CATEGORIES FOR THE WORKING PHILOSOPHER

ELAINE LANDRY

Reviewed by Chris Kapulkin and Nicholas Teh

Categories for the Working Philosopher
Elaine Landry (*ed.*)
Oxford: Oxford University Press, 2017, £75.00
ISBN 9780198748991

The publication of Saunders Mac Lane's ([1971]) *Categories for the Working Mathematician* (CWM) was a signal event in the history of category theory. This influential textbook grew out of the recognition that there had emerged a well-established body of material that one might consider 'basic' category theory, and that it provided an architecture of concepts (such as categories, functors, limits, and adjoints) that unified many areas of mathematics. One of the explicit aims of CWM was thus to transmit this meta-mathematical *lingua franca* to mathematicians in other subfields.

In the following years, the unificatory power, generality, and foundational import of category theory have made it fertile soil not only for mathematical but also philosophical thought, and so it is a pleasure to see some of this work expertly curated by Elaine Landry in her homage to Mac Lane: *Categories for the Working Philosopher* (CWP). Like CWM, CWP seeks to convey the fundamental concepts of category theory to a broad audience, especially philosophers of *X* (where *X* is a subject to which category theory has made a significant contribution or has the potential to do so). But there are some important differences in the scope and context of these respective works.

First, while CWM organizes and transmits a collection of categorical concepts that might uncontroversially be regarded as 'basic', category theory has since experienced an explosive development in various sub-areas, for example, ∞-categories in logic and algebraic topology, 'categorification' in representation theory and knot theory, and the use of topos theory and 'derived' methods in algebraic geometry. Thus, beyond a minimal core of background knowledge (for example, the contents of Borceux [1994a], [1994b], [1994c]), what counts as basic knowledge depends heavily on the sub-area of category theory with which one is concerned. Since Landry's volume is rightly intended to give the reader a flavour of new concepts and techniques that have arisen within these sub-areas, its scope is significantly more ambitious than that of CWM and the material is correspondingly less unified. A second difference is that while CWM is concerned with the relevance of category theory for pure mathematics, one finds in CWP a roughly equal division between pure (Chapters 1–10) and applied (Chapters 11–19) material, where the latter includes applications to physics, biology, and the philosophy of science.

The breath-taking scope of the material covered in CWP makes it particularly challenging to identify some core set of themes that confer a narrative unity upon the various chapters in this volume (beyond the true but unhelpful statement that the chapters all concern category theory!). Nonetheless, we believe that one can discern two such themes, which provide the reader with a useful (although by no means the only) path through this material.

Theme A—which is dominant in the pure chapters—is that the 'principle of extensionality' (PE) is the key to understanding why category theory is conceptually or foundationally illuminating, and thus drives many developments in modern category theory. We note that (as will be evident below) we adopt the sense of 'extensionality' most commonly used by category theorists, and not the—antithetical—sense that is adopted by set theorists. Roughly speaking, PE stipulates that one should not 'look inside' a mathematical object (for example, by considering the elements of a set) to determine whether two objects or constructions are identical. One way in which PE can be formalized in category theory is via the Yoneda lemma (and its higher-categorical generalizations), which asserts that two objects of a category are isomorphic (or suitably n-equivalent) whenever they are related in the same way to all objects of that category. By contrast, the use of set theory as a candidate foundation of mathematics violates PE, because of the so-called axiom of extensionality—a sense of 'extensionality' that differs from our own and that of the category theory community!—in ZFC (according to this axiom, two sets are equal if and only if they have the same elements; hence one is required to 'look inside' the sets to examine whether or not they are equal). An important immediate consequence of PE is the invariance of all categorical constructions under isomorphism in a category (and the higher-dimensional analogues of this statement). In the literature, this statement is sometimes called the principle of equivalence (nlab [2018]); below, we shall treat it as the restriction of PE to the special case of 'groupoids', namely, categories in which all arrows between objects are invertible.

