

Preface

In this essay I develop an integrated theory of knowledge, truth, and logic. I do this by tying together, and further expanding, three projects I have been working on since completing my dissertation in 1989: a model of knowledge that combines Quinean and anti-Quinean elements, a substantivist theory of truth, and a philosophical foundation for logic. What led me to embark on this extended project was the realization that its three parts were so thoroughly interrelated that I could no longer either adequately explain or continue developing any one of them without connecting it to the others. My model of knowledge requires a non-traditional correspondence account of truth, as well as a new foundation for logic. My foundational account of logic centers on issues of veridicality (truth), and the methodology that makes a veridical foundation for logic possible is a non-traditional methodology that was developed for the model as a whole. Similarly, my substantivist theory of truth is motivated by epistemic considerations associated with the model. It exemplifies the model's approach to philosophical theorizing, and it develops a non-traditional conception of correspondence that is flexible enough to encompass all fields of knowledge, including mathematics and logic. A common thread passing through all three projects is the centrality of both freedom and friction (constraint) to knowledge, including philosophical and logical knowledge. In this essay I put more emphasis on friction; in a sequel, *Epistemic Freedom*, I hope to further pursue the second aspect, freedom.

It is impossible to introduce all the ideas developed in this essay in a brief preface. Instead, I will offer a bird's eye view of the essay, focusing on some of its ideas and leaving others to the work itself.

Underlying my approach to knowledge is an interest in "the basic human epistemic situation". Among the central elements of this situation are: (i) We live in a world of which we are a part. (ii) We desire to know and understand this world not just practically but also theoretically. But (iii) our cognitive resources are considerably limited and this renders the world highly complex relative to our epistemic abilities. Nevertheless, (iv) we humans are ambitious creatures and we aspire to know the world in its full complexity. What makes our aspirations achievable (to some significant degree) is the circumstance that (v) while we are cognitively limited in some ways, we are cognitively endowed in others. Our cognitive capacities are intricate, they involve not just sensory perception but also intellect and imagination, they include the ability to participate actively in the project of knowledge: plan, initiate, design, shift perspectives, critically evaluate our epistemic strategies and results, create technological tools, and so on. In light of this complex situation, a central question of epistemology is whether and to what extent we are capable of

acquiring theoretical knowledge of the world and how we can, do, and should go about it.

This question brings to the fore two fundamental principles of knowledge: *epistemic friction* and *epistemic freedom*. Knowledge requires friction or constraints. Most importantly, knowledge is substantially constrained by its target—the world. But knowledge also requires active involvement of the theorist in both discovery and justification, i.e., the exercise of freedom. Freedom and friction, however, are not disjoint. In particular, freedom itself is an important source of constraint. Thus, the constraints of truth, justification, and substantive (informative, explanatory) theorizing are imposed on us by ourselves, through norms that, exercising epistemic freedom, we ourselves create. These normative constraints are universal, applying to knowledge qua knowledge; hence they encompass *all* fields of knowledge, from the natural and social sciences to logic, mathematics, and philosophy. It is with these two basic principles of friction and freedom that I begin the essay.

Two of the main friction requirements applicable to philosophy are (i) grounding our system of knowledge (the compilation of our most advanced theories) in the world, and (ii) substantive theorizing. By a philosophical “grounding” I mean a critical explanatory and normative theory that tackles the central epistemic question delineated above, namely, whether, to what extent, and how we humans are in principle capable of acquiring genuine knowledge of the world, and what constraints such knowledge should satisfy. Accordingly, my project is the classical project of a “foundational” theory of knowledge. Today, however, the classical foundational project is widely viewed as a dead project. This is not surprising. Traditionally, the foundational project was associated with a highly problematic—arguably self-defeating—methodology, the so-called “foundationalist” methodology, and the failure of this methodology is viewed by many as a failure of the foundational project itself. In my view, this identification is unwarranted. The foundationalist methodology is just one methodology for pursuing the foundational project; nothing prevents us from devising a new foundational methodology that avoids its pitfalls.

