Andrew Bacon, *A Philosophical Introduction to Higher-Order Logics*. New York, NY: Routledge, 2024. Pp. xviii + 464.

Bacon's *Introduction* is part textbook, part research monograph. Both parts are excellent. I have no doubt that, in philosophical circles, *Introduction* will become the standard textbook on higher-order logic. It will prove invaluable to any philosopher who wants to learn about higher-order logic for the first time, as well as to experts in the field who are looking for a comprehensive reference book. I also strongly recommend *Introduction* to anyone with research interests in higher-order metaphysics, especially to those who are interested in higher-order theories of propositions or higher-order theories of modality.

Like all good logicians, Bacon counts from zero, and Chapter 0 provides some philosophical motivations for using higher-order resources. The bulk of the book is then split into five parts. Part I (Chapters 1–3) covers the basics of typed languages and λ-abstraction. In Part II (Chapters 4–8), Bacon introduces classical higher-order logic, presents some appealing extensions of that logic, and explores their consequences for modal metaphysics. In Part III (Chapters 9–13), Bacon considers some alternative higher-order logics, and uses them to develop theories of structured propositions. Part IV (Chapters 14–17) is dedicated to the model theory of higher-order logics. Importantly, Bacon's model theory is carried out in first-order set theory, and so no model of a higher-order theory is ever 'intended', in the sense of capturing what that theory means. Rather, the model theory is of purely instrumental value: by establishing soundness and/or completeness results, we can use the model theory to show that some premises do or do not imply some conclusion in a given higher-order logic. Part V consists of two technical appendices (Appendices A and B).

That is obviously a very brief summary of *Introduction*. However, it would be impossible to give a detailed survey of the entire book in a short review. That is partly because Bacon covers an impressive range of topics, but also because much of the material is technical and challenging. Indeed, part of what makes *Introduction* such a good textbook is that Bacon takes his time to explain the material clearly and accessibly. Instead, I will use the rest of this review to discuss one theme that emerges from the book.

In Chapter 5, Bacon presents system H, which is his best candidate for classical higher-order logic. The language of H is typed (see Chapter 1). There are two basic types: names are type e, and sentences are type t. If σ and τ are types, then $\sigma \to \tau$ is the type of term that yields a term of type τ when it is applied to a term of type σ. (We apply M to N simply by writing MN.) The language of H also permits us to form complex terms via λ-abstraction (see Chapter 3): if x is a variable of type σ, and M is a term of type τ, then $(\lambda x.M)$ is a term of type $\sigma \to \tau$. The resulting λ-terms are governed by a number of principles, including:

β: Every biconditional, $M \leftrightarrow N$, where N can be obtained from M by replacing some occurrence of $(\lambda x.P)Q$ with $P[Q/x]$, i.e. with the result of substituting Q for every occurrence of x in P (provided no variable free in Q becomes bound in the process of this substitution)

Finally, H includes a pair of quantifiers for each type: $\forall_\sigma$ and $\exists_\sigma$. Intuitively, these quantifiers allow you to quantify over entities of type σ, i.e. the values of type σ terms.[1] These quantifiers behave as you would expect: $\forall_\sigma M$ implies MN whenever M is type $\sigma \longrightarrow t$ and N is type σ; and the reverse entailment also holds if N is arbitrary.

H has some striking metaphysical consequences. To give just one example, it implies that every relation has a converse, in a very strong sense:[2]

Converses: $\forall_{\sigma\rightarrow(\sigma\rightarrow t)}X\exists_{\sigma\rightarrow(\sigma\rightarrow t)}Y\forall_\sigma wz(Xwz =_t Yzw)$

(where $=_t$ is a propositional identity sign that takes formulas as arguments).[3] Let X be any relation of type $\sigma \longrightarrow (\sigma \longrightarrow t)$. By the β-principle, $(\lambda nm.Xmn)wz = Xzw$. *Converses* follows by the classical quantifier rules built into H.