Theme A can be seen as motivating two projects: A1, the use of category theory to rearticulate some branch of logic or mathematics so that one can implement PE in specific contexts; and A2, the far more ambitious programme of using category theory to build PE into the very foundations of logic and mathematics, thus making PE an 'intrinsic' feature of one's linguistic framework, as it were. On the other hand, Theme B—which is dominant in the applied chapters—is the project of abstractly conceptualizing a scientific theory (or certain aspects of it) so that the resulting description is amenable to the methods of category theory. We shall have more to say about the extent to which these two themes dovetail at the end of our review.

Based on the above hermeneutic, one way of reading the pure chapters of CWP is as motivating a progression from A1 to A2. An introduction to the logical application of A1 is given in Chapter 7, where John Bell summarizes the advances in categorical logic from its conception in the 1960s up to roughly 1990. Among the main topics addressed are Lawvere's functorial semantics and the topos-theoretic description of logic. These early developments and results relating different kinds of logics to different kinds of categories (called 'hyperdoctrines') set the stage for the remaining nine pure chapters. In Chapter 9, Kohei Kishida extends this approach to logic by using category theory to describe modal logic and its semantics. We recall that in Lawvere's view, the syntax of a logic is a category and its models are given by functors from this category to some target category. By contrast, in Kishida's re-articulation of modal logic, the syntax is a functor and the models are certain natural transformations (namely, arrows between functors). While the development of mathematical logic in the twentieth century has generally placed more emphasis on syntax than semantics, we note that category theory also has something to contribute to syntax-heavy systems; for instance, Chapter 10 by J. R. B. Cockett and R. A. G. Seely reviews how symmetric monoidal categories can be used to develop a categorical semantics for linear logic.

It is relatively uncontroversial that category theory is 'foundational' in the sense that it has provided many areas of mathematics with a convenient linguistic framework in which to reason. However, and far more controversially, category theory has also been put forward as a foundation for all of mathematics, and in particular one that does not rely on set theory. This claim has been the subject of ongoing debate: could and should a categorical foundation replace the more familiar set-theoretic one? In Chapter 5, Michael Ernst provides an overview of this debate. While the debate has previously focused on technical adequacy and the autonomy of each of these competing foundations, Ernst chooses instead to emphasize the question of which foundations best capture the practice of working mathematicians. This chapter, along with Chapters 1 and 6, can be viewed as a bridge between A1 and A2, because it reflects on the significance of the fact that PE encodes an essential insight from mathematical practice and is responsible for the fruitfulness of using category theory to re-articulate known mathematical structures. For instance, in Chapter 1, Colin McLarty surveys the various scenarios in which a working mathematician might encounter set theory and concludes that what ultimately matters to practitioners is our aforementioned PE. His main example here is the fact that there are many equivalent—but not equal—constructions of a tangent bundle and the fact that practitioners only care about these up to the structural properties that determine their equivalence. Along similar lines, Jean-Pierre Marquis argues in Chapter 6 that the notion of 'canonical maps' illustrates how category theory provides insight into the structural character of mathematics that is not provided by set theory.

While category theory can yield such insights, there are still areas of mathematics whose structural description requires us to enhance our meta-mathematical framework beyond categories. An example of such an area is… (1)-category theory itself! To illustrate this point, consider that from the set-theoretic perspective, the standard of sameness of two categories is equality; similarly, one might be tempted to say that from the category-theoretic perspective, the relevant standard is isomorphism. However, we know from the practice of category theory that we miss out on many interesting phenomena if our standard of sameness is as strict as isomorphism—in fact, it is often fruitful to relax this standard to a weaker one called 'categorical equivalence'. How can we build into our meta-mathematical description a principle that will yield the 'right' standard of sameness for categories (or higher categories)?

One attempt to do so is Michael Makkai's FOLDS (first-order logic with dependent sorts), which is presented by Marquis in Chapter 8. For instance, in the FOLDS theory of categories, two categories are equal if and only if they are equivalent. As one might expect, while FOL has a natural interpretation in sufficiently structured categories, the general interpretation of FOLDS requires the resources of higher categories. On the other hand, one can take a different tack from Makkai and use the existing system of Martin–Löf type theory as the basis for a system that captures the relevant notion of sameness. This approach—now known as homotopy type theory (HoTT)—was taken by Vladimir Voevodsky, who noticed a close connection between the notion of 'equality' in independent type theory and the notion of 'homotopy' that is familiar from topology.

