PHILOSOPHY OF LOGICAL SYSTEMS

Jaroslav Peregrin

Table of contents

1 The transmutation of logic

1.1 The evacuation of natural language

Nobody would likely deny that logic underwent an essential rebirth during the twentieth century (some might even be inclined to say that this was the *birth* of a *real* logic[1]).

We can say, perhaps with a certain amount of oversimplification, that logic, since its very beginning, has been primarily concerned with the correctness of arguments – not in the sense that its *sole* business would be to characterize correct arguments, but that everything that is its business evolved out of its concern with their correctness. And as arguments are articulated in the medium of a language that is normally a natural language, logic was always concerned with natural languages and with the fact that some sentences in these languages follow from other sentences.

What happened in the twentieth century was that logic introduced and embraced – sometimes rather hastily – various kinds of *artificial* languages. The employment of the languages looked extremely helpful and it looked to be able to elevate logic to a brand new level of rigor and clarity. Therefore, most logicians simply vacated the old, squat shack of natural language and moved to one of the shiny new lofts of the artificial ones, almost completely forgetting about their previous dwelling-place.

Logic flourished in its new dwelling. For the first time, it could be developed with a mathematical rigor (as the artificial languages were basically mathematical structures) rather than as a collection of quasi-empirical comments on either the way human reasoning proceeds (or should proceed), or on actual human languages.[2] The liberation of logic from the yoke of psychology and its fruitful alliance with mathematics provided for that explosion of results that we know from the twentieth century.

However, the change logic underwent in this way was in no way trivial, and it is also far from trivial to determine the extent to which the "new logic" only engaged new and more powerful instruments to answer the questions posed by the "old" one, and the extent to which it replaced them with new ones. And in so far as a replacement took place, it is not easy to see whether it was because the old questions were found to be ill-conceived or obsolete, or

[1] Coffa (1991, p. 113), for example, commenting on the state of the art around 1900, when Russell tabled the idea of reducing logic to mathematics, writes that at this stage "mathematics was a reality and logic a project."

[2] Which we can find in all logic treatises since Aristotle via Port-royal logic (Arnauld & Nicole, 1662) up until nineteenth century logical textbooks (like Lotze, 1874).

whether what was going on was that the shift in question was not really noticed and logic has moved somewhere without its protagonists duly reflecting it.

An optimistic picture is that nothing significant happened, that logic only sharpened its tools and delimited its subject matter with greater precision. This view is often underpinned by the opinion that the ultimate subject matter of logic is neither linguistic structures nor mental entities, but rather some ideal entities that are only imperfectly captured/expressed/represented by our parochially human means. From this viewpoint, logic, by its nature, *is* akin to mathematics – dealing with the very kind of abstract or ideal reality that is studied by algebraists. In this way, the engagement of artificial languages and mathematical tools only brings the nature of the subject matter of logic out into the open.

The point of departure of this book consists in the insight that such optimism is unwarranted, if not naive. The move from natural languages to their artificial replacement is a significant one that must be considered with due care. And to take logic as a type of mathematics is not viable. Consider the claim of Béziau (1994, p. 73):

> Universal Logic is a general study of logic in the same way as Universal Algebra is a general study of algebra. It is based on the fact that there is no One Logic or Absolute Laws of Logic, but rather a type of logical structures who are [*sic*] fundamental mother structures in the sense of Bourbaki. Logic is then an autonomous field of mathematics, with its own intuitions and concepts and which can survive and be developed without importing specific notions from other fields of mathematics.

It is certainly plausible that logic deals with certain kinds of structures and that we can have a very general, abstract theory of such structures that reveal some very general facts holding across different fields of logic.[3] However, does it follow that logic is a part of mathematics? A great deal of work in physics nowadays consists of solving systems of certain differential equations – but does that make physics into a part of mathematics? Mathematics is certainly an indispensable *tool* of physics, but this does not make physics its *part* – and so it is with logic.

