

Qualitative Theory and Chemical Explanation

Michael Weisberg^{†‡}

Roald Hoffmann and other theorists claim that we ought to use highly idealized chemical models (“qualitative models”) in order to increase our understanding of chemical phenomena, even though other models are available which make more highly accurate predictions. I assess this norm by examining one of the tradeoffs faced by model builders and model users—the tradeoff between precision and generality. After arguing that this tradeoff obtains in many cases, I discuss how the existence of this tradeoff can help us defend Hoffmann’s norm for modelling.

1. Qualitative Models in Chemistry. Roald Hoffmann and other chemists have claimed that highly idealized models are extremely important and probably indispensable for chemical explanations.¹ Such highly idealized or “qualitative” models are often contrasted with “quantitative” models, which are more predictively accurate. The distinction between qualitative and quantitative models is not about the use of numbers; both types of models can be numerical. Rather it is a distinction resting on degrees of approximation and idealization. Qualitative models contain more ap-

[†]To contact the author, please write to: Department of Philosophy, University of Pennsylvania, 433 Logan Hall, Philadelphia, PA 19104; e-mail: weisberg@phil.upenn.edu.

[‡]Many thanks to Michael Friedman, Peter Godfrey-Smith, Robin Hendry, Ben Kerr, Deena Skolnick, and Michael Strevens for extremely helpful comments on earlier drafts of this and related papers. John Brauman, Marc Feldman, Paul Needham, Joan Roughgarden, Janet Stemwedel, and Ward Watt also provided helpful comments in conversations about these themes. Most of all I would like to thank Roald Hoffmann who has been an inspiring mentor and whose reflections on the philosophical aspects of chemistry stimulated my interest in this topic. This research was partially supported by a National Science Foundation Pre-Doctoral Fellowship.

1. This argument can be found in Hoffmann 1995, 1998. One of the classic examples of the use of qualitative models in chemical explanations can be found in Hoffmann’s joint work with Woodward on pericyclic reactions (Woodward and Hoffmann 1970).

Philosophy of Science, 71 (December 2004) pp. 1071–1081. 0031-8248/2004/7105-0037\$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

proximations and are more highly idealized than quantitative models.² This contrast is especially striking in the literature about molecular and electronic structure. Increases in computational power have given chemists the ability to calculate the electronic structure of small- and medium-sized molecules with incredible accuracy using quantitative models. Theorists who emphasize the need for qualitative models acknowledge the importance and achievements of the quantitative modelling tradition, but emphasize the continuing need for qualitative models to increase our understanding of chemical phenomena. Qualitative models of molecular structure are rarely able to make highly accurate predictions about molecular geometry; however, they are thought to explain why molecules have the shapes that they do.

One of the clearest statements of the continuing need for qualitative models comes from a recent physical organic chemistry textbook. Felix Carroll writes:

Why then don't we just talk about high-level theoretical calculations and ignore the simple theory? We must choose the model that is sufficiently accurate for our computational purposes, yet still simple enough that we have some *understanding* of what the model describes. Otherwise, the model is a black box, and we have no understanding of what it does, perhaps even no idea whether the answers it produces are physically reasonable. (Carroll 1998, 27; my emphasis)

In this passage, Carroll asserts a connection between building qualitative models and an increase in explanatory power.³

Before we can analyze the connection between idealization and explanation in greater detail, we need to know what Carroll, Hoffmann, and others mean by “qualitative models” or “simple models.” These expressions refer to a number of things including:

1. literally simpler models (i.e. ones that have fewer parameters in their descriptions).
2. models with restrictions on the number of causal factors included.

2. Although I will be treating the qualitative/quantitative distinction as a distinction about kinds of models, there is another way to render it. Hoffmann and his philosophical allies sometimes apply the qualitative/quantitative distinction to the practice and goals of modelling. For example, Hoffmann has claimed that models he relies on in his research exemplify “qualitative thinking” about chemical problems. In Weisberg 2003, I discuss the goal version of this distinction in greater detail.

3. Robin Hendry (2004) argues that the use of highly idealized models in chemical contexts goes back to the earliest days of quantum chemistry. He claims that Coulson advocated the use of simplified quantum mechanical models in order to generate chemical explanations that were both plausible and chemically intuitive.

3. models appraised with lower standards of fidelity.
4. imprecise model descriptions.