My first task, therefore, is to develop a methodology for a “foundation without foundationalism”¹. The *foundational* methodology I develop in Part I differs from its *foundationalist* predecessor in being *holistic*, and from other new foundational methodologies in its emphasis and scope. I call it “foundational holism”. Instead of using rigid tools and demanding a strict ordering of the grounding process, it offers flexible, holistic tools, setting no preconceived demands either on the order in which the grounding is conducted or on the resources used. And instead of having a partly coherentist orientation and being limited to empirical knowledge, it is robustly world-oriented and universal. Furthermore, its conception of grounding is not “thin” or watered-down. Indeed, in certain significant ways foundational holism is

¹ A term adapted from Shapiro 1991.

more demanding than foundationalism. While most foundationalist methodologies give logic a “free pass” (as far as grounding it in the world is concerned), foundational holism does not. As a field of knowledge, logic requires a grounding in the world as much as any other field, though it must be grounded in an appropriate facet of the world. We may say that foundational holism puts the holistic method in service of a robust, world-oriented, universal foundational project.

It is important to clarify that foundational holism is not a holistic methodology in the sense that the smallest unit of knowledge is our body of knowledge *as a whole*. Instead, it regards our system of knowledge as a *structured network* of relatively independent units, interrelated yet preserving their own identity. Accordingly, it takes into account both similarities and differences in the conditions under which the knowledge provided by different branches of knowledge is acquired and justified, and it grants epistemologists the freedom to ground different units of knowledge in ways that reflect both their similarities and their differences.

My next step (Part II) is the development of a general model of knowledge that exemplifies the principles of epistemic friction and freedom as well as those of foundational holism. Among the distinctive characteristics of this model, in addition to holism, are its dynamic structure, its broad conception of reality, and its view of intellect and sensory perception as equal players in knowledge. My starting point is Quine’s model in “Two Dogmas of Empiricism”. Quine’s model has a significant interface with both reality and mind (periphery and center); it is a rich holistic model, with an elaborate network of connections between diverse units of knowledge; and it rejects the traditional divisions of units of knowledge into those grounded in reality and those grounded *solely* in the mind. But while I adopt these elements of Quine’s model, I renounce others.

One of my criticisms of Quine’s model focuses on its overly static structure. Although Quine’s model is more flexible than most traditional models, its structure is still exceedingly static and its two poles—periphery and center—are exceedingly narrow. Specifically, its periphery is limited to observational sentences and its center to pragmatic considerations and conventional postulations. As a result, highly abstract disciplines, such as mathematics and logic, are barred from the periphery. Untangling the metaphor, logic and mathematics, in Quine’s model, are not subject to exacting standards of veridicality; as such, they are largely frictionless. Logic and mathematics can be challenged obliquely by the world, through observational statements, but this challenge is limited to those parts of their content that impact observational statements. There is no room in Quine’s model for a veridical challenge to logic and mathematics as branches of knowledge in their own right.

This limited view of abstract knowledge is closely connected to Quine’s deep, and quite radical, empiricism—another target of my criticism. Not only is Quine’s conception of center and periphery relatively narrow, but so is his conception of mind and world. Quine never asks whether objects in the world have abstract features, nor does he ask whether humans have resources for cognizing such features.

To me, it is inconceivable that humans would reach the level of knowledge they have without a significant contribution of intellect (as something distinct both from sensory perception and from pragmatic conventions), but Quine completely neglects the role of intellect in knowledge. He considers, and rightly rejects, supernatural means of discovery and justification, but human intellect is not supernatural. Human intellect cannot be identified with either telepathy or clairvoyance, nor is it related to Greek deities, or the like. In the entire Quinean corpus there is no consideration of intellect as a crucially, or even potentially, significant cognitive resource.

My own model differs from Quine's on both these counts. First, it is a dynamic model: center and periphery are job descriptions rather than fixed locations, and each discipline moves from periphery to center and vice versa according to the task at hand. In this way, observational science is subject to conceptual, linguistic, and pragmatic norms; logic and mathematics are subject to veridical norms. Second, both my center and periphery are broader than Quine's, representing a broader conception of world and mind. Objects in the world, in my model, have both abstract and concrete features, and our system of knowledge aims at knowledge of both. My paradigm of an abstract feature is a formal property, for example, a cardinality property. Most properties of objects—including physical properties of physical objects—have cardinality properties (for example, the property of being a moon of Earth has cardinality ONE). These properties, like many physical properties, are governed by laws (regularities), and our epistemic goals include knowledge of those laws.