Once you realise that H has metaphysical consequences like this, you might choose to lean into them, or you might choose to do what you must to avoid them. In Part II, Bacon leans into them. He presents a coarse-grained theory of properties and propositions, which he calls C for *Classicism*, that identifies properties/propositions which are provably materially equivalent in H (see Chapter 6). He then develops some important consequences of C for modal metaphysics: we can use C to define a variety of key modal notions (e.g. *necessity*, *entailment*, *possible world*), and prove some important results about modality, most notably that there is a *broadest* necessity (see Chapter 7). Bacon describes this as a kind of *modal logicism*, and it is a clear demonstration of the power and utility of higher-order logic.

In Part III, Bacon tries the other choice. He explores what it would take to avoid results like *Converses*. The two obvious options would be to restrict the β-principle or the classical quantifier rules. However, Bacon argues against these options, roughly on the grounds that that principle and those rules fix what we mean by the λ-operator and by the quantifiers. Instead, he recommends restricting the formation rules for λ-terms (see Chapter 9). If you wanted to avoid *Converses*, for example, then you could prohibit binding variables out of order: $\lambda m.Xmn$, $\lambda n.Xmn$ and $\lambda mn.Xmn$ would all be fine, but $\lambda nm.Xmn$ would not be allowed.

[1] $\forall_\sigma$ and $\exists_\sigma$ themselves are type $(\sigma \rightarrow t) \rightarrow t$; for the sake of familiarity, $\forall_\sigma\lambda x.M$ is abbreviated as $\forall_\sigma xM$ (see Chapter 4).

[2] Bacon's typing is officially monadic: every function is a function from one argument. Bacon handles polyadic functions by Currying. What we might intuitively think of as a dyadic function from entities of types σ_1 and σ_2 to an entity of type τ is treated as a monadic function of type $\sigma_1 \longrightarrow (\sigma_2 \longrightarrow \tau)$. We can make this typing easier to work with by using sensible abbreviations, e.g. abbreviating $(Xw)z$ as Xwz, and $\lambda n.(\lambda m.Xmn)$ as $\lambda nm.Xmn$.

[3] In H, we can define $=_\sigma$ as: $\lambda yz.\forall_{\sigma\rightarrow t}X(Xy \leftrightarrow Xz)$.

Bacon calls a language with restricted formation rules for λ-terms *general λ-languages*. He clearly presents, in precise formal terms, a number of general λ-languages, each of which is motivated by various metaphysical scruples (see Chapters 9 and 10). He also develops *general higher-order logics* within these general λ-languages, and maps out their consequences (see Chapters 11–13).

It is impossible to fault Bacon's execution of this project. However, for my money, the project as a whole seems less promising than the project of Part II. As Bacon demonstrates, C is theoretically very fruitful. It allows us to settle a number of important questions in modal metaphysics. Moreover, these results do not matter *just* for metaphysics: we use modal reasoning throughout our theoretical and practical reasoning (Williamson 2016), and so the implications of C reach far and wide. By contrast, it was not obvious to me that any of Bacon's general higher-order logics promise the same kind of wide-reaching impact.

More importantly, I am not convinced that the best way to avoid the metaphysical consequences of H — if that is really what you want to do — is to restrict the formation rules for λ-terms. My concern is that, for this strategy to work, each restriction on the formation of λ-terms will have to be matched by a problematic restriction on the basic terms that we may introduce.

Imagine that you want to avoid any metaphysical commitment to converse relations. Bacon recommends you do this by forbidding λ-terms to bind variables out of order. This will obviously save you from committing yourself to converses by λ-abstraction: if you start with a basic predicate, S for *shorter than*, you will not then form $\lambda xy.Syx$. But λ-abstraction is not the only way of introducing the converse of S. You might just introduce a new basic predicate, T for *taller than*, on the stipulation that $\forall xy(Txy = Syx)$.[4]

The obvious reply is that this kind of stipulation is not a permissible way of introducing basic predicates. If you start with S, then you cannot introduce T. But now you face a difficult question. How do you know whether to start with S? Why not start with T instead? According to this reply, you cannot start with both, because then you would be committed to converse relations. But on what grounds could you possibly choose between them?

