Quantum Chemistry and Organic Theory

William Goodwin*

In this essay I consider whether the theory of organic chemistry is reducible to the theory of quantum chemistry. Using philosophical machinery developed by James Woodward, I characterize the understanding provided by both theories. Then I argue that there are systematic reasons to suspect that quantum chemistry is incapable of supporting some of the significant explanations, predictions, and applications underwritten by an understanding of theoretical organic chemistry. Consequently, even should quantum chemistry be 'reducible to' quantum physics in some suitable sense, there are good reasons to doubt that many of the significant results of organic chemistry could be reproduced by quantum chemistry alone.

1. Introduction. In this essay, I consider whether, or in what sense, mainstream theoretical organic chemistry is reducible to quantum chemistry. This will not only open up a fresh perspective on the general issue of the extent to which chemistry and physics are unified but also will directly resonate with a live philosophical issue for chemists—should they anticipate the continued autonomy of theoretical organic chemistry? In order to approach this question, I will first try to characterize what it means to understand theoretical organic chemistry. Only by developing a relatively rich appreciation of what mastering this theory allows organic chemists to do can one then go and to ask whether a reducing theory, such as quantum chemistry, could ever hope to "accomplish the same purposes" (Garfinkel 1991, 443) as theoretical organic chemistry.

For the purposes of this essay, I will characterize understanding organic chemistry as the ability to explain and predict some organic phenomena using the theory along with the ability to apply it in the solution of practical problems, such as designing the synthesis of a novel compound. In order to characterize the understanding provided by both theoretical organic chemistry and quantum chemistry, I will employ some of the philosophical machinery developed by James Woodward in his manipulationist account of

*To contact the author, please write to: Department of Philosophy, University of South Florida; e-mail: wgoodwin@usf.edu.

Philosophy of Science, 80 (December 2013) pp. 1159–1169. 0031-8248/2013/8005-0002$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/673734 Published online by Cambridge University Press

causal explanation (2003) because this perspective is useful in bringing out the differences between these two theories. Once these differences have been brought out, I will then argue that there are systematic reasons to suspect that quantum chemistry is incapable of supporting some of the significant explanations, predictions, and applications that constitute an understanding of theoretical organic chemistry. The upshot, then, is that even should quantum chemistry be 'reducible to' quantum physics—in some suitable sense—there are good reasons to doubt that many of the significant results of theoretical organic chemistry could ever be reproduced by quantum chemistry alone.

Explaining phenomena, according to Woodward, requires situating them within "systematic patterns of counterfactual dependence" (2003, 191). That is, we explain why something happened by describing the ways that the world might have been different so that the outcome would have changed in some specific sense. In order to do this, scientists invoke claims about the possible ways that the world might be. Sometimes scientists do this by invoking laws of nature, which are relationships of broad scope that hold true no matter how the world might be. However, all that is really required, on Woodward's view, is an appeal to some sort of modally robust generalization. Such generalizations may be limited in scope, but as long as they imply that if certain changes were made, the world would have been different in certain specific ways, this is enough to support scientific explanations. Explanations are useful because in order to explain a phenomenon, one must be able to identify interventions that have the capacity to change the phenomenon in some specific way. That is, having an explanation allows one to predict how the phenomenon would be changed if one were able to alter the conditions by such an intervention. Such predictions are at least part of what allows scientists to apply their knowledge in the solution of practical problems.

According to Woodward, the changes in circumstances that are particularly important to a scientific discipline "help to set the explanatory agenda" (2003, 262) for that field. Scientists concerned with changes of a particular sort, therefore, seek modally robust generalizations that hold up across those changes. As a result, "different choices of variables for theorizing are associated with different ways of carving up nature into possible alternatives, answers to different why-questions, and hence to different explanations" (233). On Woodward's view, it is possible for there to be scientific disciplines that investigate the same objects but which have distinct explanatory agendas with respect to those objects. Metaphorically, these disciplines would "carve up nature" into different pieces, according to the sorts of changes to which they were interested in subjecting those objects. To explain why their objects have certain characteristics (to answer why-questions), these scientists must consider what changes, among the array of those that form their explanatory agenda, might modify those characteristics should they actually be im-

https://doi.org/10.1086/673734 Published online by Cambridge University Press

plemented. This requires appeal to generalizations that hold up across some space of ways that the object might be changed. Scientists interested in different changes seemingly would require generalizations that differ in the space of alternatives that they span.

