any quantitative definition of our terms" (p.765). When the term electronegativity appeared in Pauling's notes in the early 1930s, the sole function of the concept was to serve as a means for the quantitative calculation of the degree of polarity of a bond.¹⁹ The presence of the concept of electronegativity within the chemical literature expanded significantly following Pauling's 1932 publication, published as the fourth in a series of seven papers under the heading, "The nature of the chemical bond," which would collectively earn Pauling the 1954 Nobel Prize in Chemistry and provide the foundation for his widely read textbook of the same name (1939, 1940, 1960).

To discuss the polarity of a bond in quantitative terms in his 1932 paper, Pauling needed a source of chemical data, and his choice of bond energies would for the most part be idiosyncratic to Pauling throughout the 20th century.²⁰ Pauling made the assumption that he somewhat problematically termed

¹⁹ L. Pauling, Notes, Calculations, Assorted Materials re: Electronegativity and the Electronegativity Scale, No Date [1930s]. Courtesy Ava Helen and Linus Pauling Papers, Special Collections & Archives Research Center, Oregon State University Libraries, LP Science Box 5.001, Folder 1.14.

²⁰ Though see Batsanov & Batsanov (2012) for a discussion of other minor contributors who made use of bond energies.

the additivity of bond energies: that for “normal covalent bonds”, the energy of a bond between atoms A and B (A:B) is equal to the arithmetic mean of the energies of their corresponding homonuclear diatomic molecules, A:A and B:B; that $A:B = 1/2 (A:A + B:B)$ (p. 3571).²¹ Whenever this additivity did not hold—when the bond energy of A:B was larger than predicted on the assumption that it was a normal covalent bond—then to that extent the bond possessed ionic character. Using this assumption to calculate the degree of ionic character of a bond, Pauling identified a pattern in the degree of ionic character across chemical combinations. This pattern allowed him to assign each element a rough electronegativity value, X , such that the difference between the electronegativity values of two bonded atoms could be used to calculate the approximate degree of ionic character according to the equation $\Delta_{A:B} = (X_A - X_B)^{1/2}$, where $\Delta_{A:B}$ is the additional bond energy due to ionic character, or the deviation from the bond energy expected from the postulate of additivity. The range of electronegativity values in this paper was arbitrarily set between 0 and 2 for the least and most electronegative atoms there examined, hydrogen and fluorine respectively.²²

As a method of tabulating empirical bond energies there was no doubting the value of Pauling’s insight, but Pauling’s characterization of the concept in quantum chemical terms would continue to be a source of criticism throughout his career. Pauling defined a “normal covalent bond” to be one

²¹ While “additivity of bond energies” more commonly means that the energy of a molecule was simply the sum of the energies of the bonds that composed it, here Pauling is really talking about the additivity of the energy of the bond *halves*, that the energy of the bond A:B is the sum of the energy of A· and ·B.

²² The more familiar 0-4 range would appear in Pauling’s textbook on the matter (1939). While there was no explicit justification given for the change, one can argue it makes sense based on the spread and precision of the values.

where the ionic contributions of A^+B^- and A^-B^+ possessed equal weight in the wave function describing the bond.²³ The problem was whether Pauling had the means to determine whether a bond was normal covalent independent of whether measures of bond energies satisfied his additivity assumption. His claim that the additivity postulate is supported by “a convincing body of empirical evidence” contains only a citation to an article by Pauling and Don Yost (1932), a colleague at Caltech, submitted only nine days prior to Pauling’s electronegativity paper (p. 3572). This three-page paper only discusses the bond energies of the hydrogen and alkali halides, and only so far as to show that these energies are larger than predicted by the additivity postulate on the assumption that the bonds were normal covalent. Although this is favorable evidence given that these molecules were expected to be highly polar based on considerations of crystal structure and ionization energies, no attempt is made to quantify these implicit considerations. This left the determination of the quantified degree of ionic character of a bond entirely up to the difference between measures of actual bond energies and those predicted by the additivity postulate—hardly a convincing body of empirical evidence.²⁴ The connection between electronegativity and quantum mechanics—both what this connection is as well as its importance—remained underdeveloped in Pauling’s work.

²³ This corresponds to the case where constants c_2 and c_3 are equal in the wave equation assigning electrons 1 and 2 to atoms A and B:

$$\Psi = c_1(\psi_A(1) \psi_B(2) + \psi_A(2) \psi_B(1)) + c_2(\psi_A(1) \psi_A(2)) + c_3(\psi_B(1) \psi_B(2))$$

²⁴ Oglivie (1990) has gone so far to accuse Pauling of making a tautology in this regard. However, this accusation leaves completely unexplained the widespread embrace of Pauling’s work on the concept.

Also underdeveloped was the rationale for assigning a single electronegativity value to each element. A close examination of the method by which Pauling obtains what are repeatedly called “approximate” values shows an awareness of the sensitivity of electronegativity values, as Pauling conceived them, to the chemical environment (1932, p. 3571, p. 3576). This awareness is most clear in his discussion of cases that violate the additivity assumption. Homonuclear bonds are by definition completely covalent, and yet it was clear that the strength of several homonuclear bonds were sensitive to the chemical environment. For instance, it is noted that the carbon-carbon single bond in diamond is 0.2 eV stronger than that in hydrocarbons, which would suggest that carbon atoms in diamond were more electronegative than carbon atoms in hydrocarbons. Rather than offer a method for estimating the degree of this sensitivity of electronegativity values to the chemical environment, Pauling resorts to candid remarks about the approximate nature of the singular values he was assigning each element. He writes: “I have arbitrarily chosen to ignore diamond” (p. 3574); “...it is possible to make a *rough assignment* of the atoms to positions along a scale representing their electronegativity” (p. 3576, emphasis added). Lecture notes from earlier in the same year mention deviation from additivity being only a “qualitative indicator” of the ionic character.²⁵

This talk of the approximate and practical functionality of the concept comes without any strong indication of whether the values being derived are to be applied to the atoms in the molecules (*in situ*), or to the corresponding free atoms prior to the formation of the molecules (intrinsic). One might take

²⁵ L. Pauling, Berkeley Lectures. February - March 1932. Courtesy Ava Helen and Linus Pauling Papers, Special Collections & Archives Research Center, Oregon State University Libraries, LP Science Box 2.002, Folder 2.5, p. 8.

Pauling's reliance on bond energies as evidence that these atoms are to be found in their molecular state, but this alone is insufficient. The great strength of the hydrogen fluoride bond, for instance, may be due to the difference in electronegativity between the respective free atoms that upon formation of the bond has diminished or even disappeared altogether. Five years later (1937), in his first dated notes dedicated to the calculation of electronegativity values, he again refers to the concept as a "property of an atom" without further specification.²⁶ It is likely that this distinction was not yet on Pauling's mind. Although he would soon explicitly advocate an *in situ* conception (1939) and maintain this position through to the mid-80s, for now he passed over the distinction without remark. Pauling was well aware of Robert Mulliken's (1934, 1935) work on the concept, which explicitly ascribed characteristic electronegativity values to free atoms. Pauling's 1937 notes also contain several exercises of the derivation of Mulliken's values, which largely agreed with his own, without any acknowledgement of this conceptual difference. When he took the opposing position—ascribing characteristic electronegativity values to atoms in molecules, *in situ*—he does so without any argumentation to support his divergence from Mulliken's handling of the concept, or discussion that there even was an alternative way of conceiving of the concept (1939). It was left to the broader community to make these competing positions clearer.

²⁶ L. Pauling, Manuscript Notes: "Ionic Character of Bonds", [re: the concept of electronegativity] April 11, 1937. Courtesy Ava Helen and Linus Pauling Papers, Special Collections & Archives Research Center, Oregon State University Libraries, LP Science Box 5.001, Folder 1.9.

3.2.2 Dealing with the ambiguity: two approaches carve the boundaries

Pauling's widely disseminated introduction of electronegativity became the central reference point for other work on the concept, and there was a significant amount of other contributions. These contributions often conflicted with each other either by assigning different electronegativity values to the elements—quantitatively and sometimes even qualitatively²⁷—or by assigning the concept fundamentally different behavior. I argue the source of much of these disagreements can be traced back to ambiguity in Pauling's seminal work, and that the different directions in which the concept is pulled is indicative of theorists attempting to do different things with the concept. This section begins with a summary of these contributions and the conceptual tension, and ends by gesturing how much of the subsequent debates can be made sense of by the lines drawn by the rigor- and understanding-driven camps.

Much of the fodder for debates about electronegativity comes from the plethora of ways one can reasonably measure this 'power' of an atom to attract electrons.²⁸ Soon after Pauling's 1932 discussion, several additional electronegativity scales were produced by other theoretical chemists using a variety of alternative measures. While Pauling employed measures of bond energies to derive his

²⁷ Here and generally in this thesis, 'qualitative' is used to capture 'more' or 'less' comparisons between chemical systems. For example, a qualitative difference here would be disagreement about whether hydrogen was more or less electronegative than carbon.

²⁸ For important surveys, see Prichard & Skinner (1955), Allen (1994), and Batsanov & Batsanov (2012).

electronegativity values, J. Gilbert Malone (1933)²⁹ employed dipole moments, Robert Mulliken (1934, 1935) employed spectroscopic data, Walter Gordy (1946) employed calculations of effective nuclear charge at the covalent radii and (1951) nuclear quadrupole coupling, and Robert Sanderson (1951) employed calculations of electron density. Each of these scales was in approximate agreement with Pauling's when fit to his values, with few discrepancies regarding their relative order.³⁰ In the third edition of his textbook (1960) Pauling would consider each of these investigations as evidence of the ability of his electronegativity values to roughly systematize these various atomic and molecular measurements. However, these comments would mask important conceptual disagreements between these theorists, disagreements which Pauling would publicly ignore for another several decades.

The conceptual disagreements were in regards to the what should be used to establish baseline electronegativity values, as well as how these values were affected by a change in the chemical environment. Where Pauling used molecules that exhibited a VBT single bond to establish his baseline values, Malone was more inclusive in the molecules he used to derive his values, and Mulliken, Gordy,

²⁹ This “fortuitous” scale was largely discredited by Mulliken (1935).

³⁰ Typically the scales were fit by fixing the maximum and minimum electronegativity values, and using a linear or quadratic equation to place the intermediate values. The resulting scales tended to disagree on the electronegativity values of hydrogen and metals liable to form double bonds. See Gordy (1946) for discussion.

and Sanderson relied on free atoms to establish theirs.³¹ Except Sanderson, these theorists also employed different terms to quantify the sensitivity of electronegativity values to the chemical environment. Whereas Pauling only discussed how the electronegativity of an atom would increase with positive charge and decrease with negative charge, Gordy's work using effective nuclear charge had the electronegativity of an atom increasing with an increase of double or triple bond character, while Mulliken's spectroscopic work and Gordy's nuclear quadrupole coupling work had electronegativity decreasing with an increase in p or d character of the bonding orbital. These methods of determining electronegativity changes are often at variance with each other. Increasing the double bond character of a bond will increase the electronegativity of the bonded atoms according to Gordy (1946), but may decrease it according to Pauling (1939), depending on the charge. Increasing the p character of a bonding orbital will decrease the electronegativity of the atom according to Gordy (1951) and Mulliken (1934), while it can have no effect on the electronegativity according to Pauling (1939).³² Thus while it appeared these different approaches to electronegativity calculations agreed on

³¹ Mulliken could alternatively use atoms in an unhybridized state to establish his. Note that a VBT single bond can be formed between either hybridized or unhybridized atomic orbitals, and in fact many of the simple molecules Pauling used to derive his scale consisted of bonds formed by hybridized orbitals.

³² See Batsanov & Batsanov (2012, section 2.4) for discussion

the general trend across the periodic table, their proponents held different assumptions about the source of baseline electronegativity values and how these values were affected by the chemical environment.³³

In short, there are many of different ways one might conceive and quantify electronegativity, and some confusion about which ways might be better than others, or how to go about the process of evaluating the competing proposals. As two theorists would put the confusion, “It is difficult to understand the meaning of a quantity if one does not know in what units it is to be expressed” (Iczkowski and Margrave 1961). While several different attitudes have been adopted to help make sense of the concept, two trends capture the bulk of the historical details. These trends will roughly correspond to the rigor- and understanding-driven approaches to conceptual development.

3.2.2.1 Concerns about the theoretical justification of electronegativity

One point of ambiguity in Pauling’s introduction of the concept was the importance of the basis of the concept in quantum mechanics. While the derivation of electronegativity values really only required solving a system of algebraic equations involving bond energies, he supplemented this with simple quantum mechanical reasoning. Pauling supposed that the wave function of a molecule AB could be

³³ Some theorists have argued that it was a mistake to ever assign electronegativity values to atoms *in molecules* at all (e.g., Leach 2012). Instead, such theorists argue that electronegativity should be considered a ‘transcendental’ property of atoms and that only a single fixed value should be assigned to each element, akin to nuclear charge.

approximated by the linear combination $\Psi_{AB} = \gamma\Psi(AB) + \alpha\Psi(A^+B^-) + \beta\Psi(A^-B^+)$, and that α and β would be approximately equivalent only when the electronegativities of A and B were equal. However, this connection to quantum mechanics was only stated quickly in his 1932 paper, and received an equally brief discussion in his textbook (1939). Whether this discussion is properly ‘quantum mechanical’ became a point of contention in the following decades, with both Mulliken (1940) and Coulson (1952) arguing that Pauling’s discussion was lacking in this respect, and would remain so throughout the 20th century.

The first steps along the path suggested by the rigor-driven approach was taken by Mulliken only two years after Pauling’s seminal formulation, when he suggested that the electronegativity of atoms be defined as the average of their first ionization energy and electron affinity. Mulliken started with Pauling’s supposition that the wave function of a molecule AB could be approximated by the linear combination $\Psi_{AB} = \gamma\Psi(AB) + \alpha\Psi(A^+B^-) + \beta\Psi(A^-B^+)$, and went on to argue that α and β will be equal in magnitude when $I_A - E_B = I_B - E_A$, where I is ionization energy of the respective atoms and E is the electron affinity. This condition is met when $I_A + E_A = I_B + E_B$, making the sum of ionization energy and electron affinity a way to quantify electronegativity using well defined physical terms (1934, p. 783).³⁴ Mulliken claimed to “throw new light on the physical meaning of the concept of electronegativity,” which was “obscure” and “largely empirical” in the work of Pauling (p. 782).

³⁴ Mulliken also divided the each side by two, making the electronegativity of an atom the average of its ionization energy and electron affinity. No strong reason was offered as to why division by two was included.

Whether or not Mulliken's conception of electronegativity is truly more rigorous can be contested.³⁵

But in any case, this movement of pushing the concept to be more cleanly defined in terms of the quantum mechanics of the time developed a following. From Coulson's popular 1952 textbook: "Pauling, on purely empirical grounds [proposed electronegativity values]. It was left to Mulliken to show that a better measure of electronegativity is..." (p. 133); "It should be added in this connection that Mulliken's estimates of electronegativity are fundamentally much more precise than are Pauling's, since they are based entirely on observable quantities" (p. 135).

In private correspondence, Pauling contested these characterizations of his work. While we are interested in the development of the concept and not exclusively in Pauling's contributions, it is worth taking a look at his response because it helps to identify the crux of a disagreement that persists still today.

"Page 133. Here you say 'Pauling, on purely empirical grounds, first suggested that...' The whole discussion of electronegativities in my 1932 paper was based on theoretical arguments – quantum mechanical arguments. I do not think that your sentence gives the reader the correct impression about my introduction of the electronegativity scale, the concept of the stabilization of bonds through partial ionic character, and so on.