Doing mathematics, we may study all kinds of structures – it is only when it comes to applications that we find certain kinds of structures more pertinent than others. Logic, we can say, is basically a study of overt *reasoning* (and of correct arguments as its most basic instruments) – a real-world activity in which we humans engage, and can do better or worse. Thus, any study of mathematical structures done under the heading of logic should be assessed according to its usefulness in helping us do it in a better, rather than in a worse, way.

---

[3] Such generalizations started to appear long ago. Dunn and Hardegree, Font and Jansana, etc.

1.2 Artificial languages of logic as models

The problem, as I see it, is that the artificial languages that became the new dwelling-places of logic are artificial in the sense that they were built exclusively by means of our definitions. Therefore, everything we can find out about them cannot be other than a consequence of the definitions. But how then can the investigation of the consequences of our – deliberate – definitions help us decide which arguments – namely, which steps in our reasoning – are correct?

The answer may not be too complicated: studying abstract structures constituted by definitions is what is done by pure mathematics, and mathematics can certainly help us with all kinds of problems in the real world. Of course, in order for mathematics to be able to do so it must be *applied* – a structure it studies must be fashioned as a model of a phenomenon in the real world adequate enough that the study of the structure can tell us something interesting about the phenomenon.

This invokes the idea that the artificial languages of logic can be seen as models of natural language (and of the argumentative practices of which the natural languages are vehicles), similar to natural scientists' models of natural phenomena. The models are much more exactly delimited and hence much more transparent than the languages we normally use to talk and to reason; but they are not, on the other hand, languages themselves in the fully-fledged sense. Thus, they are not substitutes for the real languages in the sense that we could shift over to them and forget about the natural languages. They can be useful only insofar as they are arguably adequate to what they are substitutes for and insofar as the results of their analyses can be projected on them.

This asymmetric situation poses, it seems to me, brand new kinds of philosophical problems that are not always properly identified. The artificial languages are always *simulacra* of the natural vehicles of reasoning and hence are – more or less – parasitic upon them. This is not to say that a lot of interesting logical work cannot be done within such artificial languages, but it is to say that this must be done with an eye on their status as models. Proving that something in an artificial language is inferable from some other thing will only tell us something – beyond the fact that this is merely the result of our definitions of the language – if we can be sure that the language, as a model, is adequate to what it models, that talking about the former can be seen as a proxy for talking about the latter.

Why are artificial languages bound to remain *simulacra* and cannot ascend to the level on which they would be on a par with natural languages? (This is a question quite crucial for this book, so it deserves to be answered explicitly). The point is that perhaps they could ascend as a matter of principle, but not as a matter of fact. Natural languages' endowment with meanings is something that develops over millennia, during which time the languages become inseparably and integrally entwined with our "meaningful" practices; the artificial languages cannot make up for this in the time frame of a few years or decades. Therefore, artificial languages, in so far as they are to be treated as meaningful, are tied to natural ones by a semantic “umbilical cord.”

And the adequacy of an artificial language to a natural one, which, to be sure, is crucial here, is not a simple concept. Of course, a model is not supposed to be an identical copy of what it models, it is supposed to be in some respects simpler, more transparent, and less vague. Thus, a certain level of *dis*similarity is assumed. For example, what should we make of a situation in which an artificial language (together with some superstructure like a calculus) that acts as the model does not have a property which we would like it to have (or, the other way around, has one we would like it not to have)? If a calculus, for example, is not decidable – what does it mean? Does this fact concern the criteria of correctness of arguments in our natural language or instead only in the artificial one – or does it have to do with the way we produce the calculus as a model of the real language?

1.3 The reflective equilibrium of logical laws

At the dawn of modern logic, some of its protagonists seem to have developed logical languages in the hope that these languages would be able to replace natural language, at least in the context of mathematical, or perhaps more generally, scientific, reasoning. This, I think, has only materialized to a very limited extent. When we engage in mathematical reasoning, for instance, it often suffices for us that we know, or believe we know, that it would be possible, as a matter of principle, to articulate the reasoning in a logical calculus (though this might be quite laborious and tedious), and it is not really something that we normally do.