In what follows, I argue that the theories of organic chemistry and quantum chemistry are a case of this sort. That is, these theoretical approaches may have an object (or objects) in common—organic molecules—but they have different explanatory agendas with respect to these objects. Consequently, the theoretical generalizations to which they appeal in explanations, predictions, and applications are distinct. Organic molecules are, as it were, 'carved-up' differently in the two approaches, and this is reflected in the fact that the space of alternatives into which they are situated differs. This does not imply that the two approaches are inconsistent, or otherwise unable to contribute to one another (which they clearly do), but it does suggest that there is something misguided about the thought that refinements in one of these theories would allow for total replacement of the other.

2. Understanding Theoretical Organic Chemistry. One way to get at the explanatory agenda of theoretical organic chemistry is to look at the sorts of why-questions that organic chemists are interested in answering. Some attempts to characterize explanation in organic chemistry along these lines have been made (see Goodwin 2003). According to this analysis, organic chemists are in large part interested in providing structural accounts of product distributions, relative rates, and some other related features of organic reactions (such as mechanistic pathways). These explanations are for the most part contrastive, because they take place against a backdrop of alternatives (the contrast class). They might attempt to explain why a reaction is stereospecific rather than leading to a mix of stereoisomers, or why the dominant product is *X* rather than *Y* or *Z*, or why one reactant undergoes a transformation much more quickly than another. Structural differences between the alternatives in the contrast class are particularly salient for organic chemists. Satisfying explanations must identify the structural features of the molecules being considered that are responsible for the contrastive fact being explained.

If this is a reasonable account of explanation in organic chemistry, then the space of alternatives into which particular phenomena are situated in this approach is a space of alternative structures. To explain why some contrastive fact is true about a particular molecule or reaction, one would have to identify what and how a structural feature might be altered so that contrastive fact would no longer hold true. If one explains, for example, that a particular reaction proceeds quickly relative to some standard because of an electron-releasing neighboring group, this entails that should the substrate be modified so that it no longer had this neighboring group, it would not pro-

https://doi.org/10.1086/673734 Published online by Cambridge University Press

ceed as quickly relative to that standard. In order to support such a claim, one has to appeal to generalizations that hold up across some range of changes in structure. For instance, in this example, one might appeal to a claim like "electron-releasing neighboring groups stabilize carbocations." This generalization holds up across changes in both the structure of the neighboring group and the structure of the cation, but it is not an exceptionless law of nature. Still, it allows one to conclude that the reaction would not proceed as quickly if the neighboring group were replaced by a nonelectron-releasing substituent, and this is the kind of thing that a scientist must know, according to Woodward, if she is to explain, rather than just describe, the phenomena she investigates. In virtue of knowing this, the organic chemist can make a range of predictions about how a class of alternative substrates (those differing from the original substrate in their neighboring groups but otherwise similar) would behave. This predictive ability underwrites the chemist's ability to manipulate and control organic reactions. It supports a rule of thumb of the following form: if you want to speed up (or slow down) a reaction that proceeds through a carbocation intermediate, then consider modifying the substrate by either adding (or removing) an electron-releasing neighboring group.

Thus far, I hope to have made it plausible that many of the explanations and predictions of organic chemists depend on generalizations about what structural features might be swapped out of an organic molecule in order to change its behavior in specific ways. I now want to try to make it plausible that these are the right sort of changes or alternatives for organic chemists to keep track of, given some of their broader objectives. One of the major ways that an achievement in organic chemistry can come to be seen as important is by being, or contributing to, novel syntheses. Accordingly, we should expect that the theory of organic chemistry is significant, at least in part, because it supports the design of total syntheses of novel compounds. Other work has been more explicit about how the theory of organic chemistry does this (Goodwin 2008, 2009), but here it will suffice to provide a brief sketch.