³⁵ It may be difficult to maintain the idea that Mulliken's conception of electronegativity is more rigorous when his derivation involves "trial assumptions to be tested empirically" (1934 p. 783). Nonetheless, it is the *idea* that his conception is more rigorous in virtue of its unambiguous quantum mechanical formulation that concerns me.

[...]

Page 135. At the bottom of the page you say that Mulliken's estimates of electronegativity are fundamentally more precise than are Pauling's since they are based entirely on observable quantities. My values are, of course, based on heats of reaction, and only on heats of reaction, and these are observable quantities. It seems to me that your statement is wrong, in that it suggests that my values are not based entirely on observable quantities."³⁶

These passages raise several questions pertinent to our present inquiry (and several others about changes to Pauling's attitudes about these points).³⁷ First is the question of what constitutes a proper 'quantum mechanical argument' for the concept of electronegativity. More concretely, what fault does

³⁶ Pauling contested this distinction between his work and Mulliken's in private correspondence: L. Pauling, Correspondence: Charles Coulson. Sept. 25, 1952. Courtesy Ava Helen and Linus Pauling Papers, Special Collections & Archives Research Center, Oregon State University Libraries, Correspondence Box 68, Folder 68.4.

³⁷ The first of these objections mark a drastic change in Pauling's attitude about the rigor of the derivation of his own electronegativity values. Both the quantum mechanical argumentation and the use of bond energies displayed in Pauling's work was sustained only by the additivity postulate, which earlier Pauling was content to admit lacked rigorous justification: "[The] additivity of the energies of normal covalent bonds has no rigorous theoretical justification and cannot be expected to hold in all cases" (Pauling & Yost 1932). Beginning with this exchange with Coulson and lasting for the remainder of his career, Pauling would instead defend the rigor of his derivation against charges that his work lacked formal rigor (Pauling & Herman, 1986; Pauling, 1988). When Coulson finally responded to these matters over nine months later, he does not mention electronegativity, nor does he make a single change in the discussion thereof in the second edition of his textbook (1961).

Coulson have with Pauling defining electronegativity in terms of α and β in the approximate wave function equation? Presumably it is that they are not observable quantities, but merely terms that appear in a first order approximation of the wave function. On the other hand, bond energies as defined by heats of reaction, which Pauling actually uses to derive electronegativity values, are observable as Pauling notes, but they are empirical values rather than secure quantum mechanical terms.³⁸ Likewise, Pauling's arguments based on the stabilization of partial ionic bonds will not get him far in this exchange, since the degree of ionic character was not properly grounded in the quantum mechanical framework of the time; this is in contrast to the ionization energy and electron affinity, used by Mulliken, which can be identified with the energy of electronic orbitals.³⁹ Additionally, the use of ionization energies allowed for a proper 'zero point' that did not exist for Pauling's comparative bond energy calculations, making Pauling's conception 'relative' where Mulliken's is often said to be 'absolute' (Mulliken 1934, Gordy 1946). All of this is, at least presumably according to Coulson, gave Mulliken's conception of electronegativity a more secure quantum mechanical foothold.

This line of reasoning—focusing on the neatness of the 'fit' with rigorous quantum mechanics—is one strategy for privileging one electronegativity scale over others, and it is prevalent through the history of electronegativity (and as we will see, other chemical concepts). It is the very same approach

³⁸ It is reasonable to assume that the heats of reaction Pauling uses are actually the combined effect of several effects (once properly understood) besides bond dissociation energies.

³⁹ While the orbitals that appear in the physical models are not precisely the true orbitals due to necessary approximations (touched on in Chapter 1), they approach this as our approximations and computational capabilities improve.

used to promote the more recent conception of electronegativity as the negative of the chemical potential. In the context of the density functional approach to chemical modeling, which currently dominates the computational scene, the negative of the ‘chemical potential’ behaves very similarly to how one would expect electronegativity to behave.⁴⁰ Because this term appears directly within the mathematics used to calculate the energy of chemical systems, many prominent theoretical chemists have promoted its identification with electronegativity (Iczkowski and Margrave 1961; Parr et al. 1978; P. Geerlings, De Proft, and Langenaeker 2003).⁴¹

In the following chapter I further explore the question of what constitutes a neat fit, and why having such a fit might be desirable. First, we will explore another branch of the history of electronegativity, where I will argue theorist are trying to do something very different with the concept and accordingly bring different standards of evaluation to the many different ways of conceiving the concept.

3.2.2.2 Concerns about the variable electronegativity values of atoms in molecules

A second feature of Pauling’s discussion that received critical attention was a tension buried within his own conception of electronegativity: his insistence that electronegativity is a property of an atom

⁴⁰ The chemical potential shows up as a Lagrange multiplier when minimizing the energy with respect to electron density; see (Parr et al. 1978) for details.

⁴¹ Interestingly, this identification is sometimes often touted as vindication for Sanderson’s principle of electronegativity equalization (1955, 1983), or the idea that upon formation of a molecule, the electrons are pushed between the constituent atoms until their electronegativities become equal. Most other conceptions of electronegativity do not behave this way, including Pauling’s.

in a molecule (1939, p.58), or what would come to be called his *in-situ* conception of electronegativity, coupled with his ascription of only a single electronegativity value to each element. The reason these two aspects of his conception harbor a tension is that together they seem to imply that an atom of a given element exists in the same state regardless of its bonding interactions, an implication that contradicts the very method Pauling used to establish electronegativity values in the first place. To address this tension in the first edition of his textbook (1939), Pauling supplemented his treatment with a discussion of how the electronegativity of an atom would vary with its charge in the molecule.⁴² His estimation of this variation was candidly approximate, and no quantum mechanical reasoning was used to support these approximations. Furthermore, nowhere did Pauling elaborate on the sensitivity of electronegativity values to the chemical environment, and the discussion was dropped from later editions of his textbook (1960).

⁴² Specifically, according to Pauling's 1939 discussion, a +1 charge on an atom in a molecule will increase its electronegativity approximately 2/3 of the way to that of the element to its right on the periodic table, and vice versa, with a negative charge decreasing its electronegativity partway to the atom on its left in the periodic table (Pauling 1939).

It is likely Pauling dropped these considerations because his procedure for obtaining electronegativity values from bond energies did not easily allow for this functionality.⁴³ However, other conceptions of electronegativity did allow for this, including Mulliken's approach using orbital energies. Since atomic orbitals are appropriately sensitive to the chemical environment, determining how an atom's electronegativity changes from one molecule to another is simply a matter of calculating the energy of the appropriate orbitals. Furthermore, the widely employed Valence Bond Theory of the time offered a convenient framework for doing just that using the idea of hybridization: carbon that is sp^3 hybridized (e.g., CH_4) could then be assigned a different electronegativity value than carbon that is sp^2 (C_2H_4) or sp (C_2H_2) hybridized (Mulliken 1935).⁴⁴ Needless to say, the chemical potential, and so the DFT conception of electronegativity, would also be sensitive to the molecular environment in a precise way (though there are some hurdles in terms of identifying the atom in the molecule, explored in the following section).

⁴³ Using bond energies to quantify the electronegativity of a given element in various chemical environments would require somehow incorporating a variety of other bonds (reflecting the various chemical environments) into the calculation; however, the molecular heats of reaction used by Pauling to obtain bond energies for simple diatomic molecules does not neatly extend to more complex molecules. In particular, there is no recipe for decomposing the molecular heat of reaction into the constituent bonds without relying on a theoretically more nuanced analysis, which Pauling sought to avoid in the first place. Additionally, his brief proposal for using atomic charge (preceding footnote) surely did not pan out with much generality.

⁴⁴ Note that the language of 'hybridization' and 'valence state' has historically been inconsistent and varies between subfields—sometimes they are equivalent and sometimes not. For our purposes they can be treated as the same.

As more nuanced procedures were developed for handling the sensitivity electronegativity to the chemical environment, concerns began to grow about the simplicity of the concept. Here is a nice expression of these concerns in a landmark survey on the concept following these developments surrounding Mulliken's conception of electronegativity:

The crucial point, as Evans [(1951)] has remarked, is that there is the possibility that every molecule must needs be treated as a special case. (Is the valence state of nitrogen the same in NH_3 , in NF_3 , and in $\text{N}(\text{CH}_3)_3$?) It seems clear that the further development of the electronegativity theory will attempt to define the range of electronegativity open to an atom. There remains the problem of choosing from the range the unique value which the atom achieves in a given molecular environment. Unless some comparatively simple criterion of choice can be found...the discomforting possibility that Evans has foreseen may become a very real one.

-Pritchard and Skinner (1955)

This passage speaks to a very different trend than those pushing to place the concept on more secure quantum mechanical grounds. There was something appealing about only having to think about a single electronegativity value for each element, and the prospect of losing that was 'discomforting'.

To appreciate this concern, it is worth looking at how the concept is employed in practice. Above it was noted that the plethora of empirical measures one might use to obtain electronegativity values was a source of some debate (which is best?)—however, in practice this can be considered a virtue. Simply by referencing an electronegativity table—even if it only assigns a single value per element—one is able to make rough predictions about an array of chemical properties, and this is precisely in

virtue of the fact that the many ways of obtaining electronegativity values are in rough agreement with each other. With these values, it is straightforward to predict molecular polarity, relative ionization energies, the strength of bonds, the location of any hydrogen bond donors, and more. While there are some errors in these rough predictions, they are for the most part qualitatively accurate, and variable electronegativity values simply amount to an unnecessary complication when the concept is employed in the foregoing manner.⁴⁵ Additionally, Pauling's values do quite well in when used for this purpose, and in some notable cases better than Mulliken's.⁴⁶ Pauling himself, in later years, evaded discussion of the 'rigor' of his conception of electronegativity by stating that the concept is clearly useful enough for this to be a non-issue (Pauling & Herman 1986, Pauling 1988). And apparently others agreed—then and now—as Pauling's values have remained in widespread use.

3.2.3 Electronegativity concluded

In the following chapter, I further explore the foregoing conceptual utility as a feature of some concepts, which I call generality of qualitative accuracy. We will also explore these opposing considerations—between finding a conception that fits well with physical theory, and finding one that allows for this simple qualitative reasoning.

⁴⁵ Again by qualitative accuracy I mean that predictions using the concept make the right comparisons: Which atom is *more* electronegative? Which molecule is *more* polar? And so on.

⁴⁶ One notable difference between the Pauling and Mulliken scales is the different ordering of hydrogen and carbon – Pauling's values have carbon more electronegative than hydrogen, and Mulliken vice-versa. In terms of predicting most molecular properties, Pauling's ordering is more appropriate.

For now, it is worth gesturing at one discussion that has been passed over in the above brief and selective history of the concept. An additional question that has repeatedly raised by theorists is *whether* a definitive scale must be identified. That is, do we need to select one measure of electronegativity as definitive of the concept, or can productive science be done with concepts that identify larger sets of measures that share a family resemblance? While many theorists have argued a single measure ought to be privileged (e.g., Parr et al. 1978, Allen 1994), other prominent theorists have argued that electronegativity is an inherently ‘fuzzy’ concept and admits no precise definition, and that this is acceptable (e.g., Coulson 1952, Gonthier et al. 2012). Notice that these two positions correspond fairly neatly with the two ‘concerns’ discussed above: those who prioritize having a neat connection to physical theory are generally inclined to seek a definite scale, whereas those who prioritize qualitative reasoning are more permissive of a plurality of conceptions of electronegativity. That is to say, organizing the literature around these two camps helps to elucidate this debate as well.

3.3 The atom in the molecule

With electronegativity behind us, we can proceed more quickly to show that many of the historical features witnessed there are present for other chemical concepts as well. Additionally, we will see how interaction between these contested concepts adds another layer to the debate.

We turn now to the concept of the atom in the molecule, or the idea that the individual atom persists after it becomes part of a molecule, and continues to possess properties of its own. Above, when we explored the idea of an atom’s electronegativity changing depending on the molecule of which it is a

part, this presumes the existence of the atom in the molecule—something that continues to exist and continues to possess the property of electronegativity. The most common property of an atom in a molecule that is discussed is its charge. For example, what is the charge of the oxygen atom in an H₂O molecule?

Determining the properties of an atom in a molecule requires a procedure for identifying it, and this is where matters become both complicated and contentious. Rather than explore the history of the concept in the manner we did for electronegativity, we will briefly examine several proposals for how this identification procedure might be carried out, as well as outline the major debate surrounding the concept that has received a good amount of attention the past two decades. Following this, I will again gesture at how my philosophical analysis (Chapter 3) will make sense of these longstanding debates.

3.3.1 Three ways to identify the atom in the molecule

The many ways to identify the atom in the molecule (AIM) can be divided into two classes: orbital-based approaches and density-based approaches.⁴⁷

⁴⁷ Though see Cramer (2004) Section 9.1.3 for additional approaches. I am setting aside two approaches discussed therein, namely what he calls the intuitive approaches and the semi-empirical approaches. These additional approaches reinforce what I will come to say about how many chemists make use of the concept of the AIM (to facilitate a conceptual network with certain qualitative functionality).

Orbital based approaches assign electrons to atoms based on the degree to which their atomic orbitals contribute to the occupied molecular orbitals. While in some sense there are as many orbital-based approaches to identifying the AIM as there are computational approaches to calculating energetics, usually only ‘natural’ bonding orbitals are used for this purpose. That is, after the molecular structure is accurately calculated by whatever means, the resulting electronic density is transformed into the more familiar atomic orbitals of Valence Bond Theory, after which assignment of the electron density is straightforward.⁴⁸ Incidentally, this approach was developed primarily by Mulliken(1932 ,1955), and the resulting atomic charges are typically called Mulliken charges. Notice that this method of identifying the AIMs presupposes that it remains sensible to talk about their atomic orbitals, even though these are not physically observable and are only used to calculate the molecular properties in an indirect way.

Density-based approaches are typically considered more rigorous because they are based on a proper physical observable: the electron density. They also have the advantage of being more easily visualized. The game here becomes developing a procedure for partitioning the electron density into the constituent atomic electron densities. The two density-based AIM procedures that have received the most attention are the topological AIM and the stockholder AIM.

⁴⁸ Alternatively, the molecular structure can be calculated using some form of *ab initio* Valence Bond Theory (discussed in Chapter 4) or a simple Hartree-Fock approach, in which case no transformation is needed. See Weinhold and Landis (2012, chapter 3). The latter option, however, leads to an undesirable sensitivity to the basis set used to calculate the molecular energetics.

According to Bader's (1990) topological partition scheme (henceforth "the topological AIM"), an atom in a molecule is identified with a three dimensional real-space region bound by a zero-density gradient in the map of electron density. All and only the electron density within these basins are attributed to the single atom whose nucleus is also within the basin, and it is the union of the nucleus and the basin of electron density that together make up the atom in the molecule. Topological AIMs are non-overlapping, separated by a definite boundary, and these boundaries often possess "sharp" edges, as in the nitrogen example below (Figure 3, bottom left).

A second commonly used partition scheme is Hirshfeld's (1977) stockholder partition scheme (henceforth "the stockholder AIM"). One identifies stockholder AIMs by assigning the molecular electron density to the individual atoms using a weighting scheme obtained from a theoretical construct called the "promolecule". The promolecule consists of superimposed *free* atoms in space with their nuclei in the location where they actually exist in the molecule, but without letting their electron density change to reflect the other atoms. One may imagine the promolecule as a Venn diagram-like theoretical entity, although the density of each atom in the promolecule actually extends over all space. In computational practice a cutoff is made, but even so, each atom will typically overlap with every other atom at every point within the promolecule. Then a weighting factor is obtained for each point in space based on the percentage of the density at that point that is due to each atom in the promolecule—that is, based on the atom's 'investment' at that point. After these weighting factors are obtained for each point in space, the promolecule is replaced with the molecular electron density, which is then partitioned into the individual atomic electron densities using these weighting factors. The more an atom 'invests' in a point in the promolecule, the greater the payout from the electron density that ends up at that point. In contrast to topological atoms, stockholder atoms are more spherical, completely lack sharp edges, and overlap with each other (Figure 3, bottom right).




