

UC San Diego

UC San Diego Previously Published Works

Title

Epistemic Friction: Reflections on Knowledge, Truth, and Logic

Permalink

<https://escholarship.org/uc/item/0fj290kv>

Journal

Erkenntnis: An International Journal of Analytic Philosophy, 72(2)

ISSN

1572-8420

Author

Sher, Gila

Publication Date

2010-03-01

DOI

10.1007/s10670-009-9202-x

Peer reviewed

Epistemic Friction: Reflections on Knowledge, Truth, and Logic

Gila Sher

Received: 1 July 2008 / Accepted: 16 November 2009 / Published online: 11 December 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract Knowledge requires both *freedom* and *friction*. Freedom to set up our epistemic goals, choose the subject matter of our investigations, espouse cognitive norms, design research programs, etc., and friction (constraint) coming from two directions: the object or target of our investigation, i.e., the world in a broad sense, and our mind as the sum total of constraints involving the knower. My goal is to investigate the problem of epistemic friction, the relation between epistemic friction and freedom, the viability of foundationalism as a solution to the problem of friction, an alternative solution in the form of a neo-Quinean model, and the possibility of solving the problem of friction as it applies to logic and the philosophy of logic within that model.

1 The Problem of Epistemic Friction

The problem of epistemic friction is the Kantian problem of providing our theories with appropriate resistance so they do not “hover idly in thin air”. Kant illustrates the problem with his “dove” metaphor: “The light dove, cleaving the air in her free flight, and feeling its resistance, might imagine that her flight would be still easier in empty space...” (Kant 1781/1787, p. A5/B8-9). Kant’s prototype of a frictionless discipline is traditional metaphysics which purports to provide knowledge of “things in themselves” through our conceptual faculty (“understanding”) alone, having no adequate tests for its ideas or means for correcting its theories.

Kant is well aware, however, that the threat of frictionless theorizing is present in science and mathematics as well. Friction in these branches of knowledge is attainable by satisfying a series of requirements, starting with general logical requirements and ending with more specific requirements. The latter are rooted in two fundamental

G. Sher (✉)

Department of Philosophy, UCSD, La Jolla, San Diego, CA 92093-0119, USA
e-mail: gsher@ucsd.edu

conditions on human cognition: (1) human cognition requires sensible intuition, (2) human cognition requires conceptual synthesis. The idea is that human cognition is both receptive and creative; its receptivity is intuitive and sensual, its creativity synthetic and conceptual. Since human intuition is sensual and human synthesis conceptual, human knowledge is restricted to the world as it is intuited through our senses and synthesized by our concepts—the so-called world of appearance. In contrast, neither the world as it is in itself nor the world of pure sense data is cognitively accessible to us. We can *think* of “things in themselves”, but we cannot *know* them.

The two elements of friction in Kant’s philosophy are thus the external world as it affects us through our senses and the structure of our own cognitive apparatus with its specific forms of intuition and synthesis. These constraints affect different sciences differently: Whereas natural science is bound by testimony of our senses directly, through experiment and observation, mathematics is bound by our senses only indirectly, through its applicability to experience; while physics is required to establish *causal* connections between physical phenomena, geometry is required to *construct* its objects in intuition and *prove* its theorems by “demonstration”.

Kant’s concern with frictionless theorizing, however, does not extend to all fields of knowledge. A notable exception is logic. Logic creates friction for other sciences, e.g., by forbidding contradictions; yet logical knowledge itself—knowledge of the law of non-contradiction, the law of excluded middle, modus ponens, etc.—appears to require no substantial friction. This is not just a matter of the analyticity of logic: non-logical analytic statements *are* subject to constraints, namely, logical constraints. But logic, according to Kant, is special. Logic is a purely negative discipline, setting limits to human knowledge without producing new knowledge. And this somehow shelters it from the danger of frictionless theorizing.