One can think of synthetic design for a novel compound as the process of progressively pruning the 'retrosynthetic tree' of the target molecule.[1] The retrosynthetic tree of a compound is a branching collection of potential synthetic pathways that is determined by the known synthetic transformations available to the designer. Such a tree can be generating by working backward from the target molecule in stages. In the first stage, all the ways that one might generate the target molecule as the product of a known re-

1. This way of talking about synthetic design is adapted from E. J. Corey (see Corey and Cheng 1989), but it is not the only way, or necessarily the best way, to think about synthetic design.

https://doi.org/10.1086/673734 Published online by Cambridge University Press

action are considered, and the structures all of the reactants that would be needed to run these reactions are described. These structures are the nodes of the first stage. In subsequent stages, all the ways that one might generate the nodes of the previous stage are considered, and a new array of nodes corresponding to the reactants for these reactions is described. The process continues until potential pathways end by generating nodes that are already synthetically accessible. The generation of this tree of potential synthetic routes requires using the taxonomic resources of structural organic chemistry. To generate the tree, one must be able to determine, based on the structure of an organic molecule, what (single-step) reactions might be used to make it. This requires not only the recognition of functional groups (or perhaps narrower structural patterns) in the molecule but also a structurally characterized set of potential reactions (e.g., alcohols can be made by the addition of water to an alkene). In this sense, the structural characterization of reactions and the functional group classification of molecules are conceptual resources required to even set up the problem of synthetic design. Once the problem is set up, the full theory of organic chemistry can be used to help make intelligent decisions about which of these routes are worth pursuing.

The selective exploration and evaluation of the retrosynthetic tree can be thought of, again following Corey and Cheng (1989), as occurring in three stages: strategic pruning, plausibility assessment, and optimization (see Goodwin 2009). Each of these stages requires assessing what would happen should a previously known synthetic reaction be applied to a novel molecule. To do this the organic chemist requires information about both how that synthetic reaction occurs (its mechanism) and how the particular environment of the novel molecule will influence or modify the mechanism and/or reaction. Synthetic reactions are discovered, and their mechanisms investigated, in chemical structures that are typically substantially different from the novel molecules for which synthetic chemists must make these assessments. Consequently, synthetic design requires evaluating how the characteristics of the known reaction would change should the structure of the substrate be modified in particular ways. In other words, these chemists must evaluate the effects of swapping out parts of the standard structure (the structure on which the synthetic reaction was initially investigated) for structural features relevant to the novel molecule being investigated. This requires appealing to generalizations that are robust across changes in structure: claims that span the space of possibilities characterized in the explanations and predictions described earlier. This is, then, one reason why the mastery of the theory of organic chemistry demonstrated in the explanations characteristic of the discipline is useful in synthetic design and why the explanatory agenda of theoretical organic chemistry is the space of alternative structures. Applying these generalizations to the reactive possibilities of novel molecules is typically more complicated than it is in standard cases of explanation or pre-

https://doi.org/10.1086/673734 Published online by Cambridge University Press

diction. Complex novel molecules often differ from reactive standard reactions in multiple overlapping, and possibly conflicting, ways, with some features that promote a reaction and some that inhibit it. Synthetically relevant novel molecules are therefore often outside the empirically investigated scope of the generalizations being used in these assessments. As a consequence, these assessments employ heuristics and produce predictions that are often only contrastive (this will work better than that), qualitative (one product will dominate another), or disjunctive (giving a range of possible results). The success of synthetic organic chemistry suggests that, nonetheless, these theoretically grounded heuristics are effective tools for synthetic design.

To summarize, then, at least some of the significant results of organic chemistry depend on structural organic chemistry for their very formulation and on theoretical organic chemistry for the plausibility assessments that lead to their solution. These assessments require appeal to generalizations that span a range of ways that organic structures might be modified, and this is at least part of the reason that the explanatory agenda of organic chemistry is focused on the space of alternative structures.