Figure 3: Two ways of identifying the atom in the molecule

Two different partition schemes as applied to a single nitrogen atom in N_2 (topological AIM bottom left, stockholder AIM bottom right). In each of the three images, the two '+'s represent the nuclei in the molecule, and the contour lines display the electron density. The two images on the bottom show the electron density assigned to the atom on the 'left' of the molecule according to the two competing schemes (Adapted from Davidson and Chakravorty 1992, p. 321).

The question then arises: how should chemists choose between these two partition schemes? While these are two historically prominent partition schemes, there are many more (see Schwarz and Schmidbaur (2012) for thorough survey). The resulting AIM's possess different properties depending on which scheme is selected; this selection also affects the manner and degree to which the AIM will

change in response to changes in the chemical environment.⁴⁹ One might wonder whether there is a simple empirical test to decide which of these AIM concepts is correct, if either. The answer to this question is itself a matter of debate, as is the more basic question of whether the atom in the molecule can be directly observed. It is in debating these and related questions that theoretical chemists have waxed philosophical. After briefly identifying the contours of the debate, I will gesture at how they can be understood as the product of the rigor- and understanding-driven camps pushing the concept in different directions.

3.3.2 Debate about the correct way to identify the AIM

The question of how this evaluation and selection process ought to be carried out has spurred acrimonious and, even within the chemical literature, philosophical debate. These debates have focused on several related questions, including whether the AIM is observable in principle (Parr et al. 2005; Matta & Bader 2006), the manner in which the atom does or does not continue to exist after it forms a molecule (Sukumar 2012), whether the properties of the AIM are “determined by nature” (Schwarz and Schmidbaur 2012), and whether it is appropriate to employ a plurality of AIM partition schemes (Bader 1999, 2011; Shahbazian 2013; Eberhart & Jones 2012; Parr et al. 2005). Amidst these multifaceted debates, the AIM has even been called a “Kantian noumenon [...], a conceptual construct ultimately unknowable by observation but conceivable by reason” (Parr et al. 2005, p. 3958; see Ayers

⁴⁹ Note that, potentially unlike electronegativity, it would make little sense to conceive of the AIM as being insensitive to the chemical environment. The very premise is that that atoms continue to exist in an altered form upon formation of a molecule.

et al. 2007 for similar comments). Additionally, it has been argued that pluralism with regard to the AIM is required for “chemical understanding” without elaborating on the nature of this understanding (Bultinck et al. 2007; Gonthier et al. 2012; Shaik 2007).

Despite the many lines of argumentation surrounding the concept, there remain two general strategies for evaluating the various proposals for identifying the AIM. As above, these will correspond to the rigor- and understanding-driven camps.

On the one hand are theoretical chemists who argue for their preferred conception of the AIM—generally a *single* conception—based on its fit with underlying physical theory. This includes Richard Bader and his collaborators, who argue that all chemical properties ought to be Dirac observable,⁵⁰ and that only topological AIMs, in virtue of their satisfying Schwinger’s (1951) principle of least action, will possess such properties (Bader 1999, 2006; Shahbazian 2013). Others argue that a single proposal ought to be universally adopted—usually Bader’s system—for the purpose of creating a consistent and rigorous theoretical system (Eberhart & Jones 2012). When theorists in this camp evaluate a proposal for how to identify the AIM, they prioritize its fit with the physical models underlying modern computational chemistry; “like all true physical concepts, that is, those determined by

⁵⁰ That is, they argue all chemical properties, such as atomic charge or electronegativity, ought to have corresponding Hermitian operators that can be applied in order to obtain the value for these properties.

observation, the choice is dictated by physics and the definition of an atom in a molecule is unique” (Bader 1999).

On the other hand, many practicing chemists argue that no single proposal for identifying the AIM fully captures the content of the concept, and that doing so requires the adoption of different, mutually incompatible AIM concepts for different classes of molecules. This kind of pluralism is a byproduct of an alternative strategy for evaluating AIM proposals: seeing how compatible the proposal and resulting AIMs are with *other* chemical concepts. For example, another concept that intuitively ought to come into play in discussion of atomic charge is electronegativity. Atoms in the molecule that have a greater “power to attract electrons to themselves” (Pauling 1939, p. 58) should, other things being equal, have a more negative charge. This, anyway, is how the concept is intended to work, and this intention can be realized if the right charge partition scheme is selected. Accordingly, one approach used to select a partition scheme is to test how well it conforms to these expectations relating to electronegativity considerations. In the following chapter, we will examine this practice in more detail; for now, it is enough to note that the end result of this practice is the use of different AIM partition schemes for different classes of molecules—whatever enables other chemical concepts (e.g., electronegativity) to function as expected (e.g., see Saha, Roy, and Ayers 2009).

The concerns of these pluralists may not appear to match those seeking a single-valued conception of electronegativity; however, in the following chapter we will examine in greater detail the practice of developing flexible conceptual networks to capture simple qualitative trends, and we will see how both this pluralism about the AIM and the single-valued conception of electronegativity contribute to the process of developing these conceptual networks. In short, the flexibility of drawing from different

conceptions of the AIM is required in order to maintain simple qualitative accuracy for the conceptual network as a whole.

More generally, in the following chapter I will elaborate on the two camps which, again, correspond to the (roughly) two sides in the foregoing debates. The rigor-driven camp seeks to establish a rigorous and clean connection between the concept of the AIM and quantum mechanics, whereas the understanding-driven camp seeks to develop the concept in such a way that it facilitates qualitative prediction. Ultimately, I argue that the two camps are trying accomplish very different tasks with their concepts, and for this reason seek to develop concepts with different features.

4.0 Two Approaches to the Development of Chemical Concepts

I do not believe in the existence of parallel universes wherein their laws are either ignored or ‘bent’ to accommodate personal points of view.

-Richard Bader 2011

It is nice to know that the computer understands the problem. But I would like to understand it too.

-Eugene Wigner quoted in Shaik 2007

4.1 Introduction

In chapter 2 we saw a host of debates about the correct way to define the chemical concepts of electronegativity and the atom in the molecule. Amidst high performance computational models that do not ‘speak’ in terms of these traditional chemical concepts—or at least not in any straightforward way—there were many proposals for how we might bridge this gap and use the models to quantify these concepts. It is in the evaluation of these proposals that contentious and often philosophical disagreements regularly arise. Understanding the contours of these disagreements in a general way is the aim of the present chapter.

It is my central argument that the theorists engaged in these debates fall rather neatly into one of two camps that disagree on the intended function of chemical concepts. I elaborate on the nature of these

camps throughout the present chapter, but in short: the rigor-driven camp aims to craft concepts that possess a neat connection with the physical theory on which computational models are based, whereas the understanding-driven camp aims to craft concepts that make simple qualitative patterns more accessible. Once their respective motivations are understood, their conflicting approaches to conceptual development can likewise be understood.

Similarly, my analysis of these debates will reveal that questions about the nature of chemical concepts, pursued in the small but growing field of philosophy of chemistry, cannot be resolved without first addressing this disagreement about the intended function of chemical concepts. For the two camps develop concepts that, while they often go by the same name, possess very different features that enable them to perform very different functions; they might reasonably be said to have different ‘natures’.

4.2 Concepts for rigorous description

On the one hand are chemists who aim to develop concepts that are motivated by the physical theory on which successful computational models are based. According to these theorists, chemical concepts ought to be clearly identifiable within the physical equations used to accurately calculate molecular structure and energetics. To the extent that traditional chemical concepts ought to stick around, they ought to be identifiable in this way; accordingly, should the concepts need to be modified in order to

be so identifiable, these adjustments are welcome. And where such identifications can be found, they should be regarded as definitive.⁵¹

I call the collective advocates of this approach to defining chemical concepts the ‘rigor-driven camp’ because I believe ‘rigor’ is a fair way to summarize the various qualities these theorists seek to build into their concepts. However, I wish to avoid certain implications this name may suggest. I do not mean that they are more careful in their conceptual practices than those outside the camp, and conversely there is no reason one cannot consider the practice of the opposing camp to be rigorous, even if quite different than that here so labeled. Rather, in choosing this label, I am relying on the common notion that rigorous theoretical work in this context involves a greater reliance on mathematics and precise quantification, in contrast to the more qualitative alternative discussed in the following section.⁵²

During the past two decades, this approach to the development of chemical concepts has largely centered around density functional theory (DFT) -based approaches to chemical computation; many proposals are made for defining chemical concepts using features of the molecular electron density,

⁵¹ Some comments on my chosen terminology. By ‘identifiable within the physical equations’, I mean able to be expressed in an equation using only terms already present in the physical equations, and using them in a straightforward way. If an identification is ‘definitive’, it rules out the possibility of alternative identifications that are incompatible with or contradict the definitive identification in any way.

⁵² See Hendry (2004) for a similar dichotomy in the early history of quantum chemistry, and which relies on similar terminology.

so many that there is now a subfield referred to as ‘conceptual DFT’ (Geerlings et al. 2003, Geerlings et al. 2014; Becke 2014). As shown in the previous chapter, it is asserted that electronegativity is nothing but the negative of the chemical potential (Parr et al. 1978), that an atom in a molecule is just a basin in the electron density (Bader 1990), and that a chemical bond is just the path in the electron density that connects nuclei and passes through the intervening saddle point in the electron density (Bader 1990, Bader 2008). Each of these features of the electron density—the chemical potential, the basins, and the ‘bond-paths’—is unambiguously defined in the physical equations and behaves in a way that is similar—though not identical—to the behavior of the corresponding chemical concept as previously understood. In this section I explore the appeal and implications of these conceptual practices.

4.2.1 Practical motivations and hurdles to their payoff

There are two major practical motivations for defining concepts in the foregoing way. First, the resulting concepts can—at least under the right circumstances—be quantified with precision. Second, they can be quantified in a manner that is consistent across all chemical systems. These are two natural consequences of having the concepts based directly on the physical details on which the computational models are constructed. The resulting concepts can be quantified with as much precision as the model is able to be run, and since the models are indiscriminate with regard to the chemistry being modeled (in the sense that the same computational scheme can be used to calculate the energies of any chemical system), the concept can thereby be applied universally across organic, inorganic, or any such class of chemical systems.

However, in practice, both of these draws can encounter trouble. First, whether or not a computational model can be used to precisely quantify a concept (again, defined in some explicit way using the theory on which a computational model is based) depends on the extent to which the accuracy of the relevant theoretical details and that of the conceptual quantification are *aligned*. Sometimes treating the relevant theoretical details accurately improves the accuracy of the conceptual quantification, and interestingly, sometimes it can impair it. This is perhaps best illustrated by way of example.

Consider the concept of atomic charge, or a specific example: what is the charge of the oxygen atom in an H₂O molecule? As previously examined, there are competing proposals for how one might quantify atomic charge. For simplicity, consider two: Mulliken's partition scheme and Bader's QTAIM scheme. Briefly put, Mulliken's scheme assigns charges to atoms by partitioning the electron density in molecular orbitals according to how much the involved atoms (and specifically, their atomic orbitals) contribute to the formation of the molecular orbital, and equally splitting electron density in the 'overlap region' (Mulliken 1955). Because this partition scheme depends on the orbitals used to construct the molecular orbitals, it is sensitive to the choice of basis set, and of course, choosing any finite basis set is an approximation needed for the sake of computational tractability. Problematically for Mulliken's scheme, more realistic atomic charges are obtained when a smaller basis set is chosen

due to unreasonable charge partitioning when a larger basis set is used.⁵³ This is a problem because, unsurprisingly, using a larger basis set improves the overall accuracy of the model with respect to the system-wide energetics it calculates. Therefore, if atomic charge is defined as per Mulliken's scheme, the accuracy of the atomic charge quantification is at odds with the overall accuracy of the computational model used to calculate it. This phenomenon—of there being a tradeoff between model accuracy and our ability to accurately quantify some concept—is common among orbital-based approaches to conceptual recovery.

However, this problem is not universal. Typically, density-based approaches to quantify atomic charge, including Bader's QTAIM scheme, are not critically affected in the forgoing manner by the choice of basis set or other approximations. Again, according to Bader's proposal, atomic charge is quantified by integrating over the basin in electron density assigned to that atom (and which he also uses to define the atom altogether). Atomic charge is calculated directly from the electron density, and since the accuracy of the calculation of electron density improves with the use of larger basis sets, so too will the accuracy of the atomic charge assignments as per the QTAIM. Likewise for other approximations made in the context of DFT: the computational tractability they purchase is traded

⁵³ Since the electron density in overlapping regions is split equally by the involved atoms, a larger basis set (esp. the inclusion of explicit diffuse functions) will lead to more overlap and therefore more charge transfer. Strange things can happen as a result; e.g., centralized atoms will have a disproportionate amount of negative charge assigned to it. Also, a comment about the use of 'unreasonable' in this context. Amidst debate about whether the charge of an atom in a molecule is even a real thing, what I really meant to say is that the charges that result are not valued by any party involved.

off against model accuracy on system-wide energetics, and these energies are determined exclusively by the electron density. This means that no approximation will simultaneously improve the accuracy of the model and diminish our ability to accurately quantify atomic charge as conceived by this density-based scheme, and so too with other conceptual identifications based directly on the electron density.

The second draw to this strategy of defining concepts—namely that their subsequent quantification can be applied universally, irrespective of the chemistry being modeled—can also run into trouble. While the models may be used to calculate the energies of any chemical system, their accuracy sometimes varies across different classes of systems. Accordingly, a conceptual quantification procedure based cleanly on the details of a particular model can likewise demonstrate variable accuracy depending on the system being calculated. DFT approaches are no exception to this: current methods of approximating electron correlation systematically underestimate long-range correlation (Klimeš and Michaelides 2013). As such, they are relatively inaccurate when it comes to calculating energies of systems where long-range correlation plays a significant role, such as extended carbon systems. While there are attempts to design add-on corrections for these models (dispersion corrections are touched on in the next chapter) these corrections constitute ‘patch work’ theory that are generally not encouraged by the advocates of the rigor-driven camp.⁵⁴

⁵⁴ The theoretical standing of these patches is a complicated issue. Even in procedures that are touted as ‘*ab initio*’ it is typical to include some manner of empirical fitting. Sometimes one procedure will even be touted as “less empirical” than another (Grimme et al. 2010), but sorting out the import and validity of such statements is set aside for another time. In any case, the DFT patches mentioned here are not neatly based on features of the electron density in the manner advocated by those who promote ‘conceptual DFT’.

These practical considerations are not the only reasons some theorists advocate for this approach. In addition, this approach develops concepts in a way that satisfies several philosophical presuppositions one might reasonably have about the way science ought to work. Before engaging with these presuppositions, we will first look at the kind of conceptual changes that are embraced in the pursuit of concepts that possess a neat connection to physical theory.

4.2.2 Willingness to embrace conceptual modification

It is important to appreciate the willingness with which theorists of this camp accept modifications to their traditional chemical concepts in order to satisfy their desire to define concepts in the foregoing fashion. Although such changes can be seen in the previous chapter, a more compact example from the existing literature is as follows.