The term “friction”, in its epistemic sense, appears in McDowell (1994). Referring to the Kantian faculty of “spontaneity” (which he roughly equates with conceptual activity), McDowell says: “We need to conceive... spontaneity as subject to control from outside our thinking, on pain of representing the operations of spontaneity as a *frictionless spinning in a void*” (*ibid* p. 11, my italics). McDowell’s conception of the problem of friction centers on Kant’s condition that cognition is grounded in the world through sensible intuition. Relating to Kant’s dictum that “[t]houghts without content are empty” (Kant, *ibid* p. A51/B75), McDowell elaborates: “[I]f our freedom in empirical thinking is... not constrained from outside the conceptual sphere, that can seem to threaten the very possibility that judgements of experience might be grounded in a way that relates them to a reality external to thought” (McDowell, *ibid* p. 5).

McDowell’s main concern, however, is not with friction itself, but with a trap we are likely to fall into when trying to establish friction: “the myth of the given.” In seeking to ground our theories in reality, McDowell says, we are tempted to postulate the existence of a brute “given.” But appealing to the given is of no use, since being under the control of an external force does not by itself constitute *justification*. McDowell’s solution (following Sellars) is: receptivity does not make an independent contribution to the cooperation between intuitions and concepts; rather, intuition already contains a conceptual element. This is a radicalization of Kant’s second dictum, “intuitions without concepts are blind” (Kant, *ibid* p. A51/B75).

My own conception of epistemic friction is similar to Kant's and McDowell's in certain respects, different in others. The problem of friction, for me, is the problem of setting adequate constraints on our system of knowledge so we avoid empty theories on the one hand and maximize genuine knowledge on the other. The problem is both a problem of design and a problem of explanation. The problem of design is the problem of figuring out how to construct our theories so they are subject to constructive friction; the problem of explanation is the problem of identifying the main mechanisms of epistemic friction, describing their principles, and critically evaluating their contribution to knowledge. The two problems are interconnected: our design depends on our understanding, and our quest for understanding is at least partly motivated by difficulties in design. So far, this is a variant of the Kant-McDowell approach. But my approach differs from theirs in several ways:

1.1 Dimensions of Friction

Both Kant and McDowell restrict the problem of friction to embeddedness of knowledge in reality. While this is undoubtedly a central aspect of friction, friction has other dimensions as well. In particular, a theory can be empty, or fail to have adequate friction, by being "trivial" or "non-substantive", i.e., by adopting exceedingly weak standards of theorizing, discovery, explanation, etc. Further types of friction consist of rational, pragmatic, aesthetic, and other constraints and desiderata (For example, a theory subject to the pragmatic requirement of simplicity is, everything else being equal, subject to greater epistemic friction than a theory not subject to this requirement). Still another type of friction issues from physical, biological, psychological, and social constraints set by our nature and environment.

1.2 Logic

Neither Kant nor McDowell seem to have regarded the problem of friction as applicable to logic. In my view the problem is applicable to *all* branches of knowledge, logic included. The problem, in fact, is especially crucial for logic due to its central role in all areas of knowledge and its close connection with truth, which is highly relevant to embeddedness in reality.

1.3 Philosophy

Kant regards the problem of friction as unsolvable for philosophy, and today the possibility of substantive and "factual" (or theoretical) philosophy is denied by many philosophers. The prevalent methodology is deflationist, minimalist, relativist or quietist, and it rules out both the need for and the possibility of friction in philosophy. No one, however, to the best of my knowledge, has given a decisive argument against the possibility of a substantive and grounded philosophy, and *prima facie* there are a number of examples in the philosophical literature of theories whose substantiveness, at least, is difficult to deny (Kant's epistemology is just one

example). This leads me to regard friction as a challenge to philosophy, or at least an open question, rather than a demonstrably unsolvable problem.

1.4 Reality

Both Kant and McDowell identify embeddedness in reality with embeddedness in *empirical* reality, and such a view might lead one to conclude that if logic and philosophy are not exempt from the embeddedness requirement, they must be reconfigured as empirical disciplines—a view quite popular today among hard-core naturalists. My own position is that embeddedness in reality is not, in principle, restricted to empirical reality. Since I don't regard either logic or philosophy as exempt from this requirement, some might say I am committed to Platonism. I believe this is not the case. Naturalism and Platonism are not the only options open to the philosopher (or the logician), and the question of what aspects of reality logic and philosophy are embedded in is answerable in principle outside both ideologies.