The origins of HoTT are explained by Michael Shulman's Chapter 3, after which Steve Awodey's Chapter 4 presents the univalence axiom of HoTT as the ultimate form of structuralism. To state the univalence axiom, consider that for any two objects, *A* and *B*, in a (possibly higher) category, one may consider a map taking 'a proof that *A* and *B* are equal' to 'an equivalence between A and B'. The univalence axiom then asserts that this map itself is an equivalence, that is, in suitable sense, every equivalence arises from a proof of equality. While it may at first seem that this axiom restricts the class of equivalences in a category, in fact it does the opposite: it extends the notion of equality to match that of equivalence. This formally forbids the user of a foundational system from making any statements that would violate PE, thus turning a philosophical principle into a foundational axiom. Upon closer examination, the univalence axiom bears close resemblance to the object classifier of an ∞-topos, as defined by Jacob Lurie ([2009]). It is believed (although not yet proven) that HoTT is the internal language of such ∞-topoi; this more elaborate set of connections figures in David Corfield's Chapter 2, which describes Urs Schreiber's novel approach to geometry, namely, doing geometry 'internally' to an ∞-topos.

To sum up, one way of reading the 'pure' chapters of CWP is as building up to an explication of how in recent mathematics PE has been implemented in ever more thorough-going ways. We now turn to the 'applied' chapters of CWP.

In Chapter 16, David Spivak makes a compelling case for the fruitfulness of category theory as a model of the models that we use in the applied sciences: this case essentially turns on using PE (in the guise of Yoneda) to articulate the relationships between such models. Furthermore, he rightly stresses a point that highlights our Theme B: a large part of the task of applying category theory to some discipline consists in understanding which aspects of a theory one should conceptualize in categorical terms, and the level of abstraction one needs to work at in order for this choice to be mathematically and scientifically fruitful. For instance, Joachim Lambek's (posthumous) Chapter 14 provides an intriguing—if idiosyncratic—illustration of this task by describing a small fragment of field theory (Dirac spinors on Minkowski spacetime) as an additive category.

For a rather more ambitious and comprehensive attempt to conceptualize aspects of physics in terms of category theory, the reader need look no further than Chapters 11 and 12, which review the work of the 'Oxford school', who apply category theory to quantum mechanics. In Chapter 11, Samson Abramsky reviews how the probabilistic data of quantum theory can be conceptualized in terms of a presheaf that assigns such data to various sets of compatible measurements. Among other things, this powerful abstraction of (part of) the structure of quantum theory allows one to provide a classification of the phenomenon of 'contextuality' (that is, the 'inconsistent' aggregation of data from the perspective of classical probability), to apply this classificatory scheme to non-quantum theories, and (in principle) to apply the full category-theoretic machinery that has been developed for sheaves (and their higher analogues) to such an analysis. In Chapter 12, Bob Coecke and Aleks Kissinger provide the first of a three-part overview of the programme called 'categorical quantum mechanics'.[1] Here we see that many of the key concepts of quantum information theory (for example, compositionality, causality, no-cloning, teleportation, non-locality) can be abstracted and formalized within the setting of symmetric monoidal categories with further structure; thus, these concepts can also be applied to non-quantum theories that share a similar categorical structure. While much of applied category theory has focused on the case of mathematical physics, Andrée Ehresmann's Chapter 15 shows that category theory can also be fruitfully applied within specific frameworks for modelling living systems and for modelling cognitive systems in biology.

Chapters 13 and 17 of CWP are directed at the broad topic of how category theory can be used to formalize relationships between various scientific theories. In order to provide a toy model of how category theory can be used to compare theories, one can simply treat the theory's models as structured sets collected into categories of various kinds. In Chapter 13, James Weatherall assumes this set-up and discusses how statements about differences between various fragments of physical theories can be re-articulated in the language of 'forgetful functors' and in particular Baez's taxonomy of 'structure', 'property', and 'stuff'-forgetting functors. We note that this topic is in fact deeply related to (the equivalence principle case of) PE: Baez's taxonomy was originally intended to apply to the homotopy theory of *n*-groupoids (Baez and Shulman [unpublished]), and the relationship between this application and the notion of theories/models (as well as 'gauge symmetry') has been discussed in both the philosophy (Dougherty [2017]; Nguyen *et al*. [forthcoming]) and the physics (Benini *et al*. [2015]; Schreiber and Shulman [unpublished]) literature. In Chapter 17, Hans Halvorson and Dimitris Tsementzis describe how 'syntactic' and 'semantic' categories can be associated with certain logical theories and proceed to use topos-theoretic techniques to discuss the relationships (in particular, equivalence) between such categories. They then consider a two-category of logical theories and discuss the sense in which this might help us understand (analogically) various relationships between actual scientific theories.