This relates to debate about the proper way to define the chemical bond. According to one advocate of the conceptual DFT program, chemical bonds ought to be understood as lines along the path of maximum density that connect neighboring atoms and which pass through a saddle point in the electron density (Bader 1990). These saddle points are local density minima along the line connecting the atoms, and local density maxima along the perpendicular. For example, five bond paths can be seen in the electron density contour map of ethene (below, left), the saddle points marked as dots. This way of defining the bond recovers the structure characteristic of ethene. However, there are scenarios, if sometimes contrived, where this way of defining the bond diverges significantly from traditional understanding of the chemical bond. For example, when an argon atom is placed inside a buckyball, according to the foregoing conception of the bond, argon forms a bond with each of the

sixty carbon atoms that make up the sphere (below, right) (Cerpa et al. 2008). This is very much in conflict with the traditional understanding of the chemical bond, according to which carbon can form a maximum of four bonds. The originator of this DTF-based conception of the chemical bond, Richard Bader, very much characterizes the rigor-driven approach; when confronted with the buckyball ‘counter example’ to his definition of the bond, he bluntly states that his definition was never meant to capture the traditional understanding of the chemical bond, which besides, he argues is “too restrictive to account for the physics underlying the broad spectrum of interactions between atoms and molecules” (Bader 2009).

This change (and others)⁵⁵ to the traditional understanding of concepts are welcomed for the sake of establishing a neater connection between the concepts and the physical models. That is, in considerations about whether a conceptual modification should be adopted, little weight is given to whether or not the redefinition would be in accordance with longstanding conceptual practices.

⁵⁵ Less contrived divergences between Bader’s bonds and traditional chemical also exist; for example, Bader’s bonds are often highly ionic even in cases that are widely accepted to be covalent (Fonseca Guerra et al. 2004). Regardless, Bader continues to embrace these changes; for example, another characteristic response: “Relating structure to a topological property of the electron density necessarily broadens the scope of the concept of structure, extending it beyond the models that were used in its development” (Bader 2010).

Indeed, disrupting these practices is sometimes touted as a virtue because they have thereby been shown to rest on unprincipled grounds.




Figure 4: Success and failures(?) of the topological AIM

Left: a contour map of the electron density in ethene (C_2H_4). Black dots indicate the saddle point (or “bond critical point”) along the bond paths connecting the atoms. (The diamond-forming lines identify, on Bader’s view, interatomic surfaces.) Figure taken from (Bader 1990).

Right: Argon forming sixty bonds paths with its buckyball cage. Figure taken from (Weisberg et al. 2011).

4.2.3 Philosophical motivations: definiteness & observability

So far, we have only discussed practical motivations for pursuing the rigor-driven approach to defining chemical concepts (as well as some hurdles to achieving these payoffs). There is a different kind of motivation also evident among members of this camp, which takes on a more philosophical flavor. This is the idea that only by defining concepts in the foregoing manner will chemical theory identify the truth of the matter, or so to speak, cut nature at its joints. Conversely, concepts that are not cleanly rooted in physical theory are misleading and detrimental to the future of the field. This sentiment can be found in many discussions of chemical concepts:

“The qualitative (‘hand-waving’) explanations [...] are not science (i.e., are nonsense) and should consequently be completely discarded. Instead the effort of chemical educators should be expended to demonstrate the myriad chemical substances and properties of real matter...” (Ogilvie 1990, p. 288).⁵⁶

“Overcoming the current obstacles [namely that our concepts appear to not be ‘real’] will require new methods to describe molecular structure and bonding, just as new concepts were needed before the phlogiston theory could be set aside” (Eberhart 2002).⁵⁷

“I believe in the existence of a single universe, one in which the laws of physics apply. I do not believe in the existence of parallel universes wherein their laws are either ignored or ‘bent’ to accommodate personal points of view. [...] The universe in which we live is more exciting and demanding of our abilities than any that one might care to invent” (Bader 2011).

These statements are not based on empirical evidence in any straightforward way, although this is a nuanced issue. While there is presumably no empirical finding that would definitely tell us what counts

⁵⁶ Or about the teaching of qualitative conceptual thinking: “Like the legendary emperor who displayed his newest suit of a material so fine as to be invisible, the authors and professors (teachers) who naively parrot these old untruths succeed only in exposing their ignorance” (Ogilvie 1990, p. 288).

⁵⁷ See also Eberhart & Jones (2013). In both, the attitude is one of acknowledging the importance of the traditional (fuzzy) chemical concepts, but insistence that the way forward is to develop ‘quantum mechanically correct’ concepts that are able to take over the work of the traditional concepts.

as science or what features theorists ought to seek in their conceptual tools, there is a sense in which empiricism speaks in favor of the foregoing points of view. Namely, if we seek to develop concepts that are observable or directly quantifiable in the physicist's sense of the term (i.e., possessing a corresponding operator ⁵⁸), then certain ways of defining concepts—including those promoted by advocates of the rigor-driven camp—are rightly privileged. This is to say, at stake in these debates is the question of *whether* the concepts are physically observable.

For instance, regarding the charge of an atom in a molecule, Bader's QTAIM definition is such that the charge of an AIM possesses a corresponding operator allowing, in principle, its 'direct' measurement. So, if we accept the QTAIM conception of the AIM, the properties of an AIM are physically observable; or put differently, if we assume good chemical concepts ought to be observable in the foregoing fashion, the QTAIM passes the test. According to the vast majority of alternative definitions, including the Hirshfeld AIM discussed in the previous chapter, no such operators exist, rendering the properties of AIMs according to these conceptions non-observable.⁵⁹ It is this fact that enables Bader to bolster his AIM proposal by asking: do we want to decide on these matters by free choice, or by observation? (Bader 2010).

⁵⁸ See Bader (2008) for the most in-depth discussion of this point; see especially section 2.2 (pp. 13717-8).

⁵⁹ And further, calculating the charge of an atom in a molecule according to the Hirshfeld conception of the AIM requires reference to non-physically motivated reference state: the promolecule. This kind of physically unmotivated calculational maneuvering is precisely what advocates of this camp disparage.

Still, the assumption that concepts ought to be observable in the foregoing manner remains just that—an assumption about what chemical concepts ought to be. While it is hard to imagine any theorist outright resisting this observability in their concepts, there remains debate about what exactly it provides, and what kind of sacrifice is appropriate in the pursuit of it. In fact, as we turn now to the opposing camp—what I call the understanding-driven camp—it will be clear that many would rather sacrifice this observability, as well as the above-mentioned practical payoffs, in order to develop concepts that achieve an altogether different aim.

4.3 Concepts for qualitative understanding

There are many chemists driven to develop concepts of a very different character, chemists who I place in the ‘understanding-driven camp’. The concepts developed and promoted by this camp lack many of the features sought after by the opposing camp: they are rarely ‘observable’ in the physicist’s sense of the term; they have many different formulations—rather than a single, definitive formulation—and these different formulations are often at odds with each other when it comes to quantitative matters; they are often not formulated using terms found in the physical equations on which computational models are based; and their formulations often make appeals to calculational schema that are not independently motivated.

Despite this list of what are arguably shortcomings, the concepts developed by this camp are valued for a single, often-stated reason: they facilitate *chemical understanding*. The nature of this understanding can be difficult to parse, and some might regard its answer to be in the domain of psychology. However, I have chosen an empirical, or perhaps pragmatic strategy instead: look to see what the

concepts—those that are valued for their ability to confer understanding—really allow chemists to *do*, and only then ask what it is in virtue of which the concepts are able to do this.⁶⁰ Briefly stated, my answer, which draws heavily from Hoffmann (1998) and de Regt (2017), is that such concepts—typically working in conjunction with other chemical concepts—allow chemists to qualitatively explain large swathes of chemical data in a generalizable way and without the need for regular reference to the results of computer calculations.⁶¹ These concepts are able to do this in part *because* of the flexibility purchased by their ‘fuzziness’ (as it is sometimes put), because this flexibility allows the concepts to work together in diverse chemical and theoretical settings. These ideas are unpacked below, where I hope to make it clear why, in pursuit of this kind of understanding, chemists of this second camp are

⁶⁰ Some theorists of the understanding-driven camp prefer the language of models (e.g., Pauling & Herman 1986, Houk and Liu 2014) and others prefer the language of concepts (e.g., Hoffmann 1998; Gonthier et al. 2012; Schwarz and Schmidbaur 2012). Though not without exception, the language of models is used to describe a particular application of one or several concepts to explain a particular chemical phenomenon. I believe the conceptual practices underlying the development of these models is the same, and while I will continue to talk of concepts rather than models, I include these modeling practices within the scope of my analysis.

⁶¹ As Hoffman puts it, “I define understanding (operationally, if faced with a really good computer program and a complex problem), as being able to predict qualitatively (this forces you to think before) the result of a calculation before that calculation is carried out. If the calculation’s result differs from what your understanding gives you, well then it’s time to think again, do numerical experiments until you rationalize (that’s also to ‘understand’) the results” (Hoffmann 1998, p. 3). Henk de Regt puts this line of thinking in a theory-oriented account of scientific understanding: “A scientific theory *T* is intelligible for scientists (in context *C*) if they can recognize qualitatively characteristic consequences of *T* without performing exact calculations” (de Regt 2017). I aim to place this idea in a concept-oriented account with additional moving parts.

led to foster concepts that break from the physical models in ways that otherwise appear stubborn and shortsighted.

Again, some clarification on my chosen terminology is in order. I call this group the ‘understanding-driven camp’ because I believe understanding is the central guiding principle in their evaluation of conceptual formulations and their implementations. Though it remains to elaborate on the nature of this understanding, I do not mean to suggest that the concepts developed by the rigor-driven camp are devoid of this feature; they too may facilitate some understanding in the same manner. The difference is in the priorities, and this difference manifests in very real ways. For now, put generally: this difference in priority manifests in the way various features of a conceptual definition are weighted in its evaluation; for example, the extent to which it is acceptable to stray from the physical equations in order to make sense of some data will be a point of contention between the two camps.

Also, some clarification on the scope of my analysis. I treat this issue in a modern context, as a debate about the proper way to relate longstanding conceptual practices and newer computational models, and as such it is a modern topic. While to some extent I examine its historical threads back through 20th century (Chapter 2), I do not aim to make any claims about its historical origin, or the extent to which it differs from any conceptual strategies that might have existed prior to the 20th century. I will only say that I believe many of the concerns that guide the two camps discussed here will be present in earlier history of chemistry, but that I do not believe the contours of the debate will be the same or as localized. That is to say, I believe the advent of computational models, at least those accurate enough to have practical application, has intensified and localized the concerns that may or may not be traceable well into past chemical and likely physical theory.

4.3.1 Practical motivations and their implications

The motivation for defining concepts in the manner preferred by the understanding-driven camp can be found in what is lacking, from their perspective, in the physical equations underlying computational models, as well as the concepts based directly thereon. This desideratum is called many different things, but I believe it is most clearly captured in the following: cognitively manageable qualitative patterns. In short, what their hard-earned concepts provide is the ability to qualitatively predict large swaths of chemical data without needing to regularly appeal to the results of computational models.⁶²

The desire for cognitively manageable qualitative patterns can be seen even in the relatively early stages of quantum chemistry, particularly clearly in the case of the concept of electronegativity.⁶³ Remember that electronegativity was put forth as a measure of the strength with which an atom of a given element attracts electrons to itself, and one major strategy for gauging this power was using a comparative study of bond energies (Pauling 1932b). Empirical bond energies were, for the most part, restricted to simple diatomic molecules, and thus electronegativities had to be assigned on the basis of few data points. The resulting electronegativities—Pauling’s values—were widely employed and valued for the simple, intuitive explanations they produced; anyone with a basic understanding of chemistry was henceforth able to qualitatively order bond energies of simple molecules and likewise with polarity

⁶² I will also speak about this desideratum in terms of a tradeoff between simplicity and the generality of qualitative accuracy. The nature of this tradeoff will become clearer when we examine the practices used to develop these concepts in the following section.

⁶³ Parts of this section are reproduced from the previous chapter, tailored here for a different narrative.

and other molecular features (see Chapter 2). At the same time, both quantitative and qualitative errors began to emerge, especially for complex molecules, and sorting them out would require a more nuanced approach to electronegativity. One such addition was the inclusion of considerations of the hybridization⁶⁴ of the atom: in the language of valence bond theory, carbon that is sp^3 hybridized (e.g., CH_4) could then be assigned a different electronegativity value than carbon that is sp^2 (C_2H_4) or sp (C_2H_2) hybridized (Mulliken 1935). This accounted for many errors, but already the concept of electronegativity was beginning to lose its simplicity. Here is a nice expression of these concerns in a landmark survey on the concept following these developments:

The crucial point, as Evans (1951) has remarked, is that there is the possibility that every molecule must needs be treated as a special case. (Is the valence state of nitrogen the same in NH_3 , in NF_3 , and in $N(CH_3)_3$?) The electronegativity concept has been useful in the past largely because it denied the need for special treatment of every individual case; thus, the electronegativity of nitrogen was considered well enough defined by a single number. This is now questioned [...]. It seems clear that the further development of the electronegativity theory will attempt to define the range of electronegativity open to an atom. There remains the problem of choosing from the range the unique value which the atom achieves in a given molecular environment. Unless some comparatively simple

⁶⁴ Note that the language of ‘hybridization’ and ‘valance state’ has historically been inconsistent and varies between subfields—sometimes they are equivalent and sometimes not. For our purposes they can be treated as the same.

criterion of choice can be found...the discomforting possibility that Evans has foreseen may become a very real one.

-Pritchard and Skinner (1955)

The key phrase here is ‘comparatively simple criterion of choice’. While the developments surrounding the concept were increasing the precisions with which it could be employed, it was at the same time reducing its cognitive manageability: it was becoming less possible to use the concept without the aid of, for example, an extensive compiled table that listed the electronegativity of atoms in all the different chemical systems in which it can be found.⁶⁵ The modern version of this concern does not involve tables, but rather computers that (re)run calculations—if one wanted to know the electronegativity of an atom in a molecule with precision of the modern DFT conception of electronegativity, one would have to run a calculation using advanced computational models.

However, losing the ability to predict these values does not alone wholly capture the core of the problem. Notice that electronegativity values are valued more for their approximate and relative

⁶⁵ While simple electronegativity tables that assign a single value to each element, such as Pauling’s, do involve a table that must be occasionally referenced, it is not difficult to internalize the important values; in any case, the difference is one of degree.

magnitudes rather than their precise or absolute magnitudes.⁶⁶ To understand why, it is important to consider the qualitative nature of the practices surrounding these concepts. By comparing the approximate electronegativity values of the different atoms in a molecule, certain *qualitative* patterns emerge. Some qualitative patterns are rather obvious, such as more electronegative atoms being more negatively charged, or more polar bonds existing between atoms with a larger difference in electronegativity; other qualitative patterns are less obvious, such as that chemically significant hydrogen bonding only occurs when hydrogen is bound to an atom that is more electronegative than it is by roughly 0.7, and that adding highly electronegative groups to strained molecules will reduce the strain. These patterns are qualitative in the sense that they are formulated in terms of ‘more’ and ‘less’,⁶⁷ and they are simple in the sense that they can be stated without a host of exceptions or rules for when they hold or need to be adjusted in some fashion.

The crucial point is that the simplicity of these qualitative patterns can be significantly impacted by how the concept is defined and thereby quantified. To return to the case at hand, the lamenting

⁶⁶ I am using these terms in what I believe is the common sense manner, but for clarification: the degree to which a value is approximate or precise can be evaluated based on how many significant figures are present; a value is relative if it is given only in comparison to others (or, which I take to be the same, if there is an arbitrary baseline, such as setting the least electronegative atom to a value of zero), and a value is absolute if it is given in terms of a reference that is independent in some meaningful way (e.g., electronegativity values put in units of energy).