In this paper I will focus on two requirements of epistemic friction: *groundedness* (*embeddedness*) *in reality or veridicality*, and *substantiveness*, especially, but not only, as they apply to logic and the philosophy of logic. I will sketch a general model of friction which is neither naturalistic nor Platonistic, and I will explain how a logic grounded in reality in a substantive manner is possible within that model.

Most philosophers today are pessimistic about the prospects of a substantive philosophical grounding of logic, and their pessimism is (or seems to me to be) rooted in two implicit assumptions: (1) a “foundationalist” assumption, and (2) an “insurmountable methodological obstacles” assumption. The foundationalist assumption says that a foundation for logic satisfying the friction requirements would have to be foundationalist, but a foundationalist foundation for logic is impossible: Logic, according to foundationalism, lies at the bottom of the foundationalist hierarchy, hence there is nothing more basic than logic and it is impossible to provide a theoretical foundation for logic. There are simply no conceptual and theoretical tools for constructing such a foundation. The “insurmountable methodological obstacles” assumption says that logic and philosophy face unique methodological problems, not faced by either science or mathematics, and these problems are unsolvable. I believe both assumptions are unwarranted. There are alternatives to logical foundationalism, compatible with epistemic friction; and the methodological problems facing logic and philosophy are neither radically different from those facing science nor unsolvable.

2 Epistemic Freedom

The problem of epistemic friction is complemented by another problem, that of *epistemic freedom*. Two major elements (aspects, types) of epistemic freedom are: (a) independence from the world, and (b) cognitive choice. Freedom of choice is positive freedom, freedom from the world is partly negative; freedom of choice is active, freedom from the world is partly passive; freedom of choice implies freedom from the world, but not the other way around.

2.1 Freedom From the World

While knowledge is affected by features of its subject-matter—the world, in a broad sense—it is also affected by other things. Among those we may distinguish three: (1) the natural structure of human cognition and the social and natural forces acting upon it; (2) the rational (transcendental, etc.) structure of human cognition, and (3) active, intentional intervention (choice, decision, design, etc.), which falls under the second category.

2.2 Cognitive Choice/Active Design

The construction of a system of knowledge is to a large degree an intentional project, involving a wide array of deliberate, voluntary choices and actions, including acts of design, conjecture, derivation, calculation, invention, ratiocination, problem solving, abstraction, generalization, definition, observation, experimentation, revision (expansion, contraction, replacement), and so on. All these are realizations of “positive” freedom.

The problem of epistemic freedom is not disjoint from the problem of epistemic friction. On the contrary, freedom introduces new forms of friction, and in certain cases it is a prerequisite of friction. Both types of freedom mentioned above generate constraints on knowledge: (a) We should not develop theories whose verification (falsification) requires brain structures, perceptual apparatus, life spans, or physiological attributes we do not possess; nor should we attempt to achieve knowledge beyond the rational and transcendental boundaries of human cognition (whatever they are). (b) Freedom of epistemic choice is essential for generating goals, norms, methodological guidelines, standards, and desiderata, all of which are important forms of friction. Although the model I will sketch below is a model of both friction and freedom, in this paper I will focus on the former. In particular, my philosophical account of logic will be restricted to friction, leaving freedom for another occasion.

3 The Illusion of Foundationalism

Although the term “epistemic friction” is new, the problem of epistemic friction has been a central problem for philosophy since its inception. The most influential strategy for resolving this problem is the foundationalist strategy, which seeks to establish human knowledge on a firm foundation of indubitable beliefs and reliable knowledge-extending procedures. The idea is that a given belief constitutes genuine knowledge iff¹ it is either foundational or obtained from foundational beliefs by reliable procedures.

Foundationalism purports to meet the two requirements of *embeddedness in reality* and *substantiveness* in a simple and straightforward manner: (a) Foundational beliefs are embedded in reality directly, through direct experience or

¹ If and only if.

intellectual intuition; non-foundational beliefs are embedded in reality indirectly, through reliable knowledge-extending relations that connect them to foundational beliefs. (b) Foundational beliefs are directly substantive—satisfying the most stringent standards of discovery and having ample explanatory power; non-foundational beliefs inherit their substantiveness from foundational beliefs through the knowledge-transmitting relations.