While each item on this list may play a role in justifying the claim that qualitative models are necessary for increasing the explanatory power of chemistry, I will be confining my discussion to the fourth item—precision. I will argue that sacrificing precision in model descriptions can often add explanatory depth to the models picked out by these descriptions. The defense of this thesis has two parts. First I will argue that there is a tradeoff between precision and generality. I then suggest that generality is an explanatory virtue. Hence, a more general set of models is a more explanatory set. Combining these claims generates the conclusion that imprecise model descriptions can pick out more explanatory sets of models.

2. Models. Philosophers defending familiar, simple forms of scientific realism have often argued for a two-place relationship between theoretical representations and the world. These two relata are connected via correspondence or truth; a good representation truthfully describes a target system in the world. According to these philosophers, understanding the practice of theorizing in science only requires understanding the structure of theories and the relationships between theories and the world.

Giere (1988), Cartwright (1983), Lloyd (1994), and other philosophers have emphasized the inadequacy of this view to capture aspects of modern theoretical practice. Giere argues that there are not two, but three relevant relata: models, model descriptions, and the world. My own account of modeling broadly follows Giere on this point.

Model descriptions describe or pick out models. In the case of mathematical models, model descriptions typically take the form of equations. This is optional, however, as models can be represented in other ways as well, such as with pictures, sentences, or computer programs. A single model description can pick out multiple models, and one model may be described by multiple descriptions.

Models themselves can be concrete or abstract. Concrete models, such as model airplanes, are physically constructed. In this paper, however, I will be discussing abstract or mathematical models. These models consist of a set of relationships between properties and many show how these properties change over time. For example, dynamic models of molecular structure show how molecules stretch, bend, and rotate in time.

Models, both concrete and abstract, need not be related to any target system in the actual world. It may be theoretically fruitful to study a model even if it does not describe an actual target system. However, chemists often build models with the intention of relating them to actual

target systems. Such models may be used to make predictions about how these systems will behave and how their component parts give rise to their behavior. When models relate to target systems, they do so in virtue of their being *similar* to these systems in certain respects and degrees (Giere 1988).

For the purposes of this paper, I will describe a very simple example of a chemical model: the ball and spring model of covalent bonds. This model can be used to determine the stretching energy of bonds and is often used in connection with spectroscopy and semi-empirical calculations of molecular structure.⁴ Although this model is used to describe the quantum mechanical phenomenon of covalent bonding, the model itself is completely classical. It treats a covalent bond as a spring and atoms as masses at the ends of the spring. In the simplest version which I will be discussing in this paper, the spring is treated as a harmonic oscillator; it is an “ideal” spring with no anharmonic character. These features should make it obvious that the model contains many kinds of idealizations, although I will only be discussing the ones having to do with precision.

Ball and spring models are described by the following model description:

$$E_{\text{stretch}} = k_b(r - r_0)^2. \quad (1)$$

The independent variable in this equation is the distance between the atoms. It is expressed as the distance between the atoms (r) minus the equilibrium distance (r_0) so that when the bond is at its equilibrium length, the model predicts zero stretching energy.

Equation (1) is an uninstantiated model description, an equation in which values are not assigned to the parameters. Instantiating a model description means adding in values for the parameters.⁵ We can instantiate the ball and spring model description in various ways by setting the parameter (k_b) to different values. For example:

$$E_{\text{stretch}} = 1.0(r - r_0)^2, \quad (2)$$

$$E_{\text{stretch}} = 1.5(r - r_0)^2, \quad (3)$$

4. For more information about the use of semi-empirical models in calculating molecular structure, see Carroll 1998; Carey and Sundberg 2000; Lowry and Richardson 1987.

5. I roughly follow Orzack and Sober (1993) in making the distinction between instantiated and uninstantiated model descriptions. The main difference between our accounts is that Orzack and Sober treat equations as models, whereas I treat them as model descriptions. Thus they distinguish between instantiated and uninstantiated models, whereas I distinguish between instantiated and uninstantiated model descriptions.

$$E_{\text{stretch}} = 2.0(r - r_0)^2. \quad (4)$$

Each instantiated description picks out a different model from the ball and spring family.

3. Precision. Precision is a property of model descriptions, not of models themselves.⁶ Theoreticians sometimes talk about models being more or less precise, but what they usually mean is that some parameter in an equation is specified more or less precisely. Model builders' use of the term "precision" is closely related to the everyday definition "fineness of specification."