⁶⁷ Some of these patterns or rules might be described as ‘semi-quantitative’ (e.g., the hydrogen bonding rule above) because they are formulated in terms of approximate numbers. For simplicity and because I do not believe there is an important difference to be found here, I will refer to them all as qualitative patterns.

above—about the possibility that “every molecule must needs be treated as a special case”—is not merely due to the need to resort to an extremely large table or a computer calculation to determine an atoms electronegativity in a given molecule. Rather, it is due to the fact that the simple qualitative patterns, which drew theorists to think in terms of electronegativity, were becoming less simple, more difficult to identify using the more physically motivated conception of electronegativity. The newer conceptions of electronegativity, which included details of the chemical environment in its calculation, were at first glance bringing only slight quantitative corrections to electronegativity values. No, the valence state of nitrogen is not exactly the same in NH_3 , in NF_3 ; it is slightly different, and so too is the electronegativity of nitrogen in the two molecules. Considering these quantitative corrections in the context of the foregoing qualitative patterns—formulated in terms of ‘more’ and ‘less’—problems quickly arise. For example, nitrogen is generally more electronegative than carbon, but how does the electronegativity of nitrogen in NH_3 compare to the electronegativity of carbon in CF_4 ? The answer depends on the magnitude of these small quantitative corrections, and without a ‘simple criterion of choice’ that can be used to determine these corrections, the question will have to be put to a computational model or extremely large table. It is the obfuscation of the qualitative functionality of the concept that is truly lamented.

It is the response to this problem that separates the two camps. On the one hand, there is the attitude of those in the rigor-driven camp: why should these corrections be lamented? If you want your chemical concepts to survive in the modern era, they should be put in terms that are compatible with modern physics. There are compelling and rather natural ways to accomplish this, terms in the physical equations that capture the essence of electronegativity (most notably the negative of the chemical potential, Chapter 2). To the extent that this reformulation of the concept brings corrections, this

should be accepted and even celebrated; things are not as straightforward as previously thought, but such is science.

And the opposing response, that of the understanding-driven camp: there is no reason to commit to the proposed reformulation of these concepts. These concepts are valued for the simple qualitative patterns they track, and accordingly they ought to be defined in whatever way—or even *ways*—that enable them to function in this manner. If these qualitative patterns cannot be identified neatly using terms found in the physical equations, this only means the latter is deficient for the task of finding them.

It is easy to appreciate how one might reasonably charge such theorists with ignoring or bending physical laws to accommodate their preconceived views, as those of the rigor-driven camp have repeatedly argued. It is more difficult, but no less important, to appreciate how the traditional use of these concepts can be not only maintained, but also further developed despite these challenges ‘from below’. As we will see, there are nuanced strategies used to develop concepts that function to identify qualitative patterns. To do this, however, the resulting concepts must possess some strange features; these features will be examined following a demonstration of how chemical concepts can be jointly developed to expand their ability to identify simple qualitative patterns.

4.3.2 Developing flexible conceptual networks

To shed further light on the nature of chemical concepts according to the understanding-driven camp, I again take a practice-oriented approach: examine how these theorists develop and use their concepts. I have already mentioned the central feature that these practices attempt to bestow upon the concepts

they develop, namely the ability to qualitatively predict large swaths of chemical data in a simple, cognitively manageable manner. In order to develop concepts capable of doing this, liberties are taken when it comes to interpreting both empirical and computational results; namely, the concepts used to interpret these results are molded to fit the job. These liberties are justified on the basis that, according to members of this camp, the concepts are not tied to any physical basis or even single set of empirical results. Put differently, the quantification of the concept is underdetermined by theoretical and empirical commitments, and the concepts can therefore be cast to track certain trends in the data.

This underdetermination is nicely illustrated by the case of atomic charge being underdetermined by dipole moment (which for most purposes can be directly measured by empirical means). This is discussed nicely in Gonthier et al. (2012, p. 4675):

Dipole moments are the most legitimate probe of atomic charges for polar molecules; however, assigning atomic charges based on measured molecular dipole moments is unrealistic. An atom inside a molecule does not possess a spherical distribution of charge, and its resulting atomic polarization may oppose or enhance the charge transfer dipole (i.e., the dipole moment deduced from the atomic charges). In other words, a small value for the experimental dipole moment is equally compatible with both large atomic charges and large atomic polarization or with small atomic charges and small atomic polarization.

Thus, what appears to be the closest we can get to empirically measuring the charge of an atom in a molecule—using dipole moments—does not commit us to a certain quantification when it comes time to assign charges to the atoms in the molecule. An assumption underlying this passage, one which

is attacked by advocates of the rigor-driven approach, is that physical theory likewise does not commit one to certain atomic charge values, that it underdetermines their quantification. For now, we will set aside this dispute to examine how, in the presence of this underdetermination, chemists of the understanding-driven camp decide on certain quantification for atomic charges.

In chapter two, we examined in detail two density-based proposals for how one might identify the atom in the molecule (AIM) for the purposes of quantifying atomic charge: the stockholder AIM (Hirshfeld 1977), and the topological AIM (Bader 1990) (see Chapter 2 for technical details). The two schemes partition the electron density of molecules into very different AIMs, as can be seen below for one of the nitrogen atoms in N_2 . In this example, given the symmetry of the molecule, they happen to agree on the charge assigned to the two nitrogen atoms (zero); however, in most cases they will assign different charges to the atoms in the molecule, and in complex systems they may even disagree on which atoms are positively charged and which are negatively charged.




Figure 5: Two ways of identifying the atom in the molecule (revisited)

Two different partition schemes as applied to a single nitrogen atom in N_2 (stockholder AIM on left, topological AIM on right). In each of the three images, the two ‘+’s represent the nuclei in the molecule, and the contour lines display the electron density. The two images on the bottom show the electron density assigned to the atom on the ‘left’ of the molecule according to the two competing schemes (Adapted from Davidson and Chakravorty 1992, p. 321).

The question then arises: how should chemists choose between these two partition schemes? Where those of the rigor-driven camp are guided by considerations about which one is more motivated by physical theory, the understanding-driven camp follows a very different approach.⁶⁸ Here, the goal is to choose a partition scheme that is consistent with other concepts that are likewise employed for the purpose of making qualitative patterns accessible. For example, another concept that intuitively ought

⁶⁸ “Unlike theories, the justification of concepts is provided by their success in applications rather than in firm proofs” (Gonthier et al. 2012, p. 4672). I hope to shed light on specifically what counts as a successful application.

to come into play in discussion of atomic charge is electronegativity. Atoms in the molecule that have a greater “power to attract electrons to themselves” (Pauling 1939, p. 58) should, other things being equal, have a more negative charge. This, anyway, is how the concept is intended to work, and this intention can be realized if the right charge partition scheme is selected. Accordingly, one approach used to select a partition scheme is to test how well it conforms to these expectations relating to electronegativity considerations.

Molecule	Stockholder AIM charge
CH ₃ CH ₂ C(OH) ₃	0.0315
CH ₃ (CH ₂) ₂ C(OH) ₃	0.0288
CH ₃ (CH ₂) ₃ C(OH) ₃	0.0295*
CH ₃ CH ₂ C(SH) ₃	0.0322
CH ₃ (CH ₂) ₂ C(SH) ₃	0.0290
CH ₃ (CH ₂) ₃ C(SH) ₃	0.0319*
CH ₃ CH ₂ C(NO ₂) ₃	0.0417
CH ₃ (CH ₂) ₂ C(NO ₂) ₃	0.0411
CH ₃ (CH ₂) ₃ C(NO ₂) ₃	0.0337
CH ₃ CH ₂ C(PO ₂) ₃	0.0435
CH ₃ (CH ₂) ₂ C(PO ₂) ₃	0.0422
CH ₃ (CH ₂) ₃ C(PO ₂) ₃	0.0376

Figure 6: Hirshfeld AIM charge test

Charges correspond to the hydrogen atoms shown in bold (adapted from Saha et al. 2009, p. 1795).

* deviations from the expected trend.

Above are the results from one such test, where the stockholder partition scheme is tested for its ability to reproduce expected ‘through bond’ effects of electronegativity.⁶⁹ All molecules considered in the test are hydrocarbon chains with one or more electronegative functional group(s) at one end of the chain. The stockholder charge of the emboldened hydrogen atoms, located on the end opposite the functional group, are shown in the right column. The expectation is that, other things being equal, the closer or more electronegative the functional group, the more positive the charge of the hydrogen. Of the dozen molecules tested, two deviated from expectation (indicated by an asterisk). For example, comparing the third row with the first two, it was expected that the hydrogen atoms become less positively charged as the electronegative hydroxyl groups become farther away; yet for some reason, the positive charge increased, and this counts as a point against the stockholder partition scheme.

Note that the absolute magnitude of the assigned charges is not particularly important in these tests. Rather, it would be enough that the charges are *qualitatively accurate* with regards to the set of molecules considered, correctly assigning charges that recapitulate the expected *trends*—and keep in mind, it is these trends on which the other concepts involved in the test (in this case electronegativity) are also based. In this way, a successful test affirms the qualitative accuracy of all concepts involved in the tests. Also note that the foregoing test only incorporates two concepts; more thorough tests would hopefully (though see below) show that the results are qualitatively accurate with regard to other concepts, such as bond strength and nucleophilicity. Performing these more thorough tests would put

⁶⁹ This effect is distinguished from ‘through space’ charge transfer, which in the present cases should be negligible given the linear structure of the molecules.

the conceptual description of these molecules in contact with additional empirical features besides polarity.

What ideally results, then, is a network of concepts that allows chemists to predict the qualitative properties of molecules and which is simple enough to be employed without needing to regularly resort to computational models.⁷⁰ This, I argue, is what it takes to possess chemical understanding.

Computational models and this conceptual network can be thought of as two frameworks for dealing with chemical systems, illustrated below. The conceptual network can be thought of as a lossy compression of the computational models, but one that preserves what is needed for chemical understanding. In this way, the two frameworks share a point of conversion, where they speak about the same things—namely, qualitative empirical features of molecules and other chemical systems—and agree in their assessment. For example, both frameworks can be used to calculate relative dipole moments, reactivities, and energies, and they should agree on these matters despite very different ways of coming these shared conclusions.

⁷⁰ In complex molecules, there will always be the need to run a calculation to get a handle on even the simplified picture that uses traditional concepts. However, as long as the concepts are working as they ought to, a single calculation submitted to the computer should shed light on many related molecules in a qualitative manner.

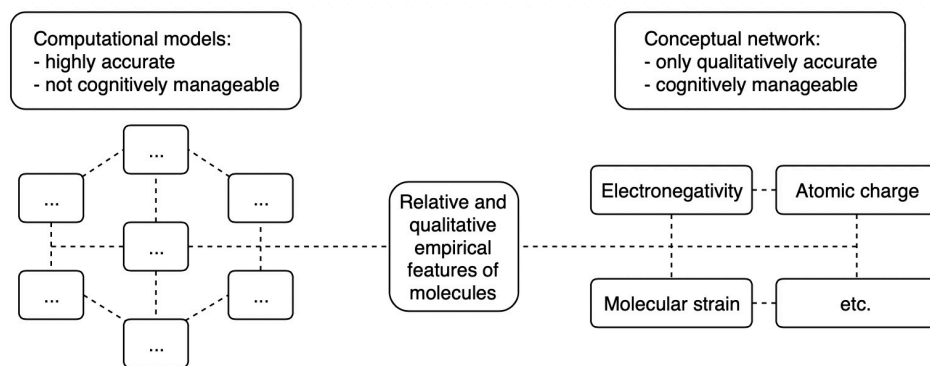


Figure 7: Computational model vs. network of chemical concepts

The framework of a computational model (left) and the conceptual framework of traditional chemical theory (right) share point of framework intersection. The concepts employed by the conceptual network are intuitive and cognitively manageable, while those in the computational model are not.⁷¹

However, sometimes these tests do not go as one might hope, sometimes proposed procedures for quantifying traditional chemical concepts do not perform as expected. Such an outcome has already been shown in Figure 6, where several assigned charges were, qualitatively speaking, the opposite of what was expected. This is to say that, to some degree—and this typically is a matter of degree—the Hirshfeld partition scheme failed to assign atomic charges in a manner consistent with the rest of the

⁷¹ The terms found in the computational model (currently represented by "...") will vary depend on what, exactly, we consider to be the computational model. We may consider them the terms found in the intractable Schrödinger equation, we may consider them the terms found in its also intractable DFT solution (as would many of the rigor-driven camp), or we may consider them the terms found in the tractable approximations to the former equations. For my purposes, it does not matter because none of these terms offer what traditional chemical concepts offer, and in virtue of which they facilitate 'chemical understanding'.

operative conceptual network (with a focus here on electronegativity). In such cases, work is put into developing a new partition scheme (or generally speaking, into a new technique for quantifying traditional concepts). This may be a slight modification to the calculational procedure, as has been done several times for the Hirshfeld scheme (mentioned in Chapter 2), or it may be an altogether different approach.

In the absence of a quantification technique that performs with qualitative accuracy across all chemical systems—and we often find ourselves here—another strategy may be adopted in the meantime: different quantification techniques may be adopted for different classes of molecule. For example, if one technique performs well for most organic molecules, and another for most inorganic molecules, then simply use the techniques within their respective domains. In fact, this was precisely the outcome and recommendation of Saha et al. (2009) in the study introduced above. This practice of adopting a pluralistic approach to the quantification of chemical concepts certainly has its downsides (examined below) but makes sense in light of the ultimate aim of these investigations: to extend the generality and qualitative accuracy of the conceptual network that forms the basis of chemical understanding. Achieving this aim may require some strange maneuvering when it comes to developing the concepts within the network, but according to the understanding-driven camp the price is acceptable.

We now turn to an examination of the nature of the concepts that form this network.

4.3.3 Willingness to embrace the strangeness of ‘fuzzy’ chemical concepts

In the previous section we got clearer on what it is that facilitates chemical understanding: a network of concepts that are tailored to track qualitative patterns in chemical data. We also saw how this

tailoring is carried out, developing and testing a multitude of proposals for how one might connect the chemical concepts to the physical models, and adopting whichever method(s) best tracks the relevant qualitative patterns dictated by other concepts within the network. While tracking the important qualitative patterns is the goal, the concepts that are able to do this possess other features that are not intrinsic to their intended functionality, but which are accepted as a means to an end; they might be considered side effects. These side effects include some rather strange features, which we examine here.

We have already seen how the practices examined in the previous section can lead to one such feature: *pluralism*. In the pursuit of being able to quantify their concepts in a way that is consistent with the relevant parts of their conceptual network, it is common to employ different quantification schemes for different domains of chemistry. As these quantification schemes are what make the connection between the chemical concepts and the physical models, this practice precludes the possibility of offering a singular physical definition of the chemical concept. This is not merely because there is more than one way to say the same thing; the different schemes for quantifying these concepts offer different numbers, both in terms of the absolute values as well as how values assigned to different molecules compare to each other. Indeed, it is precisely because they offer different numbers that they are applied to different domains of chemistry.

Another related feature of many of these concepts is their *unobservability*. As discussed above, for many chemical concepts there are proposals for how to relate them to the physical models in such a way that renders the concepts physically observable in the quantum mechanical sense of having corresponding operators (see Bader 2005 for mathematical details in this context). Where these proposals are championed by those of the rigor-driven camp, those of the understanding-driven camp

are not averse to concepts being physical observable so much as they are indifferent to it. Use is made of quantification schemes that correspond to physical observables as well as those that do not, and this difference does not weigh into their evaluation, which again, focuses on their ability to track qualitative trends in the data. In defense of this practice, these concepts have been called Kantian *noumena*, “conceptual constructs ultimately unknowable by observation or unique definition but conceivable by reason” (Parr et al. 2005). Note, though, that the context of this philosophical discussion strongly suggests this point is put forward only as a defense, and not a driving force of these practices.