Foundationalism has problematic consequences for logic and philosophy, and these, I believe, point to basic structural flaws in the foundationalist methodology. Due to limitations of space, I will restrict myself here to logic.

Because of its extreme generality, basicness, and normative force, logic is commonly placed at the base of the foundationalist “pyramid”. This, however, creates problems for logic: logic can provide (or partake in providing) a foundation for other fields of knowledge, but no field of knowledge (or combination of fields of knowledge) can provide a foundation for logic. The result is either that logic is altogether ungrounded, or that logic is grounded in something other than a field of knowledge, i.e., logic has a non-theoretical foundation. The former alternative is a *reductio ad absurdum* of the foundationalist approach: The main support for a given theory or discipline, according to foundationalism, comes from below, from theories or disciplines more basic than it; but if the most basic theories and disciplines, those on whose soundness the integrity of the entire structure of knowledge rests, are themselves devoid of foundation, the entire system is unfounded.

Defenders of foundationalism might argue that grounding (explanation, justification) must stop at some point; why not at the foundation? Granting the practical necessity of stopping the process of grounding at some point, not all points are equal. It would be harmless for foundationalism to stop justification at some “remote” discipline, i.e., a discipline D such that an error in D has few ramifications for the rest of the system, but not at a foundational discipline like logic, whose integrity is crucial for the entire system.

It is a predicament of foundationalism that the foundational problem is unsolvable for it: infinite regress is not permitted, grounding the foundation is structurally impossible, and stopping the process of grounding short of the foundation would not solve the problem.

What about a non-theoretical foundation for logic? Is such a foundation possible? Three contenders for such a foundation are: (1) intuition, (2) obviousness, and (3) conventionality. Let us briefly examine these possibilities:

3.1 Intuition

The idea that logic is grounded in intuition is advocated by Gödel. Take, for example, the logical law (rule of derivation) Modus Ponens. The proposition stating this law, according to Gödel, “*can directly be perceived to be true*” (Gödel 1953–1959, p. 347), and this perception is “executed” by something like a sixth sense, analogous to the five “recognized” senses. The two kinds of perception differ in that “while through sense perception we know particular objects and their properties and relations, with [logical] reason we perceive the most general (namely the “formal”) concepts and their

relations” (*ibid.*: 354). Rational perception reveals a “second reality... completely separated from” space–time reality, if just as objective (*Ibid.*: 353 fn.).

The view that logical knowledge is grounded in rational intuition is, however, either too extreme to evade the pitfalls of unrestrained Platonism or too moderate to provide the requisite support for foundationalist epistemology. There is no need to elaborate on the former here, but the latter requires some explanation. Several contemporary philosophers (e.g., Alston 1976a) attempt to salvage foundationalism by arguing that foundational knowledge need not be infallible, need not be self-sufficient, and need not provide a foundation for our entire system of knowledge, first- as well as higher-order. When it comes to logic, Gödel himself, in fact, adopts a rather moderate Platonism: he acknowledges the existence of errors in rational intuition, regards rational intuition as but one component in logical knowledge and admits non-intuitive ways of grounding logic, for example, theoretical and/or practical success.

But while weaker versions of foundationalism might be successfully incorporated in *non*-foundationalist epistemologies, they appear to undermine foundationalism as such. Consider a foundationalist theory that admits the possibility of error in our intuition of basic principles and accordingly assigns to these principles probabilities smaller than 1. Suppose it assigns probability .9 (of transmission of truth) to some rule of inference, while lowering the cap on bona fide items of knowledge to, say, probability .8. Using a best case scenario—initial premises have probability .999...—three successive applications of this rule will carry us from established items of knowledge to items well below our threshold—clearly, an unacceptable result.