3. Explanation in Quantum Chemistry. Whereas the explanatory agenda of theoretical organic chemistry was revealed by characterizing the sorts of structural changes that were relevant to the field, it is easiest to characterize the explanatory agenda of quantum chemistry by looking at the sorts of generalizations or laws that figure in its explanations. According to Friesner, the core technology of quantum chemistry is "computational solution of the electronic Schrodinger equation; given the positions of a collection of atomic nuclei, and the total number of electrons in the system, calculate the electronic energy, electron density, and other properties by means of a well defined automated approximation" (2005, 6648). This core technology has undergone marked development in recent years and, as a result, increasingly large systems can be modeled to greater levels of detail and accuracy. The hope, at least according to Friesner, is that eventually "fully explanatory and predictive models" will be constructed for the "complex, condensed phase processes that govern most of the natural world" (6652). If we allow for the sake of argument that Friesner is correct, and that quantum chemistry does indeed develop along the path that he envisions, it is still worth considering what kinds of explanations (i.e., which why-questions) this technology will enable the chemist to provide and which it will not.

Most fundamentally, the "core technology" of quantum chemistry can tell you the value of an output variable (the electronic energy or electron density, for example) when provided with a nuclear arrangement and an electron count. So far this is just a description, since from Woodward's perspective one must be able to say how things would be different after cer-

https://doi.org/10.1086/673734 Published online by Cambridge University Press

tain sorts of interventions in order to have an explanation. The sorts of interventions about which this core technology can provide direct information would seem to be limited by its input variables, namely, the nuclear arrangement and electron count. If this is right, then the quantum chemist can explain, for example, the electronic energy of a system in the sense that she can describe how that energy would change should the nuclear arrangement (or electron count) be altered. The space of possibilities across which the electronic Schrodinger equation ranges is therefore the space of alternative nuclear arrangements (and electron counts, perhaps). Keeping track of how the electronic energy or electronic density changes with changes in nuclear arrangements is very useful to chemists. Many features of individual molecules and chemical transformations can be reconstructed from such changes. For example, the equilibrium geometry of a molecule can be found by searching a subspace of the space of alternative nuclear arrangements corresponding to the range of different bond lengths and angles likely available to the molecule. The lowest energy nuclear arrangement in such a subspace would then correspond to the equilibrium geometry. Similarly, reactions can be explored by mapping how the electronic energy changes as the nuclear arrangement is shifted from that of the reactants to that of the products by a variety of different sequences of changes in nuclear arrangement. It is because of applications like (but not limited to) these that quantum chemistry has been able to make increasingly important contributions to all areas of chemistry, including theoretical organic chemistry. Still, though this developing technology allows chemists to explain many things by keeping track of changes resulting from manipulations of the nuclear arrangement, there is no reason to think that every why-question, or application, that chemists are interested in can be addressed by keeping track of such changes.

Most important, from the perspective of this essay, is that the space of alternative nuclear arrangements is distinct from the space of alternative structures. The array of distinct structures with which the organic chemist is professionally concerned differs from the target molecule not only in the arrangement of the nuclei but also in the number and type of those nuclei. Consider, for instance, an organic chemist who explains why a solvolysis reaction run on a substrate (say X) with an electron-releasing group (say R) adjacent to the substituted carbon proceeds more quickly than the standard reaction without the neighboring group. The answer would likely appeal to the fact that reactions of this type proceed through a carbocation intermediate, and therefore their transition states can be stabilized by a neighboring, electron-releasing group. This explanation has implications that go well beyond the particular structures (both X and R) that are actually mentioned in the original question. For example, by putting forward this explanation, the organic chemist implicitly suggests that it does not really matter which electron-releasing group is adjacent to the substituted carbon—any of them

https://doi.org/10.1086/673734 Published online by Cambridge University Press

would be expected to have qualitatively similar effects. The structural difference that makes a difference, the organic chemist is saying, is that the adjacent group has a certain character (it is electron releasing), and that character can be realized in a multiplicity of ways. Furthermore, an explanation like this would typically imply a lot about the structural features that do not make a difference. For example, the particular structures considered may be acting as representatives of hydrocarbons with a good leaving group attached to a tertiary carbon (which would be the type of structure that undergoes solvolysis through a carbocation intermediate). If this were the case, then this particular explanation would be expected to be portable (ceteris paribus) to any other relevantly similar solvolysis reaction run on a structure in this broad class. It is the broad implications of an explanation like this that make establishing such an explanation useful to the organic chemist, for ultimately he or she will want to make reasonable assessments of how novel molecules will behave by invoking the principles behind such explanations.