More generally, those of the understanding-driven camp are far more permissive when it comes to giving a quantification scheme the light of day. When confronted with a proposal for how one might quantify traditional concepts in terms of modern computational models, there is little that would prevent its evaluation with regard to capturing the intended qualitative trends. This also includes permitting the use of quantification schemes that make use of *ad hoc reference structures*. A prominent example of this is the Hirshfeld AIM partition scheme and its subsequent iterations (Hirshfeld 1977; Vanpoucke et al. 2013). The ‘promolecule’ that is required for calculating atomic charges according to this scheme is a hypothetical molecule with atoms whose electrons do not interact in any way, but which is an essential reference point in order to carry out the calculations needed to determine Hirshfeld charges (see chapter 2 for details). Not only does this structure not exist in nature, but it is not theoretically motivated for any reason other than its role in a scheme that decompose molecules

into atoms that are pleasing to chemists.⁷² Further, when problems are found with the Hirshfeld scheme, changes are proposed to this reference structure and the calculation based thereon in order to find a better quantification scheme (Bultinck et al. 2007, Vanpoucke et al. 2013).

Accepting these conceptual features—pluralism, unobservability, and the use of ad hoc reference structures—allows chemists the flexibility that is often needed to develop a conceptual network capable of tracking important qualitative trends using simple concepts. It is precisely this flexibility that those of the rigor-driven camp criticize, arguing that these practices only serve to affirm presuppositions about the behavior of these concepts and their role in explaining chemical phenomena. However, I hope to have shown that the two camps are playing very different games, with very different standards of success.

⁷² One might think this statement is unfair—after all, the promolecule might be conceived as an approximation to the molecular electron density where certain electron-electron interactions have been ignored. However, which electron-electron interactions are ignored is carefully selected (namely those from different atoms) in order to influence the result, that is, to get pleasing atoms in the molecule.

Alternatively, one might say there is theoretical motivation to consider the promolecule because it minimizes the information loss, as expressed by the Kullback-Liebler measure, during the formation of the molecule (Kullback 1959, Parr et al. 2005). But again, this information loss is measured between the electron density of the promolecule and the actual molecule, and hence its significance rests on that of the promolecule.

5.0 Extending the Implications and the Future of Philosophy of Chemistry

The aim of this final chapter is to explore the future of philosophy of chemistry. I argue that, in light of the above discussion of chemical concepts, drawing on chemical practice to feed investigations of the ontological status of chemical concepts—a rather common exercise among philosophers of chemistry—is misguided. I go on to suggest several lines of philosophical inquiry that promise fruit and which are genuinely aligned with scientific practice in chemistry. These include exploring what concepts built for qualitative understanding can and cannot do.

5.1 Introduction

Philosophy of chemistry has not flourished—both in terms of its ability to attract students and its contributions to broader issues in philosophy of science—and I turn now to prospects for its future. Above I hope to have made a compelling case that theoretical chemists working in the computational era encounter problems relating to conceptual development that are philosophical in nature. They are tasked with developing concepts without a clearly stated account of what, precisely, these concepts are supposed to accomplish, and the absence of clearly stated standards of evaluation has sustained the longstanding debates explored above. In discussing the practice of computationally guided conceptual development—where the goal is to develop a network of concepts that are qualitatively accurate across large swathes of chemistry in a way that is consistent with one another—I hope to have clarified some standards of conceptual evaluation that one might reasonably hold.

However, setting my specific proposal aside, the problem area I have identified has implications for philosophy of chemistry. The implications are in part restrictive: appreciating the fact that there is genuine debate among chemists about what chemical concepts are intended to do means that certain approaches to philosophy of chemistry cannot get off the ground. More specifically, the presence of this debate among chemists means that investigations into the ontological status of chemical concepts cannot be properly grounded in scientific practice without first engaging the very same issues that divide chemists, and for the very fact that these issues have ontological implications. The implications are also, I hope to show, fruitful: appreciating the very same debate among chemists opens up several avenues worthy of philosophical consideration that hold the potential to be helpful to chemists and offer philosophers of chemistry an opportunity to develop a unique contribution to philosophy of science.

I begin this chapter with the critical component, arguing that perhaps the largest ongoing literature in the philosophy of chemistry—about the nature of the chemical bond—is fundamentally misguided for the above stated reasons, and that philosophers of chemistry should instead focus their attention on the philosophical issues that are motivated within chemical practice. I proceed to identify two open-ended questions worthy of philosophical pursuit. One is immediately practical: characterizing the limitations of concepts built for qualitative understanding. By abstracting away from longstanding conceptual debates among chemists, these limitations can be identified. The second question is more abstract, but likewise attuned to philosophical training: characterizing the auxiliary functions of chemical concepts. I hope to show that these fuzzy, ill-defined concepts that lack a certain sense of rigor are able to do additional things for the chemical community in virtue of these same features.

5.2 Misguided philosophical pursuits and the nature of the chemical bond

I have made it most of the way through a dissertation on the nature of chemical concepts without making much contact with the current literature in the philosophy of chemistry that focuses on just that: the nature of chemical concepts. I would now like to broach this literature in order to show how the foregoing chapters suggest an alternative direction for philosophy of chemistry. Having waited so long, the groundwork for my case is already set.

Witness the debate among philosophers of chemistry about the nature of the chemical bond. This literature begins with the observation that the concept of the chemical bond has been difficult to articulate precisely in quantum mechanical terms (Hendry 2008)—(which generally speaking was our starting point, too)—and proceeds to explore a variety of ways the bond might be conceived and how these conceptions stand in the face of quantum mechanics. However, the alternative conceptions of the bond under consideration are not what you might expect given our historical foray of chapter 2.⁷³ Instead, it is the ‘structural’, ‘energetic’, and ‘interactionist’ conceptions of the bond that are vetted for their compatibility with quantum mechanics (Hendry 2008; Weisberg 2008; Hettema 2008; Weisberg, Needham, and Hendry 2011; Esser 2018). To elaborate, the structural conception of the bond asserts

⁷³ I have not focused on the concept of the chemical bond, in part to make a case for the generality of this phenomenon, but by now it should come as no surprise that chemists are engaged in debate about the proper way to conceive of the bond just as they are about the proper way to conceive of electronegativity or the atom in the molecule. See Gonthier et al. (2012) for orientation.

that bonds are “material parts of the molecule that are responsible for spatially localized sub-molecular relationships between individual atomic centers” (Hendry 2008, p.917); the energetic conception holds that facts about bonding are “just facts about energy changes between molecular or super-molecular states” (Hendry 2008, p.919); and the newer interactionist conception holds that “Bonding is a special physical relationship between pairs of atoms in a molecule where displacement (within limits) leads to particular restorative responses within the molecular framework” (Esser 2018 pp.16-17). The game, then, is to see which of these conceptions can be maintained given what quantum mechanics says about chemical systems; for example, the structural conception apparently struggles to account for the fact that greater delocalization of electrons stabilizes the molecule (Weisberg 2008).

I argue this philosophical undertaking is fundamentally misguided both in the question it seeks to answer and in the manner it draws from chemistry to answer it. Note first that it is the philosophers who brought the three options—the candidate conceptions of the chemical bond—to the table; debate among chemists about the proper conception of the chemical bond is not put in terms of the ‘structural’, ‘energetic’, and ‘interactionist’ conceptions of the bond, and it is unclear how the bonding models chemists do discuss are supposed to map onto these distinctions.⁷⁴ But even more problematically, chemists are in fact engaged in substantial debate about the proper way to conceive of the bond. Amidst debate about whether the chemical bond can be identified with a single bonding model (see Bader 2009 in favor; Grimme et al. 2009 against), surveys of when different conceptions

⁷⁴ It strikes me as difficult to find a bonding model discussed in the chemical literature that does not involve ‘structural’, ‘energetic’ and ‘interactionist’ components. Why one of these components should be deemed the true nature of the bond is puzzling.

of the bond might be appropriate (Gonthier et al. 2012), and discussion of bonds being the ‘unicorns’ of theoretical chemistry (as Frenking and Krapp 2007 have colorfully argued), how are philosophers supposed to pick a bonding model to use as their scientific reference point? And note the great impact this decision makes. If one draws from Richard Bader (1990) for an understanding of where the concept stands in chemistry (as does Esser 2018), one may get a very different impression than if one draws from Roald Hoffman (1998). The alternative of doing a robustness analysis (Weisberg 2008) likewise does not do justice to the disagreements among chemists. In the preceding chapters, I hope to have shown that these disagreements are in fact substantial, and precisely about what kind of a thing chemical concepts are.

I am not suggesting that philosophers of chemistry must await consensus among chemists before there can be grounds for a philosophy of chemistry—just the opposite. The problem is to pursue the investigation of the ontological status of the bond in way that presupposes a resolution to ongoing (and dare I say philosophical) debate among theoretical chemists. Greater still, these are pressing questions for chemistry that philosophers are trained to address. More and clearer dialogue (and if possible, agreement) surrounding how theoretical chemists make progressive conceptual developments in the era of computational chemistry...this is both philosophical and active.

Perhaps the problem has been that philosophers seeking to do philosophy of chemistry were under the impression that if there was going to be a philosophy of chemistry, then they would have to bring the philosophy. Prior to weighing the merit of the structuralist conception of the chemical bond, Weisberg (2008) comments:

“Most commonly, chemists adopt a pragmatic stance and simply demand that bonding theories be useful for making predictions and aiding in the synthesis of new molecules. The underlying ontological status of the bond holds little interest. Although useful to everyday chemical practice, this approach is deeply unsatisfying to the philosopher of chemistry because it leaves unanswered fundamental questions about the nature of the chemical bond” (Weisberg 2008, p.2).

What I find deeply unsatisfying is the idea that the nuanced epistemological considerations examined in the previous chapters of this dissertation should be deemed mere pragmatics. They are the substance of modern theoretical chemistry, and merit the attention of philosophers—not least because they have clear implications for questions about ontology, reduction, and pluralism. Furthermore, being problems native to chemistry, they are a foothold for philosophers of chemistry to prove that chemistry deserves attention beyond being testing ground for ideas generated elsewhere in philosophy of science.

In what follows, I continue the exploration of concepts built for qualitative understanding by assessing their limitations and strengths. This analysis is preliminary, and my greater aim is to develop the case that the epistemological maneuvers involved in developing concepts for qualitative understanding carry with them rich philosophical implications about chemistry.

5.3 Limitations of concepts built for qualitative understanding: the strain case study

The goal of the practice of computationally guided conceptual development, as motivated in Chapter 2 and outlined in Chapter 3, is to develop a network of simple concepts that each possess qualitative

accuracy across large swathes of relevant chemistry. Although the concepts in this network might address very different features or behavior of chemical systems, they must be consistent with each other if chemists are to use them to tell qualitative stories. In order to obtain this consistency, concepts are occasionally quantified by several different procedures (and we observed this kind of pluralism above with regard to the concept of the atom in the molecule). In this way the concept is engaged in a multifaceted balancing act: on the one hand, no more complicated than necessary in order to pick out a qualitative trend, and on the other hand, defined in such a way (or ways) that its qualitative behavior is compatible with related concepts.

While I am resistant to quick turns to prescriptive philosophy of science, my account does suggest implications that have practical consequences for conceptual development in chemistry. Below I explore one such implication: some questions that deal with concepts built for qualitative understanding are inherently resistant to a clear answer due to limitations of such concepts. To explore this idea, I make use of a brief case study that centers on the barrier to rotation of ethane. This case study is useful because the question being asked—whether the barrier to rotation is due more to strain or more to hyperconjugation—involves the direct interaction of two concepts, and for this reason the question is conceptually precarious. Additionally, there are fewer concerns about the ‘real’ way to quantify these concepts, which simplifies the case study in a way that serves our present aim.⁷⁵

⁷⁵ For example, “Nature does not know steric hindrance, only energy changes” (Fernández, Frenking, and Uggerud 2010).

5.3.1 The barrier to rotation of ethane

Ethane consists of two methyl moieties connected by a carbon-carbon single bond, and because single bonds admit rotation, the ethane molecule is capable of conformational change. The two conformations we care about are the staggered and eclipsed conformations (Figure 8) for it is the energy difference between these conformations that determines the barrier to rotation.

Since the 1930s, experimental and theoretical data have agreed on the magnitude of this barrier: the staggered conformation is about 2.9 kcal/mole more stable than the eclipsed conformation (Pitzer & Kemp 1936). The initial explanation for this rotational barrier was strain: in the eclipsed conformation, each C—H bond is ‘cis’ with another bond on the neighboring atom, and these filled atomic orbitals will resist close proximity due to a) the Pauli exclusion principle, and b) classical Coulombic repulsion. It was generally accepted that “strain” captures the sum of these two forces. Soon after, the idea of hyperconjugation was explored as a possible alternative explanation by Robert Mulliken (1939). In short, hyperconjugation occurs when there is partial bonding interaction between bonding and antibonding orbitals. This stabilizing interaction requires a suitable orbital orientation, and in ethane this only occurs in the staggered conformation, potentially explaining the barrier to rotation in this way.




Figure 8: Two conformations of the ethane molecule.

Left, hyperconjugation: in the staggered conformation, electron density from a C—H σ bond may be donated to the σ^* (antibonding orbital) of a C—H bond on the adjacent carbon; this sharing requires the σ and σ^* be ‘anti’ to each other, as they are in the staggered conformation, and lowers the overall energy of the molecule.

Right, strain: in the eclipsed conformation, no significant hyperconjugation is possible, and eclipsed C—H bonds are sterically stressed.

Corresponding Newman projections are shown below each 3D model.

In 1939 Mulliken judged that hyperconjugation in ethane—even in the staggered conformation—would be negligible, and for the most part textbooks would agree with him throughout the 20th century. Still, Frank Weinhold and others (Brunck & Weinhold 1979, Reed et al. 1988, Reed & Weinhold 1991) persistently pushed the idea that hyperconjugation was a more significant factor, and

many others remained uncertain.⁷⁶ Then in 2001, Pophristic and Goodman published results showing that hyperconjugation wholly accounted for the barrier to rotation, and that, in fact, there was *no* strain present in the eclipsed conformation. Since then, there has been continued debate (Weinhold 2001, Schreiner 2002, Bickelhaupt et al. 2003),⁷⁷ but little consensus (see Mo 2011 for discussion).

So it is that after nearly 90 years of knowing the magnitude of the barrier to rotation, there remains debate about the proper way to use chemical concepts to account for it. This persistence may be partly due to the very same debate about the proper function of chemical concepts explored in the previous chapter; needless to say, there are a multitude of ways both strain and hyperconjugation can be quantified (see Gonthier et al. 2012 for orientation). However, in this case, I argue there is a more

⁷⁶ "Despite the extensive studies which have been made on a large number of molecules, it is still true that the theory of the origin of these potential barriers is not very far advanced" (Wilson 1959); "While it is easy enough to show that the barrier is not due to these special effects, it is more difficult to sort out of the mass of information a simple explanation of what the barrier *is* due to" (Cignitti & Allen 1964); "Despite the many theoretical papers published, no general agreement has yet been reached on an explanation for such barriers" (Lowe 1973). "So far, these attempts have led to rather negative conclusions, that is, many effects have been found to be not responsible for the barrier, but it has not been possible to pinpoint a simple explanation of what causes the barrier" (Veillard 1974). See (Bruck & Weinhold 1979) for this collection of quotes and citations therein.

⁷⁷ "When will chemistry textbooks begin to serve as aids, rather than barriers, to this enriched quantum-mechanical perspective on how molecular turnstiles work?" (Weinhold 2001); "The convenience of visually pleasing explanations for intricate chemical phenomena must not outweigh proper physical descriptions" (Schreiner 2002).

instructive story to be told, one that highlights an important limitation of concepts developed for qualitative understanding.