Next, let us consider the suggestion that intuition is not self-sufficient, that it is just one constituent in the grounding of basic principles. This suggestion is very attractive (for example, it neutralizes Sellars’ charge that foundationalist theories fall hostage to the “myth of the Given”), but its viability within foundationalism is questionable. According to foundationalism our system of knowledge is a linear or tree-like structure, where every item of knowledge is grounded in reality by a finite chain, descending from less basic to more basic elements, until reaching the solid rock of brute reality. If intuition is just one medium of grounding knowledge in reality, the foundationalist must identify other media; yet no (serious) accounts of appropriate media have been forthcoming.

Finally, recent foundationalists (e.g., Alston 1976b, 1983) have drawn a line between first- and second-order knowledge, claiming foundationalism has the responsibility of grounding first-order but not second-order knowledge. This move contradicts foundationalism’s claim to provide a foundation for our *entire* system of knowledge, and it is especially unsatisfactory from the perspective of epistemic friction, according to which embeddedness in reality is a *universal* requirement, hence applicable to higher-order as much as to first-order knowledge. This requirement could, in principle, be met by a two-pronged foundationalism, providing distinct foundations for first- and higher-order theories, but I know of no bifurcated foundationalism that fulfills this requirement with respect to higher-order theories.

3.2 Obviousness

Epistemic “obviousness” can be interpreted either as pointing to a special faculty of “intuitive” knowledge or as common-sense obviousness. Since we have already dealt with the former, let us now turn to the latter.

The idea that logic is grounded in common-sense obviousness is open to multiple criticisms. First, judgments of obviousness are clearly fallible. Indeed, the development of a systematic science as opposed to everyday opinion is at least partly a reaction to the fallibility of common-sense obviousness. Second, the idea that logic is obvious is a vague idea. The unclarity of this idea is not just inconvenient; it threatens the claim that logic requires no theoretical grounds. To understand the sense in which logic is obvious (and sanctioned by its obviousness), we have to understand (a) what features of logical truths/inferences make them obvious (what features distinguish them from non-obvious truths/inferences and from obvious truths/inferences which are not logical), and (b) what the normative force of this obviousness is. In short, to understand the view that logic is grounded in the obvious we have to develop a *theory* of the kind of obviousness that logic is grounded in. But this would undermine the view that logic does not require a theoretical foundation.

Finally, grounding logic in our common-sense judgments of obviousness brings to the fore the problem alluded to above: On the one hand foundationalism’s attitude toward knowledge is exceptionally critical, on the other hand it is remarkably uncritical. On the one hand foundationalism rejects any item of knowledge that lacks a foundation, on the other hand it sets outrageously low standards of approbation (mere “common-sense obviousness,” i.e., mere *appearance* of obviousness) for items constituting the “foundation”. The result, however, is not a “stand off” between the two forces. Due to the non-symmetric nature of foundationalism, low standards at the bottom mean low standards for the entire structure, while high standards at the top do not have the same effect.

3.3 Conventionality

Another solution to the problem of grounding logic in reality is conventionality. Logical theory, according to, e.g., Carnap, is conventional rather than factual, and as such its lack of groundedness in reality is innocuous, i.e., cannot cause error in knowledge. To introduce error into our system of knowledge a theory must say (or imply) something false about the world, but logical theory says nothing about the world—it merely introduces a set of linguistic conventions for saying things about it.

This solution, too, is open to multiple criticisms:

3.3.1 *Conventionality Trivializes Foundationalism*

In Russell’s words: “The method of ‘postulating’ what we want has many advantages; they are the same as the advantages of theft over honest toil” (Russell 1919, p. 71). Essentially the same criticism was issued by Quine (1935): If logical

theses can be established by convention, why can't any thesis whatsoever (physical, mathematical, biological, or what have you) be so established?

3.3.2 *Logical Conventionalism Violates Foundationalism's Injunction Against Infinite Regress and Circularity*

In Quine's words: "the difficulty is that if logic is to proceed *mediately* from conventions, logic is needed for inferring logic from the conventions" (*Ibid* p. 104).