My conception of mind is also broader and more open-ended than Quine's. Far from being purely pragmatic, it includes intellectual resources of discovery and justification, working in tandem with sensory perception. My paradigm of intellectual knowledge, too, is relatively broad. I call it "figuring out". Figuring out, in the everyday sense of the word, is something that humans engage in at any age (babies constantly figure out things), in any context, and on any level of abstraction (from figuring out why a given computer crashed to figuring out whether a given logical system is complete). It is also something that goes beyond mere sense perception, involves discovery, and is not limited to pragmatic considerations. The theoretical investigation of figuring out is a job both for psychology and for philosophy. The latter job is within the scope of *Epistemic Freedom*. Here I emphasize the centrality of figuring out (or something like it) for all knowledge, abstract as well as empirical.

One of the main instruments of epistemic friction is a standard of truth, one that measures the success of theories with respect to the *correctness* of what they teach us about the world. Such a standard is required by our model of knowledge, and in Part III I develop a theory of truth that introduces such a standard. This theory is substantivist (rather than deflationist), and its standard of truth is a correspondence standard, albeit a non-traditional correspondence standard.

Today, many philosophers are attracted to the deflationist approach to truth. But in my view truth is too substantial, too complex, and too important to be adequately

approached in a deflationist manner. Having found the reasons put forth in support of deflationism unconvincing, I ask: Is there anything inherently problematic about truth, anything that stands in the way of substantive theorizing about it? My answer to this question is positive. There is something inherently problematic about truth, or rather about the way we, philosophers, approach it. On the one hand, truth applies to an enormous array of highly diverse and often complex cognitions (thoughts, sentences, theories); on the other hand, philosophers traditionally expect the theory of truth to take the form of a single and simple definition or definition schema. But there is no reason to surmise that truth applies to all cognitions in the same way or always in a simple manner. One way to deal with this problem is to overlook the complexity of truth so as to preserve a neatly unified conception of its theory. This way leads to a frictionless theory. Another is to recognize the complexity of truth, and in particular the tension between *unity* and *disunity* (diversity) inherent in it, and to look for ways to resolve this tension.

In our time, the tension between unity and disunity has been studied primarily in connection with science (rather than with truth). The subject matter of science—nature as a whole—is a broad and diverse subject matter, and the tension between unity and disunity poses a serious problem for scientists. But this problem is not unsolvable: “every science needs for its healthy growth a creative balance between unifiers and diversifiers” (Dyson 1988: 47). This solution has nothing to do with the empirical nature of science and as such it is applicable to all fields, including philosophy. My solution to the problem of the unity and disunity of truth follows the same line of reasoning: the theory of truth is best viewed as a system of substantive principles of varying degrees of generality. Some of these principles are unifying in nature, others are attuned to the diversity of truth. All, together, offer a “fruitful balance” between unity and diversity.

To arrive at these principles I approach truth from a cognitive rather than from a linguistic perspective. Instead of starting with the totality of true sentences and asking “What is common to all these sentences?”, I start with the basic human cognitive situation and ask “Under what conditions does truth arise as a significant standard for human cognitions?” My answer is that truth arises at the juncture of three fundamental modes of human cognition which I call “immanence”, “transcendence”, and “normativity”. In the immanent mode we direct our cognitive gaze at the world (thinking in the way one thinks when one stands *within* a theory); in the transcendent mode we move beyond our immanent cognitions to a standpoint in which we hold both these cognitions and the world in view; and in the normative mode we ask certain critical questions concerning the relation between the two. I call the principle that truth arises at the juncture of these three modes of thought “the fundamental principle of truth”. (It is worthwhile to note that the appeal to transcendence does not involve commitment to a “God’s eye view”: transcendence, here, is a move from one human standpoint to another, one which is more powerful in certain respects, yet still decidedly human.)

My second core principle of truth is a non-traditional *correspondence* principle. Like the traditional correspondence principle, it sets non-trivial demands on our cognitions with respect to their relation to the world; but unlike traditional correspondence it does not determine in advance what pattern this relation can, does, and should exhibit. In thinking about correspondence I am guided by two ideas. Given the value we place on correctness, we need a robust correspondence standard. But as a standard for creatures in a complicated cognitive situation, fraught with difficulties and requiring considerable ingenuity and creativity, it must be flexible and adjust to the ways we tackle diverse cognitive challenges. I call this type of correspondence “manifold correspondence”.