Finally, Landry's Chapter 18 discusses the uses and abuses of category theory within the metaphysics of science, especially with respect to the position known as 'radical ontic structural realism' (ROSR). She argues (persuasively, in our view) that any attempt to use only the categorically described mathematical structure of a physical theory to argue in favour of ROSR will founder, because in order to succeed, such an argument requires an appeal to the object-level physical structure of phenomena. However, Landry also argues that the conceptual resources of category theory (in particular PE) vindicate a certain sort of Hilbertian mathematical structuralism.

Professor Landry is to be congratulated on putting together a stimulating volume that introduces a broad audience to so many of the key conceptual, foundational, and philosophical ideas driving contemporary work at the intersection of philosophy and category theory. In closing, we make two small observations that may be helpful to the reader. First, despite the importance in contemporary category theory of ∞-categories (which one might think of as 'going all the way with PE'!), discussion of this concept seems to be largely absent from CWP (apart from the chapters of Marquis, Shulman, and Corfield, which gesture at ∞-categories). Second, there seems to be something of a disjunction between the pure and applied parts of CWP: Recall that the narrative of the pure part seems to yield a clear, overarching moral: an ever-deeper implementation of PE is required to comprehend mathematical practice, thus culminating in our Theme A2. However, it is less than clear that the applied discussions cohere enough to provide a unified moral of this kind. Could there, for instance, be a physical or philosophical or (in the applied realm) methodological principle to motivate an analogous moral, or that yields a more trenchant analysis of Theme B (which seems very much to be a theme at the level of *technê*)? And can one say something systematic and principled about the relationship between Theme B and Theme A? Discussion of these points would further clarify the importance of category theory for the philosophy of science and for the philosophy of the specialized sciences.

Chris Kapulkin
Department of Mathematics
University of Western Ontario
kkapulki@uwo.ca

Nicholas Teh
Department of Philosophy
University of Notre Dame
nteh@nd.edu

References

Baez, J. C. and Shulman, M. [unpublished]: 'Lectures on n-Categories and Cohomology'.

Benini, M., Schenkel, A. and Szabo, R. J. [2015]: 'Homotopy Colimits and Global Observables in Abelian Gauge Theory', *Letters in Mathematical Physics*, 105, pp. 1193–222.

Borceux, F. [1994a]: *Handbook of Categorical Algebra, 1: Basic Category Theory*, Cambridge: Cambridge University Press.

Borceux, F. [1994b]: *Handbook of Categorical Algebra, 2: Categories and Structures,* Cambridge: Cambridge University Press.

Borceux, F. [1994c]: *Handbook of Categorical Algebra, 3: Categories of Sheaves*, Cambridge: Cambridge University Press.

Dougherty, J. [2017]: 'Sameness and Separability in Gauge Theories', *Philosophy of Science*, 84, pp. 1189–201.

Lurie, J. [2009]: *Higher Topos Theory*, Princeton, NJ: Princeton University Press.

MacLane, S. [1971]: *Categories for the Working Mathematician*, New York: Springer.

Nguyen, J., Teh, N. J. and Wells, L. [forthcoming]: 'Why Surplus Structure Is Not Superfluous', *British Journal for the Philosophy of Science*.

nLab [2018]: 'Principle of Equivalence'.

Schreiber, U. and Shulman, M. [unpublished]: 'Quantum Gauge Field Theory in Cohesive Homotopy Type Theory'.

Notes

[1] This first part treats purely quantum processes, whereas the second part includes classical data and its interaction with quantum data, and the third part treats the concepts of observables and complementarity by means of 'internal' Frobenius and Hopf algebras.