3.3.3 *The Introduction of Even Seemingly Innocent Conventions is Not Risk Free*

Take Prior's (1960) example, according to which we conventionally introduce a new binary logical connective, "tonk", and two rules of inference characterizing it, (1) $\Phi \vdash \Phi \text{ tonk } \Psi$, and (2) $\Phi \text{ tonk } \Psi \vdash \Psi$. This may seem an innocuous convention, but its use would wreak havoc in our system of knowledge. Belnap (1962) proposed a solution to the problem, namely, restricting ourselves to *conservative* conventions (conventions that do not affect those parts of our system of knowledge that do not involve the new vocabulary). But this would be of no help. Belnap's solution is useful in, and is only intended for, a situation in which logical conventions are introduced into a *pre-existing* logical system which presumably is (or has a sufficiently rich subsystem which is) grounded in something other than convention. In this situation the pre-existing part of the system would constrain the use of convention to expand it. But the conventions the foundationalist is interested in must ground our *entire* logical system; so there is no pre-existing (sub-)system that could constrain them. In particular, one cannot ground an entire logical system in conventions without using non-conservative conventions.

A natural response to the problems of foundationalism is to give up the foundationalist approach altogether. But giving up foundationalism is not a simple proposition. Non-foundationalism commonly takes the form of coherentism, relativism, deflationism, or naturalism, each of which is highly problematic from the point of view of epistemic friction. The problems with the first three are straightforward: neither coherentism nor relativism satisfies the groundedness requirement, and deflationism does not satisfy the substantiveness requirement. The situation with naturalism is more involved. Naturalism acknowledges only empirical knowledge, but this is highly problematic with respect to logic and philosophy: First, it is questionable whether logic and philosophy can be reduced to empirical disciplines. Second, naturalism's treatment of the problem of friction in logic and philosophy is another case of "theft over honest toil": naturalism "solves" the problem of friction in philosophy simply by eliminating those parts of philosophy that are not amenable to empirical friction, and it ignores friction in logic altogether (Naturalism's treatment of mathematics is also problematic, but I will not go into this here).

In the next section I will sketch an outline of a model of knowledge designed to satisfy the friction requirement with respect to all disciplines, including logic, while designating a place for epistemic freedom, to be worked out elsewhere. The model is neither foundationalist nor coherentist, and it is not relativist, deflationist or

naturalist either. It incorporates many themes from Quine while rejecting others, and as such it may be viewed as a “Neo-Quinean” model.

4 A Neo-Quinean Model of Knowledge

Quine’s “Two Dogmas of Empiricism” (1951) is a fork in the road. One of its paths, the naturalistic path, has been thoroughly traveled; the other path is still largely unexplored. I believe the second fork of “Two Dogmas” is deeper and more interesting from the point of view of epistemic friction, and in this section I will briefly explore and develop it. My goal, however, is not to work out the second fork of Quine’s model as Quine himself would have done it. My aim is to develop an epistemic model that stands on its own and show how it balances the dual requirements of friction and freedom.

The antithesis of Quine’s “Two Dogmas” model is Carnap’s positivist model. The latter is a two-tier, thoroughly dualistic model, dividing our system of knowledge into two separate parts, Science and Meta-science. The natural and social sciences reside in Science; logic, mathematics, and the legitimate areas of philosophy in Meta-Science. Science is world- or fact-oriented, empirical, subject to standards of truth, and governed by norms of evidence. Meta-science is language-oriented, conventional, not subject to standards of truth, and governed by purely pragmatic norms. The model is characterized by a series of dichotomies: the analytic versus the synthetic, the external versus the internal, the factual versus the conventional, and so on. While some of these dichotomies are not epistemic, they have substantial epistemic ramifications. Take, for example, the analytic-synthetic dichotomy (AS). AS is a linguistic dichotomy, but it is not just any linguistic dichotomy—it is a *semantic* dichotomy having to do with a central epistemic notion, the notion of *truth* or *truth condition*. We can explain the epistemic import of AS as follows: AS postulates a *semantic* division between statements and theories whose truth is grounded in matters of fact and those whose truth is grounded in something other than fact—meaning, convention, etc. This semantic bifurcation *induces* an *epistemic* bifurcation of statements and theories into those whose acceptance, justification, and revision is based on factual or veridical standards and those whose acceptance, justification, and revision is based on conventional or pragmatic standards. To justify a synthetic statement we look for the kind of evidence that is pertinent to fact, to justify an analytic statement we look for the kind of evidence that is pertinent to convention. As a result, a system of knowledge based on AS is inherently dualistic: one part of it is governed by factual norms, the other part by purely pragmatic norms.