In the previous section, we stressed the similarity of the artificial languages of logic and the mathematical models employed by the sciences. Here we have an important *dis*similarity: there is supposed to be a feedback (though in reality perhaps only a very meager one) from the artificial languages on our actual reasoning. Hence, while natural phenomena do not care about our constructing their models, reasoning, as a social phenomenon, can be influenced by our establishment of its models, for the models can, in this case, act as *norms*.

Thus, though the picture of the artificial languages of logic as models of our actual reasoning is, as I put forward, very illuminating, we must keep in mind that, when it comes to its details, we need to be careful not to overstretch the analogy. The point is that, unlike in the case of the natural sciences, there may be a feedback loop: the model we build may influence what it models, and since, of course, this might prompt an update of the model it could, in principle, initiate a kind of self-propelling spiral.

The fact that the model may act as a norm means that it can be used to *correct* the phenomenon that it models (in our case reasoning); that is, to do away with the discrepancies between it and the phenomenon modeled by means of tampering with the latter. On the other hand, it can be taken as a (useful) model only insofar as it is adequate to the phenomenon, which is to say if there are no relevant discrepancies or if the discrepancies are negligible. Thus, we have these two different requirements which appear to pull in opposite directions (if they are not directly inconsistent): we must fashion the model to reflect the actual process of reasoning, while at the same time we are supposed to adjust the process of reasoning so as to comply with the model.

In Peregrin & Svoboda (2017), we argue that this gives rise to a "dialectical" process aiming at what has come to be called the *reflective equilibrium*[4] – an equilibrium between empirical facts concerning reasoning and our theoretical articulation of what we perceive as its rules that results from the back-and-forth movement between considering the arguments which are ("intuitively") taken for correct and the tentative explicit articulation of the corresponding rules. Thus, *logical laws* worthy of the name are born out of this very process – out of the process of a theoretical articulation of rules implicit to our reasoning practices that aim at the reflective equilibrium.

Let me add that this does not cancel the asymmetry between natural languages and the artificial languages of logic that I stressed in the previous section (though it, in a sense, alleviates it). Though there is a "two-way traffic" between the two kinds of languages, the artificial ones maintain themselves as *simulacra* and not as languages that could take over the role of the natural ones as direct vehicles of reasoning. Instead, they function as models to which we can outsource a lot of theoretical problems connected with reasoning, but can do this only to the extent they are seen as adequate models of the natural languages.

1.4 The nature of logical laws

Languages of formal logic, as we are going to present them in detail in the upcoming chapters, are first and foremost expedients of our effort to articulate the criteria of correctness of arguments and consequently of the explanation of the fact that some sentences follow from other sentences. However, we do not explain *why* it is the case that something follows from something else, we explain it only in the sense that we systematize the cases of following-from and in this sense we present every individual case as a piece in a larger mosaic.

To avoid misunderstanding, such a systematization makes room, of course, for explanations of cases of following-from in terms of other, simpler cases. We can explain that a conclusion follows from some premises in that we can decompose the step from the premises to the conclusion into a chain of simpler and more obvious steps. (This, needless to say, is what we do when we prove theorems in mathematics.) But we accordingly only reduce the more complex cases to simpler ones, we do not explain why the simpler ones hold.

[4] The term was coined by Rawls (1971) in the context of ethics, though the same considerations for logic were already presented earlier, though not under this name, by Goodman (1983). See Brun (2014; forthcoming) for an overview.

This, again, is similar to natural laws. The laws of gravity explains why a stone, released from my hand, falls to the earth. It explains it not in the sense that it would disclose some hidden essence of the stone that would make it fall down; it does so in that it subordinates it to a general law, where the law is just a generalization. In effect, it tells us that this particular stone falls to earth because it is what things always – at least insofar as we have been able to find out – do.

We cannot say *why* the law of gravity holds; we can only state *that* it holds and given this, we can use it to explain individual cases. And we consider rules of logic as articulated by the calculi in a similar way: we articulate the rules as generalizations from the individual, correct arguments. Thus, as we find out that plus minus all the arguments of the form

$$\frac{A \qquad \text{If } A \text{ then } B}{B}$$

in English are correct, we accepted

$$\frac{A \qquad A \rightarrow B}{B}$$

as a logical law. Are, then, the laws of logic just empirical generalizations?