By contrast, though the quantum chemist might be able to provide a detailed potential energy surface for the reaction in question (and even for the standard reaction to which it is being compared) from which the rate of the reaction can be derived (and shown to be higher than the standard), this does not by itself tell the chemist what structural features might be changed in order to modify the outcome. The potential energy surface conveys how electronic energy changes with the nuclear arrangement, but it tells you nothing about what would happen to the rate if whole sections of molecule *X* were replaced by some other substituent (that would amount to shifting to an entirely different potential energy surface). To do that, you need generalizations that hold up across changes in structure, which is what organic chemists are in the business of providing. Without such generalizations, each structure is an island, which quantum chemistry may describe with mathematical precision but which contains no clues about what to expect from other structures. As Jensen puts it: "Although sophisticated electronic structure methods may be able to accurately predict a molecular structure or the outcome of a chemical reaction, the results are often hard to rationalize. It therefore becomes difficult to apply the findings to other similar systems" (1999, 347). Of course any particular question about how an alternative structure would behave could be answered by producing a new potential energy surface, but this is just more description without systematic attempts to uncover how rates and structures (or distinct potential energy surfaces) are related. These attempts might conceivably take a form different than the structural approach of modern organic chemistry (such as frontier molecular orbital theory or Marcus theory; see Jensen 1999, chap. 15), but they would seem to require at least some form of theorizing over and above the core technology of quantum chemistry.

https://doi.org/10.1086/673734 Published online by Cambridge University Press

To reinforce this point, it is useful to try to imagine how a quantum chemist with a really fast computer but with no knowledge of organic chemistry might try to solve the problem of coming up with a plausible synthetic design for a complex, novel compound. Our chemist, let us imagine, is given a chemical structure, which amounts to a space of possible three-dimensional arrangements of a set of particular nuclei and some fixed number of electrons. It is no problem to come up with an equilibrium geometry and even, perhaps, a set of possible molecular rearrangements corresponding to other minima on the potential energy surface. Solvent molecules might be added to develop a sense of the geometry in condensed phases. Before the process of synthetic design could even begin, however, the chemist would have to identify some set of chemical reactions that might be capable of producing the target molecule. By stipulation, this is a novel compound, so there is no history of past reactions run on the target molecule. Instead, somehow, a bunch of reactions run on different molecules are going to have to supply the array of possibilities for making the target molecule. It is hard to avoid the conclusion that this will require some notion of structural similarity facilitating the application of past chemical experience to this new synthetic problem. Organic chemists have developed the concept of the functional group (and related structural classifications of reactions) to play this basic role, and it is not clear how our computer chemist could avoid something like it. Without some guidance from known organic chemistry, the quantum chemist would be at a loss as to which expanded potential energy surfaces in which to embed the target structure in order to evaluate potential reactions leading to its production.

Perhaps it is just too much to ask the computer chemist to come up with the retrosynthetic tree on her own, and instead, we could grant her the retrosynthetic tree and just ask her to do the pruning. But here, again, the computer chemist faces the difficulty that pruning the synthetic tree requires making assessments of how reactions run in very different circumstances will behave when applied in a new structural environment. Organic chemists approach these questions by characterizing the mechanisms of reactions and the effects of structural substitutions on those reactions, and this is frequently enough to make plausible assessments. Our quantum chemist, on the other hand, must approach the plausibility of each potential reaction in the retrosynthetic tree as a new, distinct question. After all, the particular relevant potential energy surface has never been mapped before. Quantum chemistry will not tell you which features of previously explored potential energy surfaces are likely to recur in the current case, at least not by using its core technology. (It is interesting to note that computer programs designed to capture the logic of chemical synthesis are not built around the core technology of solving the electronic Schrodinger equation; instead, they consist