Such an epistemic system is highly problematic with respect to both friction and freedom: Theories in the “conventional zone” fail to satisfy the requirements of groundedness in reality and substantiveness, while theories in the “factual zone” fail to satisfy the requirement of (active) freedom.

Quine’s revolution consists in rejecting Carnap’s bifurcation. Instead of the positivist dichotomies, Quine introduces a series of unificatory theses: the Negative Analytic-Synthetic thesis or NAS (my terminology), the Inseparability of Language

and Theory, Antireductionism, Universal Revisability, Holism, Interconnectedness, etc. These are augmented by a few additional theses: the Center-Periphery thesis (CP), Pragmatism, Underdetermination, Realism, Scientific Empiricism, etc. The pivotal thesis is NAS. Like AS, NAS is, on the surface, a linguistic thesis, but its deep content is epistemic: Our system of knowledge is a unified system, comprising logic, philosophy and mathematics alongside the social and natural sciences. Each statement and theory in the system is both factual and pragmatic, both governed by veridical norms and governed by pragmatic norms: the norms of (correspondence) truth, evidence, and justification on the one hand; those of explanatory power, fruitfulness and simplicity on the other.

The model appears to satisfy both the requirement of universal friction and that of universal freedom: *Every* discipline is subject to the constraints of friction—groundedness in reality and substantiveness—and, at the same time, enjoys the privileges of freedom—the freedom to manipulate its theories and their statements according to pragmatic preferences. Logic, in this model, must agree with reality, physics has great latitude in resolving its theoretical and even experimental problems.

NAS, in particular, plays an important role in rendering friction and freedom universally satisfiable. This can be clearly seen by considering (what, on my construal, are) the central epistemic principles underlying it:

4.1 The Complementary Principles of the Unpredictability of Nature and of Maximizing the Maneuverability of Our Cognitive “Forces”

AS gives rise to a traditional methodology of verification and falsification. Synthetic truth is factual and as such requires factual support; analytic truth is non-factual, and as such is not receptive to such support. Synthetic items of knowledge are world-oriented; analytic items of knowledge are language- or mind-oriented. The former are *vulnerable* to confutation by Nature, the latter are *not*. This approach, however, Quine argues, is unwarranted. No theory, history teaches us, is immune to revision based on fact, and any theory can be either saved from revision or subjected to one based on pragmatic considerations. Logic may be revised based on experience (as quantum logicians require), and physical experiments may be disregarded based on the overall effectiveness of our world theory. More importantly, in my view, this traditional methodology imperils our knowledge by making unwarranted assumptions on the security of certain regions of knowledge. Metaphorically, we can put it in this form:

The analytic–synthetic dichotomy creates a false line of defense against nature. Nature, the analytic–synthetic dichotomy induces us to believe, is in principle incapable of threatening the analytic zone of our knowledge. But nature might (and some say, does) find ways to encroach upon this zone. The analytic–synthetic policy of complacency in the analytic zone, careful measures for establishing the correctness of our theories in the synthetic zone, is therefore unwarranted. We do not know in advance where nature will choose to strike next, and by restricting our defenses to the synthetic front of

knowledge, are we not creating an epistemic Maginot line? As a matter of prudent strategy it is incumbent upon us to maximize the maneuverability of our cognitive resources, and just this is accomplished by NAS (Sher 1999a, pp. 504–505).

4.2 The Principle of the Normative Insignificance of the Genesis of Theories

One may try to defend the traditional methodology by first noting that some items of knowledge are incorporated into our system of knowledge by postulation, and then arguing that these items, due to their genesis, are necessarily immune to factual refutation. But we have already seen how this argument fails in at least one field, namely logic, i.e., how logical truths cannot be grounded in mere postulation (convention). Quine (1954), however, has an additional argument against truth by postulation. To say that a truth is warranted by postulation, Quine says, is to commit a genetic fallacy. In the course of building our corpus of knowledge we may resort to postulation (e.g., in order to temporarily close holes in our theories), but the act of postulation has no justificatory value. Eventually we have to justify all our theories both factually and pragmatically, and their genesis has nothing to do with this justification.