We have already seen that logical rules differ from ordinary empirical generalizations in that they can be used as *norms* of reasoning. In fact, there are two ways in which the rules can be considered normative. First, they are normative in the sense that what they generalize are already *rules*. Their instances are *correct* arguments – they are correct not because we logicians proclaim them as correct, they are correct before any logical theory is in play, they are such because human linguistic practices are essentially rule-governed. Logic is thus supposed to pick up rules that already govern our reasoning, though in their natural form they are implicit, and it is only logical theory that makes them explicit.

Moreover, the laws of logic are normative also in the second sense, in that they can be used to correct actual argumentation. This is because what we do when we articulate them is not only to bring their implicit predecessors into the open, but rather also to "finalize," streamline, and extrapolate them to the regions where they were not really operative. This is what happens during the process of moving toward a reflective equilibrium.

Given this, we cannot really say that logic addresses a reality behind the surfaces of natural languages. It builds on some very general rules (proto-rules?) that are constitutive of our linguistic practices, but employs its huge apparatus to transform these implicit rules into something clear, unambiguous, and potentially binding. The calculi and other tools of modern logic we are going to discuss in the upcoming chapters are expedients of this enterprise, they are not a means of naming a hidden reality excavated from the depths of language.

In this sense, I think that the status of the laws of logic is similar to that of natural laws (despite all the other dissimilarities): they do not bring to light a hidden reality behind appearances, they bring a system to the appearances (to "save" them, as the Ancients would put it[5]). Of course, the system cannot be chosen deliberately, it must *fit* the appearances – but there is no reason to think that it was there, as such, all the time already before we established it.

1.5 Philosophy of logical systems

Traditional philosophy of logic targeted the "logical part" of the vocabulary of natural language, which seemed to be responsible for the correctness of a crucial part of the arguments. There were problems that logic shared with philosophy of language (such as: how do the expressions of natural language, especially the logical ones, manage to mean what they do?), and there were problems of the conceptual apparatus engaged by logic (for example, the nature of truth). And, of course, there were the philosophical questions surrounding very central problems of logic, such as the criterion of correctness of arguments and thereby also the laws of logic. Later there came discussions about the possibilities of rejecting some of the traditional laws of logic and about alternative logical systems that might ensue from their rejection.

However, after the center of gravity of logic shifted from natural language to the artificial ones, the problems multiplied. Now, for example, the question of which arguments are correct in natural language decomposes into a question of which ones are correct in an artificial language (which is usually a matter of a calculus or a system of formal semantics) and which natural language arguments are adequately represented by the artificial language ones. It is necessary to clarify the status of the artificial languages, their relationship to our actual reasoning and to the natural languages, and the criteria that would let us say that the former is an adequate model of the latter.

One set of problems concerns the newly established artificial languages. Artificial languages, or theories in the languages, may have various properties (such as consistency, completeness, decidability, etc.) and whether a particular language has such a property can be more or less easily determined by mathematical methods. In contrast to this, in the case of natural language the presence or absence of such a property is often not only not so easily found out, but it is even not clear how to make sense of them. Take, for example, such a property as consistency. Tarski (1944) was convinced that a natural language is not consistent because it contains all the ingredients needed to put together liar sentences. In contrast to this, consistency was usually supposed to be a condition *sine qua non* for an artificial language to be useful for

[5] See Lloyd (1978).

logic; and in any case its presence or absence is an unambiguous matter, usually it is subject to a mathematical proof.

The, the question is what such a proof of consistency of a formal language can tell us about natural language. One possibility is that it tells us nothing, because the artificial language was built precisely to improve on natural language in that it be provably consistent. Another possibility is that its consistency also indicates that the natural language is consistent, at least if cleansed from some problematic cases, some such ballast as, for example, the liar sentences, which are not essential to it. Or there is a possibility that it tells us that we should dispose of the amorphous and potentially vicious natural language in favor of one that is exactly delimited and provably unproblematic.