To provide an example of a non-traditional pattern falling under “manifold correspondence” I turn to mathematics. Here, too, my approach is a bit uncommon. Instead of starting with existing mathematical theories and asking what facets of the world they correspond to, I start with the world and ask: is there any aspect of the world whose knowledge requires something like a mathematical theory? And instead of using current language, which historically developed in a partly haphazard way and for multiple purposes, as a strict yardstick, I ask: how, given our cognitive resources and limitations, would we best employ the language we have to develop theories of this aspect? My answers are, first, that there *is* a particular facet of the world whose knowledge requires a mathematical discipline, namely, its *formal* facet; and second, that the linguistic expressions we use to describe this formal facet correspond to it in a “composite”, indirect way. First-order arithmetic, on this account, is a theory of finite cardinality properties rather than of numerical individuals in the world. Its ability to reach its target in the world is due, to a large degree, to the ingenious, yet disciplined, exercise of epistemic freedom by laypeople and mathematicians.²

Some of the thorniest challenges facing any foundational methodology, model of knowledge, and theory of truth are posed by logic. Logic seems to resist many of the requirements imposed on all other disciplines, and this tempts us to grant it special allowances. Logic is often released from the friction requirements of substantiveness and grounding-in-reality, from the need to face the world in the periphery, from a robust standard of truth, and so on. Instead, logic is frequently regarded as purely conceptual, true by convention, trivial, obvious, a mere game, a matter of choice of language, a matter of convenience, and so on. This is not surprising, given its fundamentality, generality, abstractness, and formality. But exempting logic from central epistemic requirements is also problematic. After all, due to its crucial—theoretical as well as practical—role in our system of knowledge, it is especially important to make sure that logic does not malfunction, does not introduce error into, and potentially compromise, the integrity of our entire system. This means that logic—as much as, and even more than, any other discipline—requires a systematic

² The route to full-fledged mathematics involves expanded principles of a similar kind.

veridical foundation. In Part IV, I delineate an outline of a theoretical foundation for logic. This foundation conforms to the universal principles of epistemic friction, explains why, and in what sense, logic has a foothold in the periphery, and subjects logic to a robust standard of truth. In a nutshell: a claim of logical consequence is true if it corresponds to (an appropriate instance of) an appropriate law governing the world; false otherwise.

Among the advantages of this account is its ability to offer a substantive explanation of the characteristic traits of logic: its considerable generality, topic neutrality, basicness (fundamentality), especially strong modal force, normative power, etc. The account explains these traits in terms of a highly informative notion: *invariance*. Roughly, these traits are due to the strong degree of invariance of those features of the world that logic is grounded in (which, in turn, are responsible for the special strength of those laws that true statements of logical consequence correspond to).

Another advantage of the account is its ability to unify logic and mathematics without reducing one to the other. This it does by providing a *common grounding* for both while distinguishing their *job descriptions*. Both logic and mathematics are grounded in the *formal facet of the world*, but while mathematics *studies* the formal, logic *develops a method of inference* based on it. The formal itself is explained in holistic, semi-Aristotelian terms. This account shares the methodological advantages of traditional logicism while avoiding one of its major (though largely overlooked) methodological shortcomings. Whereas logicism enables us to solve two difficult problems—the foundational problem of logic and the foundational problem of mathematics—by reducing them to one problem—the foundational problem of logic—it leaves logic itself without a grounding. In contrast, our account offers a joint grounding for logic and mathematics without ignoring their difference.

But the most important advantage of our account is its ability to ground, and provide a substantive explanation of, the *veridicality* of logic: explain why and how logic works in the world, under what conditions its claims are true or false, what distinguishes correct and incorrect logical theories, and so on. In this way it brings logic in line with all other disciplines, while recognizing—and explaining—its unique characteristics.

The essay ends with a forward look at *Epistemic Freedom*. Epistemic freedom is inherently complicated: without it, we are unable to achieve, or even pursue, our epistemic goals; but in exercising it we introduce risk and uncertainty into the pursuit of these goals. The present volume emphasizes the complementarity of friction and freedom, construes freedom as a central element of the basic human epistemic situation, and integrates it into its model of knowledge, its account of intellect, its theory of truth, and its logical foundation. Still, much work remains to be done. We need to study the scope and limits of epistemic freedom, investigate its tension and cooperation with epistemic friction, learn how to integrate it into successful epistemic strategies, guard against its excesses, and so on. Most importantly, we need to better understand its working, its role in discovery and justification, the cognitive faculties exercising it, and its manifestations in all fields of philosophical study.