Comparing the amount of precision in the specification of individual parameters is relatively straightforward. Compare, for example, the following two instantiations of the ball and spring stretching model:

$$E_{\text{stretch}} = 1.0 \pm 0.1(r - r_0)^2, \quad (5)$$

$$E_{\text{stretch}} = 1.01 \pm 0.01(r - r_0)^2. \quad (6)$$

The first instantiation (5) is less precise than the second (6). Whereas the first description picks out all of the models with spring constants between 0.9 and 1.1, the second picks out only those models with spring constants between 0.99 and 1.01. In other words, the second description picks out a small subset of the models picked out by the first description.

This example suggests a general method of comparing the precision of two model descriptions, which will work when the two sets of models "overlap" as they do in (6) and (5):

A model description D_1 is more *precise* than a model description D_2 if D_1 picks out a proper subset of the models picked out by D_2 .

One can add further mathematical sophistication to this relational definition of precision, but for the purposes of this paper, it will be sufficient to adopt this simple comparative definition. Giving a more general and non-relational definition turns out to be complicated and is beyond the scope of this paper.

I can now be a bit more exact about how I individuate models and model descriptions. Model descriptions can be uninstantiated or instantiated, while models themselves are determinate sets of relationships be-

6. One might worry that precision is neither a property of models nor model descriptions and argue that it is a property of sets of data. The strict statistical definition of "precision" does make it a property of data sets. However, there is also a use of the term "precision" having to do with theoretical representations such as model descriptions. It is this use which I am trying to capture in this paper.

tween properties. Uninstantiated model descriptions pick out families of models. If the parameters are set with some imprecision, such that they can take a range of values, then the description will pick out a subset of the models picked out by the uninstantiated description. When model descriptions are instantiated with high degrees of precision, they pick out smaller sets of models. If we instantiate the parameters with completely precise values, then the instantiated description will pick out a single model.

4. Generality. We now turn to generality, which is a property of the relationship between models and target systems. Since I have not given a full account of this relationship, I will have to rely on an intuitive notion of “applying to” in order to discuss the tradeoffs involving generality. A model *applies to* a target system when it accurately describes the structure and dynamics of the system according to the standards set by the model builder or model user. Although we would need to know more about a model user’s standards in order to make absolute judgments about whether a particular model applies to a particular target system, for the purpose of this paper it will be enough to make comparative judgments between two sets of models for which the same standards are being applied.

The term “generality” is often used ambiguously by theoreticians. The ambiguity consists in two different senses of generality that we might recognize. One sense of generality has to do with how many *actual* target systems a model applies to. The other has to do with how many *logically possible* target systems a model applies to. Let us distinguish between *a-generality* and *p-generality* which track these two senses. A-generality is a measure of the number of *actual* target systems a particular model applies to. P-generality is a measure of how many *logically possible* target systems a particular model applies to.

5. Tradeoff between Precision and Generality. The easiest way to explain the relationship between precision and generality is to begin with an intuitive explanation and then develop this explanation in more detail. The intuitive idea is that when a model description is more precise, it picks out a smaller set of models. This smaller set of models applies to fewer target systems. If, on the other hand, the model description is imprecise in the way it picks out models, then the description will pick out a larger set of models. Since the set of models is larger, these models will apply to many more target systems. I think that we can understand the tradeoff between precision and generality roughly in this way, but to further articulate this point, we will have to consider precision’s relationship to p-generality and a-generality individually.

5.1. *Precision versus P-generality.* In order to assess the relationship between precision and p-generality, let's consider two instantiations of the ball and spring model:

$$E_{\text{stretch}} = 2.0 \pm 0.1(r - r_0)^2, \quad (7)$$

$$E_{\text{stretch}} = 2.01 \pm 0.01(r - r_0)^2. \quad (8)$$

The second instantiation (8) has a greater degree of precision than the first instantiation (7). Unlike the first instantiation which merely bounds the spring constant k_b between 1.9 and 2.1, the second instantiation bounds k_b between 2.00 and 2.02. The set of models picked out by the second instantiation is a proper subset of the models picked out by the first, hence the second is more precise than the first.

Although we cannot say exactly how many logically possible systems the two sets of models described by (8) and (7) apply to, we can make a comparison between them. Since the models described by (8) are a proper subset of the models described by (7), for any logically possible target system that a model in the (8) family applies to there will also be a model in the (7) family that applies to the same target system, but not vice versa. Thus the (7) family is more p-general than the (8) family. Many more logically possible bond vibrations exhibit the behavior described with a model picked out by (7) than by (8).