Another set of problems concerns the very nature of the relationship between natural and artificial languages. The original idea was that the relationship be quite straightforward – logical expressions of the former being replaced by their more rigorous variants and the extralogical ones being abstracted away. But we will see that, especially on the level of predicate logic, this gave way to the artificial languages acquiring structures largely independent of those of natural languages. As a result, "logical analysis," as a process of finding a formula appropriate for a natural language sentence, became a nontrivial kind of "art."

In this way, the artificial languages of logic acquired lives of their own and lost their obvious dependence on natural language. Again, one possibility would be to see this as a process of emancipation of the artificial languages, which were on their way toward superseding the natural ones. Another possibility is to insist that artificial languages cannot but be mere *simulacra* and that their "independent life" remains highly limited. In this case, there still may be a choice of attitude: their study may be seen either as merely an indirect way of studying natural language, or they may be seen as being in some sense autonomous.

Anyway, it seems that the emergence of artificial languages, and the logicians' relocation of most of their problems into the context of these languages, creates new kinds of philosophical questions that are both unprecedented and yet worth being taken seriously.

1.6 Genealogy of the artificial languages of logic

The key to the philosophical problems of modern logic thus appear to be an understanding of the artificial languages logic employs and of what appears to be their "independent life." The idea pursued in this book is that they were formed out of the rib of natural language in a process during which logical analysis or logical formalization became more than a mere replacement of words and expressions of natural language by their more rigorous regimentations.

The language of propositional logic, as the simplest artificial language, originated mostly from such replacements, and even its employment for the purposes of logical formalization its operators can often be considered as mere proxies for their natural language counterparts. (Though in many cases it can also be employed in creative manners where its operators' roles become much more than such proxies.) But especially along the way between this language and the language of predicate logic there came about an important change: the latter have come to rest on some syntactic rules (concerning quantifiers) which have no direct counterparts in natural language. Hence, finding a formula regimenting a given natural language sentence necessitated correlating two nontrivially different syntactic structures, opening up a space for various alternative solutions.

In this way, the language of predicate logic ceased to be directly tied to natural language; hence it acquired a life of its own. And as the structure of such a language resulted from the effort to put together an efficient calculus, to understand why the structure came out as it did means to understand the calculi and the problems to which their constitution responded (as well as the new problems which the constitution brought about). This is the task to be undertaken in the following chapters: to anatomize the most basic logical calculi to show how the structure of the language that underlies them has come into being.

A further step in the emancipation of the artificial languages of logic from natural ones came with the establishment of formal semantics in the writings of Tarski, Carnap, and others. This introduction is sometimes depicted in such a way that it was only *via* it that the formal language truly became a *fully-fledged* language, with not only a "syntax" but also a "semantics." I think that to endow a language with a "real" semantics is not at all so simple as assigning its expressions some set-theoretic constructs. I think then that we should see the establishment of formal semantics as an elaboration and refinement of the technical apparatus we put together to create (useful) models of natural languages as vehicles of our matter-of-fact reasoning, rather than as promoting our artificial languages to the level of "genuine" languages.

Hence, the leading idea of this book is that the formal languages of contemporary logic, which now largely determine the direction of logical investigations, evolved out of the effort of logicians to make useful models of those structures of natural language that are important from the viewpoint of argumentation. I believe we can understand them better if we reconstruct this evolution. This is not to say that those who developed the formal languages usually saw themselves as quasi-natural scientists building increasingly self-standing models of the activity of reasoning. Undoubtedly, many of them saw themselves as penetrating to deeper structures of thought or reality by discovering hidden "logical forms"; but here we do not want to reconstruct their intentions or their own understanding of what they were doing, but rather what we believe was really happening.

This means that in the upcoming chapters I will survey the story of the constitution of the common artificial languages of logic that, on the one hand, is well-known but which we must, on the other hand, retell in a new way that is not so familiar. Thus, I beg the reader for understanding my rationale for repeating things that are all too familiar to them; that rationale being that I want to rearrange the mosaic that these familiar things constitute so as to open a new vista on the nature of modern, formal logic and on the new philosophical problems it opens up.