https://doi.org/10.1086/673734 Published online by Cambridge University Press

of heuristics for structural analysis; see Corey and Cheng 1989.) Furthermore, there would be no obvious guidance as to how much detail would be required in order to make realistic assessments of the plausibility of a reaction. For instance, solvents play a crucial role in many organic reactions, and so some reactions that are not synthetically useful in one solvent may be useful in another. Would each potential reaction therefore have to be evaluated on a range of potential energy surfaces reflecting the different solvents that might influence the synthetic plausibility of the reaction? At best, if organic chemists gave the quantum chemists the array of distinct potential energy surfaces worth investigating, the quantum chemists might, by brute force, be able to plow through these almost endless options and make some assessments of routes that might be synthetically plausible. To approach synthetic design in this way would, however, be the worst sort of purely empirical random groping. While she might be able to describe the reactions that have accessible activation energies, she would have done nothing to embed these facts into the "systematic patterns of counterfactual dependence" that are the hallmark of a developed scientific understanding of a range of phenomena.

4. Conclusion. In this essay, I argue that the theories of organic chemistry and quantum chemistry are theoretical approaches that have objects in common but that have different explanatory agendas with respect to those objects. Whereas the theoretical generalizations of interest to organic chemists span the space of alternative structures, the core technology of quantum chemistry generates explanations spanning the space of alternative nuclear arrangements. These theories interact in useful ways, and they should generally be thought of as complements to one another rather than as competitors. I tried to show why it would be misguided to expect that eventually quantum chemistry would become the mathematical theory for the whole of chemistry. Synthetic design, which is where the theory of organic chemistry really comes into its own, requires generalizations that span the space of alternative structures. While the quantum chemist can be extremely useful in answering particular questions that come up in synthetic design, she is poorly equipped to generate and evaluate potential synthetic pathways, except in a brute and nonexplanatory way.

This result is philosophically interesting for at least two reasons. First, the failure of quantum chemistry to support synthetic design is a clear example of how distinct sciences (or perhaps distinct levels of the same science) may 'decompose' the same objects or phenomena in different ways and thereby support generalizations that are stable across different sorts of interventions. When the space of possibilities that are relevant to the significant problems of one science do not line up with those possibilities tracked in another science, then it is hard to see how the sciences could result in the

https://doi.org/10.1086/673734 Published online by Cambridge University Press

same capacities to manipulate and control the world. Second, and more generally, this example demonstrates the fruits of considering the issue of reduction not just in light of the explanations and predictions generated by a scientific theory but also in the context of the applications—the potential for manipulation and control—facilitated by the theory. While the predictions and explanations generated by a theory typically play a central role in its justification, it is often the applications of a theory that are central to its significance. By focusing on what makes the truths generated by a field of inquiry valuable, and thus (in the case of organic chemistry) on how the field confronts novel cases of application, new light is shed on the autonomy of the sciences.

REFERENCES

Corey, E. J., and X. M. Cheng. 1989. *The Logic of Chemical Synthesis*. New York: Wiley.

Friesner, R. A. 2005. "Ab Initio Quantum Chemistry: Methodology and Applications." *Proceedings of the National Academy of Sciences* 102 (19): 6648–53.

Garfinkel, Alan. 1991. "Reductionism." In *The Philosophy of Science*, ed. Richard Boyd, Philip Gasper, and J. D. Trout. Cambridge, MA: MIT Press.

Goodwin, William. 2003. "Explanation in Organic Chemistry." *Annals of the New York Academy of Sciences* 988:141–53.

———. 2008. "Implementation and Innovation in Total Synthesis." *Foundations of Chemistry* 10:177–86.

———. 2009. "Scientific Understanding and Synthetic Design." *British Journal for the Philosophy of Science* 60:271–301.

Jensen, Frank. 1999. *Introduction to Computational Chemistry*. Chichester: Wiley.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

https://doi.org/10.1086/673734 Published online by Cambridge University Press