When we put these results together, we can see that there is a tradeoff between precision and p-generality. As we make model descriptions more precise, the models that these descriptions pick out will apply to fewer logically possible target systems. Precision and p-generality cannot be simultaneously increased. Thus we can conclude that there is a tradeoff between these two properties.

5.2. *Precision versus A-generality.* Since it is the world along with our standards of evaluation that determines the a-generality of a model, one might be skeptical that there is any universal relationship between precision and a-generality.

It is worth articulating this criticism in more detail, although I think it overlooks some important conclusions that we can draw about the relationship between precision and a-generality. As we have just seen, there is a tradeoff between precision and p-generality. We know, however, that all of the actual target systems constitute a very small part of possibility space. If actual systems are evenly distributed in possibility space, then the tradeoff between precision and p-generality might also hold for precision and a-generality automatically. However, we do not have any way of knowing a priori that systems are distributed relatively evenly in this

possibility space and hence we have no reason to conclude that precision trades off against a-generality. Thus the criticism states that the world will have to tell us in each particular case whether this tradeoff obtains.

It is correct to insist that only the world can tell us how a-general some model is. However, some facts about the world may help us determine what kind of relationship can hold between precision and a-generality for particular types of models.

This point can be illustrated with the help of a vivid example. Imagine we were building models of ecosystems. Say that we knew the world we lived in had extremely complex and heterogenous ecosystems, meaning that many different kinds of ecosystems were found in the world. If this was true, as we made our model descriptions more precise, they would pick out a set of models that applied to fewer and fewer target systems. If the set of ecosystems was fairly homogenous, then a very small set of models described by very precisely instantiated equations could apply to many ecosystems.

Say we lived in a desert world, where essentially the same ecosystem was repeated over and over again. In whatever way we carved up the world into discrete ecosystems, a very small set of models with a correspondingly precise model description would apply very generally in this world. On the other hand, in a world of varied and complex ecosystems like our own, a small set of models described by highly precise equations will only apply to a small number of ecosystems. Thus one factor that a-generality depends on is the homogeneity of the set of target system the model is intended to apply to.

A-generality also depends on a property that I call *scope*. *Scope* is defined as the aspects of target system(s) intended to be represented by a model. For example, even within models of ecosystems, we might restrict our scope to the sizes of populations alone. Here we would track intrinsic growth and death rates, predator-prey interactions, and other factors associated with the size of particular populations in an ecosystem. For example, a model with broader scope might include information about the foraging behavior of particular populations.

From the point of view of a model user with her intended scope, a set of target systems can appear more or less similar depending on which factors are actually included in the scope of the model. For example, in our world, basic physical phenomena such as gravity and electromagnetism affect objects in the same way everywhere. Even highly specific models of how these phenomena affect particular kinds of target systems can have fairly general applicability. On the other hand, target systems in the world look a lot more heterogenous if a model builder's scope includes the kinds of features often of interest to chemists—such as the

differences one finds as one moves down a group on the periodic table or as one starts with methanol and adds CH_2 's.

These considerations suggest that in certain cases, where either the choice of intended target systems or the choice of scope ensures homogeneity among target systems, precision does not trade off against a-generality. Only a very careful investigation of the extent to which different kinds of scopes affect homogeneity among target systems could determine whether precision trades off against a-generality for a particular set of target systems. I think that this is an extremely worthwhile project, but it will require careful empirical study of exactly how complex different target phenomena in our world really are. Another benefit of this detailed study is that we could learn the extent to which the tradeoff between precision and a-generality is a global phenomenon.

Even without this study, however, we can make a weak claim about the global relationship between precision and a-generality. We can claim that precision is an *attenuating factor* of a-generality. By "attenuating factor," I mean that increases in precision makes the achievement of a-generality more difficult, but not impossible.

Many of the issues I discussed in connection with the tradeoff between precision and a-generality are relevant to this claim about attenuation. Of particular relevance is the relationship between scope and a-generality. For any small set of models and set of intended target systems, there is likely to be some scope that will make this set of models very a-general. However, it is far more likely that any given scope will result in these models lacking a-generality. Recall that small sets of models are picked out by highly precise model descriptions. This means that for many scopes, highly precise model descriptions will pick out models that lack a-generality. Thus we can conclude that increasing precision attenuates a-generality.