To be clear, my worry here is epistemological, not metaphysical. I am not complaining that it would be metaphysically arbitrary for relation S to exist without relation T, or *vice versa*. It *would* be metaphysically arbitrary, and maybe that would be uncomfortable, but so be it. My concern is that, when setting up your general λ-language, *you* would need to choose between including S or T in your stock of basic predicates. But you could never give any reason for choosing one over the other. So the problem is not that you do not know how to speak God's language. It's that you do not even know how to speak your *own* general λ-language. And not knowing how to speak a language is importantly different from not knowing which theory within a

[4] In a general λ-language, you cannot straightforwardly take $\forall_\sigma$ to be of type $(\sigma \rightarrow t) \rightarrow t$, as you could in H (see fn. 1). Bacon offers two alternative treatments in §§9.4 & 10.3.

given language is true. When you can't pick between two theories, you can always settle for their disjunction. But if you do not know whether to speak a language featuring S or a language featuring T, you cannot speak a 'disjunction' of the two languages. At best, you could theorise *about* those two languages without expressing any preference between them, but you would still need to pick a language *in* which to do the theorising. Picking *that* language had better not require you to choose between a relation and its converse.

Fortunately, there are ways out of having to choose between S and T, even if you are leery of converse relations. Here are three possible strategies.

Strategy 1. You could decline Bacon's advice to maintain the classical quantifier rules, and adopt a *free* higher-order logic instead. In a free logic, you cannot automatically existentially generalise on a predicate; that is permitted only on the assumption that the predicate expresses a relation. So you could happily admit S and T as two basic predicates in a free logic. You would just need to insist that at most one of them expresses a relation. Of course, it may still be an unanswerable metaphysical question which, if either, of these predicates expresses a relation, but you can tolerate that, just as an agnostic can tolerate using the name 'God' without ever knowing whether it refers to anything.

Strategy 2. You could grant that S and T both exist, but then claim that $S = T$, even though $\neg\forall xy(Sxy = Txy)$. This would obviously require restricting Leibniz's Law,[5] but you could motivate this restriction by denying that predication is a transparent context. In Sxy, the predicate S does not merely express a relation; it *also* applies that relation to its arguments in a particular way. In Txy, the predicate T expresses the very same relation as S, but applies it to its arguments the other way round. (Following Williamson (1985), Bacon (p. 246) makes a similar suggestion about the natural language predicates 'loves' and 'is loved by'.) You would still technically be committed to a converse relation, but that really would just be a technicality: S would be *its own* converse.

Strategy 3. You could grant that S and T both exist, and that $S \neq T$, but claim that at most one of them is *natural*. You would still be committed to converse relations in some thin sense, but you could insist that this is not the kind of commitment that matters for metaphysics. You are metaphysically committed only to *natural* entities. So you are metaphysically committed to at most one of S or T. You may not be in a position to say which, if either, relation is natural, but that is not your problem. You can leave the world to make its metaphysical choices. What matters is that in *your* theorising, you are not forced to choose between using S and T.

I think that all three of these strategies are viable (although they each come with their own problems to solve). But, crucially, they also all eliminate the need to impose any restriction on the formation rules for λ-terms: if you adopted a free logic, you would just need to deny that S and $\lambda xy.Syx$ both exist; if you restricted Leibniz's Law,

[5] I am here thinking of Leibniz's Law as the following scheme: $\forall xy(x = y \rightarrow (\phi \leftrightarrow \phi[x/y]))$. Restricting this scheme will require either rejecting the definition of identity given in fn. 3, or revising the background logic of H.

you would just need to claim that $S = \lambda xy.Syx$; and if you introduced a privileged class of natural relations, you would just need to deny that S and $\lambda xy.Syx$ are both natural.

It seems to me, then, that there are better ways of avoiding a commitment to converses than by restricting the formation rules for λ-terms, and the same goes for the other potentially problematic commitments of H. However, I want to be clear that these critical comments are not meant to detract from the value of Part III. Bacon presents a variety of general λ-languages and logics with great formal clarity, and carefully spells out their metaphysical consequences. Whether or not you ultimately choose to use these tools, their development is an important contribution to higher-order metaphysics. When that contribution is taken together with everything else the book has to offer, it is easy to see that *Introduction* will have a lasting impact on the field.

References

Bacon, Andrew (2024). *A Philosophical Introduction to Higher-Order Logics*. New York, NY: Routledge.

Williamson, Timothy (1985). 'Converse relations'. *The Philosophical Review* 94, pp.249–262.

— (2016). 'Modal science'. *Canadian Journal of Philosophy* 46, pp.453–492.