5.3.2 Intractable questions: an implication of working with concepts developed for qualitative understanding

My foregoing account of chemical concepts sheds light on why this debate has resisted a clean resolution, and why questions such as these—whether strain or hyperconjugation is more responsible for the barrier to rotation—are problematic when concepts are developed in the manner of the understanding-driven camp. In short, this type of question demands a quantitative functionality of the concepts that was sacrificed during the course of their development.

Note first, if it is assumed that concepts behave in the manner the rigor-driven camp thinks they should, this question is straightforward, and there ought to be a clear answer. According to this manner of conceptual development, concepts should have a clean and singular connection to the underlying physical models, making their quantification unambiguous and as accurate as the model itself. As such, answering the question about which concept is more responsible for the barrier to rotation in ethane should be as simple as carrying out these quantifications and comparing the resulting numbers.⁷⁸

⁷⁸ To be more specific, it should be a matter of quantifying strain in both the staggered and eclipsed conformations with the most physically motivated conceptions, and comparing the difference to the corresponding difference for hyperconjugation.

Matters are not so simple when concepts are developed to serve the aims of the understanding-driven camp. According to their preferred mode of conceptual development, the function of a chemical concept is first and foremost to track a qualitative trend in chemical data. It has done this when it can be used to generate a correct ordering of the relevant chemical data, allowing chemists to use the value of the concept to qualitatively predict some chemical properties or behavior (as electronegativity values can be used to predict relative bond strengths, relative polarity, etc.). In achieving this end, simplicity is sometimes opted for over precise quantification (also evident in the history of electronegativity). Besides this desideratum, additional constraints for conceptual development come from analyses of how its behavior qualitatively interacts with other chemical concepts that are believed to capture related properties (as AIM concepts were evaluated based on their qualitative compatibility with electronegativity considerations).

This mode of conceptual development leads to difficulty in answering the foregoing question about the barrier to rotation in ethane. The concept of strain allows for the qualitative assessment that the eclipsed conformation will be more strained, and therefore higher in energy, than the staggered conformation; similarly, the concept of hyperconjugation allows for the qualitative assessment that the staggered conformation will be more conjugated, and therefore lower in energy, than the eclipsed conformation. This is to say, both concepts corroborate the same qualitative assessment: the staggered conformation is more stable. In order to distinguish their relative contributions to this energy difference, however, it is required that we have an accepted quantitative measure of these concepts, which is precisely what is lacking when concepts are developed along the lines of the understanding-driven camp.

One might argue: yes, the understanding-driven camp prioritizes the generality of qualitative accuracy, but their concepts do allow for quantification—this is precisely what the various conceptions of their concepts do. Therefore, assessing the relative contribution of strain and hyperconjugation is as simple as figuring out which conception is appropriate for the respective concepts, and then simply comparing numbers. But in this case, how is one supposed to determine the appropriate conception of the respective concepts? In general, this would be done by determining which conception best tracks expected trends in other qualitative concepts (as we saw Saha et al. 2009 doing in Section 3.3). But what set of molecules will be used to make these assessments? Acyclic hydrocarbons would be a natural choice, but strain and hyperconjugation are present in all of them, making for a rather intractable situation.⁷⁹

One might consider this conceptual conundrum a strike against the conceptual development advocated by the understanding-driven camp. Is it not reasonable to expect our concepts to be applicable with quantitative precision? Perhaps it is, and perhaps the effort to develop energy decomposition analyses (see Fernández 2018 for introduction) that are aligned with the rigor-driven camp is a worthwhile effort that may pay off eventually. However, it is worth asking what is lost in giving up questions of this sort. Are chemists not in a fine position to use the concepts of strain and hyperconjugations for qualitative comparisons between different chemical system, as they currently

⁷⁹ The alternative of doing an energy decomposition analysis—and most attempts to settle this debate fall into this category—generally struggles to satisfy the aims of the understanding-driven camp, and for this reason struggles to generate widely accepted answers to questions such as this about the barrier to rotation in ethane.

can for the two conformations of ethane? At the very least, however, it is important to appreciate the different elements at play in the question of the barrier to rotation of ethane so that these longstanding debates might become less acrimonious and the true contours of the debate become clearer.

5.4 Auxiliary functions of concepts built for qualitative understanding

The discussion so far has focused on the practice of developing chemical concepts, especially the less familiar approach of the understanding-driven camp. In the previous section we saw one downside of concepts so developed: questions that rely on their quantitative comparison cannot be expected to have clean answers, and are liable to longstanding disagreements unless properly recognized as problematic. The present section briefly explores several upsides of concepts developed for qualitative understanding. Besides being qualitatively accurate across large swathes of chemistry, the concepts developed by the understanding-driven approach are better able to i) facilitate meta-discussion about computational models, and ii) guide the development of more accurate and more interpretable computational models.

I broach these topics not to provide a definite account of how the concepts are able to achieve this broader functionality, but as an indication of several philosophical problem areas native to chemistry on which philosophers of chemistry ought to focus their attention. These auxiliary functions of

concepts developed for qualitative understanding can also be viewed as points in favor of the understanding-driven camp's approach to conceptual development.⁸⁰

5.4.1 Concepts able to facilitate discourse about the many models of computational chemistry

The conceptual flexibility employed by the understanding driven camp allows them to maintain familiar conceptual behavior: bonds exist where they always have according to structural chemistry, electronegativity remains a qualitative predictor of atomic charge, and so on. Where a critic might charge them with merely affirming their presuppositions about the behavior of these concepts and their role in explaining chemical phenomena, this consistency with past theory has significant benefits in terms of both maintaining straightforward relevance of past chemical research and maintaining a relatively simple conceptual educational program.⁸¹

⁸⁰ Though perhaps these functions are not 'auxiliary' at all. In the given context, the several virtues of concepts developed by the understanding-driven camp appear to be coextensive, making it difficult to disentangle their respective appeal to chemists. Because the virtues discussed in the previous chapter have received the most attention, I treat them as the primary appeal, though with some uncertainty on my part.

⁸¹ It would be a mistake to say, though it sometimes is said, that chemists of the understanding-driven camp fail to extract 'understanding' from the computational models or the concepts of the rigor-driven camp due to their being unfamiliar with quantum mechanics (e.g., see Ogilvie 1990, Bader 1999, and Matta et al. 2006). This simply is not the case, though it is indicative of the divisiveness of the issue.

Perhaps more interestingly, concepts of this kind also enable meta-discussion about computational models. To appreciate this, first note the wide variety of models currently in use (gestured at in Chapter 1). To approximate the Schrödinger equation—the central task of most computational models in chemistry—approaches based in density functional theory have become widespread in the past decade, but there remain many others (see Cramer 2004 ch.7-8). Furthermore, DFT-based approaches encompass many different implementations that make different approximations, and these models cannot be neatly ordered by their accuracy because model accuracy is often a function of the class of chemistry being explored. On top of all of this, the choice of basis set provides additional variation in how the model crafts the shape of the wave function, and this choice is non-trivial both in terms of performance and whether the model provides a reliable means of rescuing chemical concepts.⁸²

With this plethora of models at play among the chemical community, it is important—perhaps essential—to have some way to discuss how the different models perform in terms of chemical concepts, for these concepts often frame the research being carried out. Perhaps one is investigating how a certain complex might be stabilized by charge transfer: one should seek a model that generates reasonable atomic charges; or perhaps one is investigating the role of dispersion interaction in facilitating a supramolecular process: one should seek a model that accurately captures this interaction. It is for this reason that model performance is often evaluated in terms of their ability to capture

⁸² One striking example is the use of a model to generate Mulliken (orbital-based) atomic charges. More realistic charges are generated when the model uses a small basis set, which in almost all cases sacrifices model accuracy in other ways (see Cramer 2004 section 9.1.3.2). Again it is worth noting that what is deemed ‘realistic’ involves presupposition about the way the concepts ought to behave.

behavior or energies that correspond to the same concepts that frame investigation. For example, DFT-based approaches are known to be inaccurate for systems where dispersion interaction energy is important (Grimme and Steinmetz 2013).⁸³

It is the flexibility of concepts developed for understanding that allows for this meta-discussion about computational models. One notable downside of defining concepts in the manner of the rigor-driven camp is that they are—by their very definition—tied to certain modeling practices. For example, if bonding interaction was strictly identified with the localized orbitals of valence bond theory, or even with the delocalized orbitals of molecular orbital theory, then we would be unable to read off bonding interactions from models based on density functional theory, which do not possess the appropriate two electron orbitals. Put differently, the natural way to recover these concepts from one computational model will differ from the natural way to recover it from another, and some ‘fuzziness’ is required to relate the performance of different models using the same chemical language.⁸⁴

⁸³ In fact, it is rather surprising that patterns appear in this regard: these concepts are not fed into the model design—hence the difficulty of getting them out at the end—and yet model performance can often be neatly summarized in terms of these concepts.

⁸⁴ Not unlike the ‘productive ambiguity’ of the concept of the gene according to Keller (2002).

5.4.2 Concepts able to guide the development of more accurate and more interpretable computational models

This dissertation has focused on the difficulty of recovering chemical concepts from computational models. However, this difficulty can be avoided if the chemical concepts are built into the computational model—that is, presumed in the very calculational procedure on which the model is based—and there are two distinct reasons they are sometimes so presumed.

First, chemical concepts can guide the patchwork correction of empirical computational models. Model error is often articulated in terms of chemical concepts; for example, and as noted in the previous section, DFT-based approaches are known to be inaccurate for systems where dispersion interaction energy is important. If we are dealing with an empirical model, one where parameters are ‘fit’ to match empirical data, then an interesting opportunity presents itself: we can augment the empirical model with parameters that are designed to capture this concept, and these parameters can then be fit to empirical data in way that, ideally, generalizes well (see Grimme, Ehrlich, and Goerigk 2011 for an example focusing on dispersion corrected models). Interestingly, it is our traditional understanding of how the concept is supposed to work that suggests the structure of this parameterization, and so how to create more accurate empirical models.

There is a second approach to designing computational models that focuses even more on traditional chemical concepts: *ab initio* valance bond theory (Raimondi et al. 1985; Shaik and Hiberty 2008). The idea here is that we have freedom in the way we approximate the Schrödinger equation, and we can spend this freedom to develop understandable models. Where most advances in computational modeling have added complexity to the way molecular orbitals are calculated, *ab initio* VBT seeks to

maintain local atomic orbitals and the host of intuitive concepts rooted therein (discussed in Chapter 1). To overcome the poor accuracy of traditional VBT, flexibility is added to the behavior of the atomic orbitals (for example, see the 'breathing' atomic orbitals of: Hiberty et al. 1994; Hiberty and Shaik 2002). These modifications to atomic orbitals may modify the concepts based on them, but so long as their qualitative behavior remains intact, *ab initio* VBT offers a powerful tool for telling understandable stories in familiar terms (see Patil and Bhanage 2016 for a striking example).

Note how strange this practice would appear to someone committed to the conceptual development of the rigor-driven camp. If our concepts ought to be cleanly and singularly rooted in physical theory, *ab initio* VBT represents a last-ditch effort to preserve a theoretical framework that physics has already told us to abandon. Or more poetically, these practices involve the positing of “parallel universes wherein [physical] laws are either ignored or ‘bent’ to accommodate personal points of view” (Bader 2011).

However, I hope to have offered a compelling alternative view of conceptual practices in the computational era, one where understanding takes a central, and sometimes even guiding role.

6.0 Conclusion

I have shown that there is substantial and ongoing debate among theoretical chemists about the proper function of chemical concepts. Some theorists believe their primary function is to neatly summarize important elements of the underlying physical models, and others believe their primary function is to make simple qualitative patterns accessible.

I elucidated the different strategies for conceptual development theorists of these respective camps adopt, and showed that the resulting concepts possess a variety of different features besides those that guide their development. Concepts developed by the rigor-driven camp are physically observable, but often impose changes to traditional conceptual behavior. Concepts developed by the understanding-driven camp may make use of multiple quantification schemes in order to maintain consistency with the conceptual network in which they reside, and for this reason may lack a quantification with general significance.

The presence of these conflicting modes of conceptual development means that philosophers of chemistry must be careful when drawing on chemical practice to fuel philosophical inquiries. However, it also offers several lines of philosophical investigation that are properly aligned with active concerns among chemists. Thus, I hope to have identified a promising direction for the future of philosophy of chemistry.

Bibliography

- Averbakh, Yuir. 2012. *A History of Chess : From Chaturanga to the Present Day*. Russell Enterprises.
- Bachrach, Steven M. 2007. *Computational Organic Chemistry*. Wiley-Interscience.
- Bader, R. 1999. "Can There Be More than as Single Definition of an Atom in a Molecule?" *Canadian Journal of Chemistry-Revue Canadienne De Chimie* 77 (1): 86–93.
- Bader, Richard. 1990. *Atoms in Molecules : A Quantum Theory*. Clarendon Press.
- Bader, Richard F. W. 2009. "Bond Paths Are Not Chemical Bonds." *The Journal of Physical Chemistry A* 113 (38). American Chemical Society: 10391–96.
- . 2010. "Definition of Molecular Structure: By Choice or by Appeal to Observation?" *The Journal of Physical Chemistry A* 114 (28). American Chemical Society: 7431–44.
- Bader, Richard F.W. 2005. *The Quantum Mechanical Basis of Conceptual Chemistry. Monatshefte Fur Chemie*. Vol. 136.
- . 2008. "Nearsightedness of Electronic Matter as Seen by a Physicist and a Chemist." *Journal of Physical Chemistry A* 112 (51): 13717–28.
- . 2011. "On the Non-Existence of Parallel Universes in Chemistry." *Foundations of Chemistry* 13 (1): 11–37.
- Batsanov, S.S.a, and A.S.b Batsanov. 2012. *Introduction to Structural Chemistry. Introduction to Structural Chemistry*.
- Becke, Axel D. 2014. "Perspective: Fifty Years of Density-Functional Theory in Chemical Physics." *Journal of Chemical Physics* 140 (18).
- Bergmann, D, and Juergen Hinze. 1996. "Electronegativity and Molecular Properties." *Angew. Chem. Int. Ed. Engl.* 35: 150–63.
- Brock, W. H. (William Hodson). 1993. *The Norton History of Chemistry*. W.W. Norton.
- Brush, Stephen. 1999a. "Dynamics of Theory Change in Chemistry: Part 1. The Benzene Problem 1865–1945." *Studies in History and Philosophy of Science Part A* 30 (1). Pergamon: 21–79.
- . 1999b. "Dynamics of Theory Change in Chemistry: Part 2. Benzene and Molecular Orbitals, 1945–1980." *Studies in History and Philosophy of Science Part A* 30 (2). Pergamon: 263–302.