6. Generality and Chemical Explanation. We now return to the connection between qualitative models and chemical explanation. Recall the passage I quoted from the physical organic chemistry textbook, which suggested a norm for model building: If you want to generate chemical explanations, then you should build qualitative models. Since we are only considering one aspect of qualitative modelling, precision, we could read this norm as suggesting that one could rationally sacrifice some precision in model descriptions in order to gain explanatory power. I believe that my analysis of the tradeoff between precision and a-generality gives us some of the resources we need to explain this assertion. The key comes from associating generality with explanatory power.

Many philosophers have thought generality is an important desideratum for scientific explanation. This is true of philosophers who base their

entire account of explanation on unity and generality (Kitcher 1981; Friedman 1974), logical empiricists (Hempel 1965), and some modern causal theorists (Strevens 2004). I will not defend the connection between explanatory depth and generality in this essay, but I do want to note that many theorists giving quite different accounts of explanation have seen generality as an important explanatory desideratum.

There are also reasons internal to chemistry for considering generality to be an explanatory virtue. Chemists are often interested in trends across similar but slightly different target systems. For example, chemists might want to explain the similarities in molecular structure for a set of alcohol molecules. General models can be used to explain the structure of these molecules in a unified way, allowing us to make relevant comparisons between them using the same basic framework. If we build highly precise models which are not very general, then their explanatory value for comparing trends would be lost.

Given that generality is an explanatory virtue in chemistry, we can justify part of the model building norm defended by Hoffmann, Carroll, and others. Sacrificing precision typically allows us to gain generality. This sacrifice is therefore a good way to increase an explanatory virtue, namely generality. If we have the goal of trying to offer an explanation of the structure and dynamics of a target system, then it is rational to idealize by sacrificing precision to gain generality and thereby increase the explanatory power of our models.

7. Conclusions. In this paper, I have discussed an important component of modern theoretical chemistry which focuses on the construction of highly idealized models. Qualitative theory connects the construction and use of idealized models with an increase in explanatory power. I argued that precision trades off with p-generality such that one cannot simultaneously increase both of these properties. I also argued that precision attenuates a-generality, meaning that a high degree of precision makes the achievement of a-generality more difficult. These relationships are important because they can help us defend a model-building norm that associates explanatory depth with the sacrifice of precision, an aspect of qualitative theorizing.

REFERENCES

- Carey, Francis A., and Richard J. Sundberg (2000), *Advanced Organic Chemistry A*, 4th ed. New York: Kluwer Academic, Plenum Publishers.
- Carroll, Felix A. (1998), *Perspectives on Structure and Mechanism in Organic Chemistry*. Pacific Grove, CA: Brooks/Cole.
- Cartwright, Nancy (1983), *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Friedman, Michael (1974), "Explanation and Scientific Understanding", *Journal of Philosophy* 71: 5–19.

- Giere, Ronald N. (1988), *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
- Hempel, Carl G. (1965), "Aspects of Scientific Explanation", in *Aspects of Scientific Explanation and Other Essays*. New York: Free Press, 331–496.
- Hendry, Robin (2004), "The Physicists, the Chemists, and the Pragmatics of Explanation", *Philosophy of Science* 71: 1048–1059.
- Hoffmann, Roald (1995), *The Same and Not the Same*. New York: Columbia University Press.
- (1998), "Qualitative Thinking in the Age of Modern Computational Chemistry: Or What Lionel Salem Knows", *THEOCHEM—Journal of Molecular Structure* 424: 1–6.
- Kitcher, Philip (1981), "Explanatory Unification", *Philosophy of Science* 48: 507–331.
- Lloyd, Elisabeth A. (1994), *The Structure and Confirmation of Evolutionary Theory*. Princeton, NJ: Princeton University Press.
- Lowry, Thomas H., and Kathleen S. Richardson (1987), *Mechanism and Theory in Organic Chemistry*, 3rd ed. New York: Harper and Row.
- Orzak, Stephen Hecht, and Elliott Sober (1993), "A Critical Assessment of Levin's 'The Strategy of Model Building in Population Biology' (1966)", *Quarterly Review of Biology* 68 (4): 533–546.
- Strevens, Michael (2004), "The Causal and Unification Accounts of Explanation Unified—Causally", *Nous* 38: 154–179.
- Weisberg, Michael (2003), *When Less Is More: Tradeoffs and Idealization in Model Building*. Ph.D. dissertation. Stanford, CA: Stanford University.
- Woodward, Robert B., and Roald Hoffmann (1970), *The Conservation of Orbital Symmetry*. Weinheim/Bergst.: Verlag Chemie.