- Bultinck, Patrick, Christian Van Alsenoy, Paul W. Ayers, and Ramon Carbó-Dorca. 2007. "Critical Analysis and Extension of the Hirshfeld Atoms in Molecules." *Journal of Chemical Physics* 126 (14).
- Campitelli, Guillermo. 2013. "How Computers Changed Chess." *The Conversation*. 2013. <https://theconversation.com/how-computers-changed-chess-20772>.
- Cerpa, Erick, Andreas Krapp, Alberto Vela, and Gabriel Merino. 2008. "The Implications of Symmetry of the External Potential on Bond Paths." *Chemistry - A European Journal* 14 (33). Wiley-Blackwell: 10232–34. <https://doi.org/10.1002/chem.200800934>.
- Coulson, Charles. 1961. *Valence*. 1st ed. Oxford University Press.
- Cramer, Christopher J. 2004. *Essentials of Computational Chemistry Theories and Models. Essentials of Computational Chemistry*. Vol. 42.
- Davidson, Ernest R., and Subhas Chakravorty. 1992. "A Test of the Hirshfeld Definition of Atomic Charges and Moments." *Theoretica Chimica Acta* 83 (5–6): 319–30.
- Dvoretsky, M. I. (Mark Izrailevich), and Kenneth P. Neat. 2002. *School of Chess Excellence. 3, Strategic Play*. Edition Olms.
- Eberhart, Mark. 2002. "Quantum Mechanics and Molecular Design in the Twenty First Century." *Foundations of Chemistry*, 201–11.
- Eberhart, Mark E., and Travis E. Jones. 2013. "The Two Faces of Chemistry: Can They Be Reconciled?" *Foundations of Chemistry* 15 (3): 277–85.
- Esser, Stephen. 2018. "QTAIM and the Interactive Conception of Chemical Bonding," November. <http://philsciarchive.pitt.edu>
- Evans, Author M G. 1951. "Introduction to the Discussion A Discussion on Bond Energies and Bond Lengths" 207 (1088): 1–5.
- Fernández, Israel. 2018. "Energy Decomposition Analysis and Related Methods." In *Applied Theoretical Organic Chemistry*, 191–226. WORLD SCIENTIFIC (EUROPE).
- Fernández, Israel, Gernot Frenking, and Einar Uggerud. 2010. "Response to the Comment on "The Interplay between Steric and Electronic Effects in SN2 Reactions."" *Chemistry - A European Journal* 16 (19). Wiley-Blackwell: 5542–43.
- Ferreira, Diogo R. 2013. "The Impact of Search Depth on Chess Playing Strength." *ICGA Journal* 36 (2): 67–80.
- Fonseca Guerra, C., J. W. Handgraaf, E. J. Baerends, and F. M. Bickelhaupt. 2004. "Voronoi Deformation Density (VDD) Charges: Assessment of the Mulliken, Bader, Hirshfeld, Weinhold, and VDD Methods for Charge Analysis." *Journal of Computational Chemistry* 25 (2). Wiley-Blackwell: 189–210.

- Frenking, Gernot, and Andreas Krapp. 2007. "Unicorns in the World of Chemical Bonding Models." *Journal of Computational Chemistry* 28 (1). Wiley-Blackwell: 15–24.
- Gavroglou, Kōstas, and Ana Simões. 2012. *Neither Physics nor Chemistry : A History of Quantum Chemistry*. MIT Press.
- Geerlings, P., F. De Proft, and W. Langenaeker. 2003. "Conceptual Density Functional Theory." *Chemical Reviews* 103 (5): 1793–1874.
- Geerlings, Paul, Stijn Fias, Zino Boisdenghien, and Frank De Proft. 2014. "Conceptual DFT: Chemistry from the Linear Response Function." *Chemical Society Reviews* 43 (14). The Royal Society of Chemistry: 4989.
- Gonthier, Jérôme F., Stephan N. Steinmann, Matthew D. Wodrich, and Clémence Corminboeuf. 2012. "Quantification of 'Fuzzy' Chemical Concepts: A Computational Perspective." *Chem. Soc. Rev.* 41 (13): 4671–87.
- Grimme, Stefan, Jens Antony, Stephan Ehrlich, and Helge Krieg. 2010. "A Consistent and Accurate *Ab Initio* Parametrization of Density Functional Dispersion Correction (DFT-D) for the 94 Elements H-Pu." *The Journal of Chemical Physics* 132 (15): 154104.
- Grimme, Stefan, Stephan Ehrlich, and Lars Goerigk. 2011. "Effect of the Damping Function in Dispersion Corrected Density Functional Theory." *Journal of Computational Chemistry* 32 (7). Wiley-Blackwell: 1456–65.
- Grimme, Stefan, Christian Mück-Lichtenfeld, Gerhard Erker, Gerald Kehr, Huadong Wang, Helmut Beckers, and Helge Willner. 2009. "When Do Interacting Atoms Form a Chemical Bond? Spectroscopic Measurements and Theoretical Analyses of Dideuteriophenanthrene." *Angewandte Chemie International Edition* 48 (14). John Wiley & Sons, Ltd: 2592–95.
- Grimme, Stefan, and Marc Steinmetz. 2013. "Effects of London Dispersion Correction in Density Functional Theory on the Structures of Organic Molecules in the Gas Phase." *Phys. Chem. Chem. Phys. Phys. Chem. Chem. Phys* 15 (15): 16031–42.
- Haghi, A. K. 2012. *Modern Trends in Chemistry and Chemical Engineering*. Apple Academic Press.
- Heitler, W., and F. London. 1927. "Wechselwirkung Neutraler Atome Und Homöopolare Bindung Nach Der Quantenmechanik." *Zeitschrift Für Physik* 44 (6–7). Springer-Verlag: 455–72.
- Hendry, Robin Findlay. 2008. "Two Conceptions of the Chemical Bond." *Philosophy of Science* 75 (5). The University of Chicago Press/Philosophy of Science Association: 909–20.
- Hendry, Robin Findlay. 2004. "The Physicists, the Chemists, and the Pragmatics of Explanation." *Philosophy of Science* 71 (5): 1048–59.
- Hettema, Hinne. 2008. "A Note on Michael Weisberg's: Challenges to the Structural Conception of Chemical Bonding." *Foundations of Chemistry* 10 (2). Springer Netherlands: 135–42.

- Hiberty, Philippe C., Stéphane Humbel, Carsten P. Byrman, and Joop H. van Lenthe. 1994. "Compact Valence Bond Functions with Breathing Orbitals: Application to the Bond Dissociation Energies of F₂ and FH." *J. Chem. Phys.* 101 (7): 5969–76.
- Hiberty, Philippe C., and Sason Shaik. 2002. "Breathing-Orbital Valence Bond Method - A Modern Valence Bond Method That Includes Dynamic Correlation." *Theoretical Chemistry Accounts* 108 (5): 255–72.
- Hirshfeld, F. L. 1977. "Bonded-Atom Fragments for Describing Molecular Charge Densities." *Theoretica Chimica Acta* 44 (2): 129–38. <https://doi.org/10.1007/BF00549096>.
- Hoffmann, Roald. 1998. "Qualitative Thinking in the Age of Modern Computational Chemistry—or What Lionel Salem Knows." *Journal of Molecular Structure: THEOCHEM* 424 (1–2): 1–6.
- Hoffmann, Roald, Sason Shaik, and Philippe C. Hiberty. 2003. "A Conversation on VB vs MO Theory: A Never-Ending Rivalry?" American Chemical Society.
- Houk, K. N., and Peng Liu. 2014. "Using Computational Chemistry to Understand & Discover Chemical Reactions." *Daedalus* 143 (4): 49–66.
- Iczkowski, Raymond P., and John L. Margrave. 1961. "Electronegativity." *Journal of the American Chemical Society* 83 (17): 3547–51.
- Jansen, Martin, and Ulrich Wedig. 2008. "A Piece of the Picture-Misunderstanding of Chemical Concepts." *Angewandte Chemie International Edition* 47 (52). Wiley-Blackwell: 10026–29.
- Jensen, W. B. 1996. "Electronegativity from Avogadro to Pauling: Part 1: Origins of the Electronegativity Concept - Journal of Chemical Education (ACS Publications and Division of Chemical Education)." *Journal of Chemical Education*.
- Jensen, William. 2003. "Electronegativity from Avogadro to Pauling: II. Late Nineteenth- and Early Twentieth-Century Developments." *Journal of Chemical Education* 80 (3). Division of Chemical Education: 279.
- Keller, Evelyn Fox. 2002. *The Century of the Gene*. Harvard University Press.
- Klimeš, Jiří, and Angelos Michaelides. 2013. "Perspective: Advances and Challenges in Treating van Der Waals Dispersion Forces in Density Functional Theory," January.
- Kullback, S. 1959. *Information Theory and Statistics*. New York: Wiley.
- Lai, Matthew. 2015. "Giraffe: Using Deep Reinforcement Learning to Play Chess," September. <http://arxiv.org/abs/1509.01549>.
- Langmuir, Irving. 1919. "The Arrangement of Electrons in Atoms and Molecules." *Journal of the Franklin Institute* 187 (3). Pergamon: 359–62. [https://doi.org/10.1016/S0016-0032\(19\)91097-0](https://doi.org/10.1016/S0016-0032(19)91097-0).
- Levy, David N. L. 2009. *Computer Chess Compendium*. Ishi Press.

- Lewis, David. 1973. *Counterfactuals*. Blackwell Publishers.
- Lewis, Gilbert. 1916a. "Steric Hindrance and the Existence of Odd Molecules (Free Radicals)." *Proceedings of the National Academy of Sciences of the United States of America* 2 (10): 586–92.
- . 1916b. "The Atom and the Molecule." *Journal of the American Chemical Society* 38 (4). American Chemical Society: 762–85.
- London, F. 1928. "On the Quantum Theory of Homo-Polar Valence Numbers." *Z. Phys* 46 (455).
- Marvin Cohen. 2003. *Walter Kohn: Personal Stories and Anecdotes Told by Friends and Collaborators - Google Books*. Edited by Matthias Scheffler and Peter Weinberger. Berlin: Springer-Berlag.
- Matta, Cherif F, and Richard F W Bader. 2006. "An Experimentalist's Reply to 'What Is an Atom in a Molecule'" *The Journal of Physical Chemistry A* 110 (19): 6365–71.
- Mulliken, R. S. 1955. "Electronic Population Analysis on LCAO–MO Molecular Wave Functions. I." *The Journal of Chemical Physics* 23 (10): 1833–40.
- Mulliken, Robert S. 1932. "Electronic Structures of Polyatomic Molecules and Valence. II. General Considerations." *Physical Review* 41 (1). American Physical Society: 49–71.
- . 1935a. "Electronic Structures of Molecules XI. Electroaffinity, Molecular Orbitals and Dipole Moments." *The Journal of Chemical Physics* 3 (9): 573–85.
- . 1935b. "Electronic Structures of Molecules XI. Electroaffinity, Molecular Orbitals and Dipole Moments." *The Journal of Chemical Physics* 3 (9). American Institute of Physics: 573–85.
- Nimzowitsch, Aron, and Lou Hays. 1991. *My System : 21st Century Edition*. Hays Pub.
- Ogilvie, J. F. 1990. "The Nature of the Chemical Bond—1990: There Are No Such Things as Orbitals!" *Journal of Chemical Education* 67 (4): 280.
- Parr, Robert G., Robert A. Donnelly, Mel Levy, and William E. Palke. 1978. "Electronegativity: The Density Functional Viewpoint." *The Journal of Chemical Physics* 68 (8): 3801–7.
- Patil, Amol Baliram, and Bhalchandra Mahadeo Bhanage. 2016. "Modern Ab Initio Valence Bond Theory Calculations Reveal Charge Shift Bonding in Protic Ionic Liquids." *Physical Chemistry Chemical Physics* 18 (23). The Royal Society of Chemistry: 15783–90.
- Pauli, W. 1927. "Zur Quantenmechanik Des Magnetischen Elektrons." *Zeitschrift Für Physik* 43 (9–10). Springer-Verlag: 601–23.
- Pauling, Linus. 1931a. "The Nature of the Chemical Bond. Application of Results Obtained from the Quantum Mechanics and from a Theory of Paramagnetic Susceptibility to the Structure of Molecules." *Journal of the American Chemical Society* 53 (4). American Chemical Society: 1367–1400.

- . 1931b. “The Nature of the Chemical Bond II. The One-Electron Bond and the Three-Electron Bond.” *Journal of the American Chemical Society* 53 (9). American Chemical Society: 3225–37.
- Pauling, Linus. 1932a. “The Nature of the Chemical Bond III. The Transition from One Extreme Bond Type to Another.” *Journal of the American Chemical Society* 54 (3). American Chemical Society: 988–1003.
- . 1932b. “The Nature of the Chemical Bond IV. The Energy of Single Bonds and the Relative Electronegativity of Atoms.” *Journal of the American Chemical Society* 54 (9). American Chemical Society: 3570–82.
- . 1939. *The Nature of the Chemical Bond and the Structure of Molecules and Crystals : An Introduction to Modern Structural Chemistry*. 1st ed. Cornell University Press.
- . 1940. *The Nature of the Chemical Bond and the Structure of Molecules and Crystals: An Introduction to Modern Structural Chemistry*. 2nd ed. Ithaca, NY: Cornell University Press.
- . 1960. *The Nature of the Chemical Bond and the Structure of Molecules and Crystals: An Introduction to Modern Structural Chemistry*. 3rd ed. Ithaca, N.Y.: Cornell University Press.
- Pritchard, H. O., and H. A. Skinner. 1955. “The Concept of Electronegativity.” *Chemical Reviews* 55 (4): 745–86.
- Raimondi, M., M. Simonetta, and G.F. Tantardini. 1985. “AB Initio Valence Bond Theory.” *Computer Physics Reports* 2 (4). North-Holland: 171–216.
- Regt, Henk W. de. 2017. *Understanding Scientific Understanding*. Vol. 1. Oxford University Press.
- Robert G. Parr, *,†, ‡ and Paul W. Ayers, and Roman F. Nalewajski§. 2005. “What Is an Atom in a Molecule?” American Chemical Society .
- Saha, Soumen, Ram Kinkar Roy, and Paul W. Ayers. 2009. “Are the Hirshfeld and Mulliken Population Analysis Schemes Consistent with Chemical Intuition?” *International Journal of Quantum Chemistry* 109 (9). Wiley-Blackwell: 1790–1806.
- Sanderson, R. T. 1983. “Electronegativity and Bond Energy.” *Journal of the American Chemical Society* 105 (8): 2259–61. <https://doi.org/10.1021/ja00346a026>.
- Sanderson, R. T. 1955. “Relation of Stability Ratios to Pauling Electronegativities.” *The Journal of Chemical Physics* 23 (12): 2467.
- Sardesai, Rajdeep, and Viswanathan Anand. 2013. “Viswanathan Anand Video Interview with CNN.”
- Schwarz, W. H. Eugen, and Hubert Schmidbaur. 2012. “Observations and Descriptions versus Explanations-An Example: Does Nature, Does Theory Know About Steric Hindrance?” *Chemistry - A European Journal* 18 (15). Wiley-Blackwell: 4470–79.

- Shaik, Sason. 2007. "Is My Chemical Universe Localized or Delocalized? Is There a Future for Chemical Concepts?" *New Journal of Chemistry* 31 (12). The Royal Society of Chemistry: 2015.
- Shaik, Sason S., and Philippe C. Hiberty. 2008. *A Chemist's Guide to Valence Bond Theory*. Wiley-Interscience.
- Silman, Jeremy. 2014. "A Chess Engine Is Not Your Friend." 2014. <https://www.chess.com/article/view/chess-engines-are-not-your-friend>.
- Vanpoucke, Danny E P, Patrick Bultinck, and Isabel Van Driessche. 2013. "Extending Hirshfeld-I to Bulk and Periodic Materials." *Journal of Computational Chemistry* 34 (5): 405–17.
- Weinhold, Frank, and Clark R. Landis. 2012. *Discovering Chemistry with Natural Bond Orbitals*. Wiley.
- Weisberg, Michael. 2008. "Challenges to the Structural Conception of Chemical Bonding." *Philosophy of Science*, 75: 932–946.
- Weisberg, Michael, Paul Needham, and Robin Hendry. 2011. "Philosophy of Chemistry." *Stanford Encyclopedia of Philosophy*. Stanford University. <https://plato.stanford.edu/entries/chemistry/>.