4.3 The Principle of the Non-fixity of Concepts

Traditional methodology assumes that theories change but concepts are fixed. Concepts have a fixed life of their own—fixed relationships, fixed constituents, etc.—independent of theories; and it is their independent, unchangeable features that give rise to analytic truth. Against this view Quine introduces his thesis of the non-fixity of concepts, or the inseparability of concepts from theory. Quine traces the origins of this thesis to Frege's dictum that only in a sentence does a term have a meaning, but I think they can also be traced to a model-theoretic approach to terms prevalent among twentieth-century logicians. According to this approach, the content of all (non-logical) terms is relative to models, and models themselves represent ways we could think of the world as being. Generalizing to concepts (analogous of terms) and theories of the world (analogous of models), we arrive at the view that (1) it is in the context of a theory that a concept acquires its content, and (2) change in theory (tantamount to change in models) involves change in concepts. Putnam's theory of "law cluster concepts" adds another dimension to this thesis. Scientific concepts, Putnam argues, do not always have "a single defining character or a single defining law" (Putnam 1962, p. 53); instead, they are often "constituted... by a cluster of laws which... determine [their] identity" (*ibid* p. 52). Whether a given concept will actually change in the course of history is not determined in advance; but nothing can ensure a concept's constancy. Now, while it is true that at every stage in the development of our system of knowledge some concepts must stay fixed, this is no ground for reclaiming analyticity. Different concepts may be held fixed at different stages, i.e., there are no inherently fixed concepts.

It is easy to see that each of these principles incorporates both elements of friction and elements of freedom. The first principle (or pair of principles) requires

that each and every theory in our system of knowledge be ready to respond to challenges from reality, and at the same time allows great latitude in how the system as a whole responds to such challenges. The second principle acknowledges freedom of postulation, but requires that each postulate will eventually be grounded in reality. And the third principle allows choice in creating (changing, replacing, developing) concepts, but subjects this choice to veridical standards through intrinsic connection between concepts and theories.

NAS and the related principles, however, do not, by themselves, constitute an adequate model of knowledge. NAS gives rise to a holistic model of knowledge, but holism faces two potentially fatal difficulties: lack of structure, and disconnection to reality.

4.4 Structureless Holism

Quinean holism purportedly has two faces: on the one hand Quine conceives of our total system of knowledge as the smallest unit of epistemic significance; on the other hand, he views our system of knowledge as a web of interconnected elements. There is an exegetical question concerning the first face of Quine's holism, gleaned from such claims as "I am now urging... that even in taking the statement as unit we have drawn our grid too finely. The unit of empirical significance is the whole of science" (Quine 1951, p. 42). But some critics took it literally, and it is important to distinguish the conception of holism proposed in this paper from that conception. The problem with that conception is that if the smallest unit of epistemic significance is our system of knowledge as a whole, then there is *at most one* significant unit of knowledge, namely, our system of knowledge in its entirety. But this, the critics have pointed out, is an untenable position. Thus, speaking of Quine's holism, Dummett says:

[I]f a total theory is represented as indecomposable into significant parts, then we cannot derive its significance from its internal structure, since it has none; and we have nothing else from which we may derive it (Dummett 1973a, p. 600).

And Glymour adds:

No working scientist acts as though the entire sweep of scientific theory faces the tribunal of experience as a single, undifferentiated whole (Glymour 1980, p. 3).

The point is that if the minimal unit of epistemic significance is "total theory" or "the entire sweep of scientific theory", then knowledge is an "all or nothing" affair: either we construct our system of knowledge all at once, or we do not have a system of knowledge at all; either we test it all in one fell swoop, or we do not test it at all; either we replace it in its entirety, or we do not change it at all; either we grasp (learn, communicate) it "in a single spasm of seamless cognition" (Fodor and Lepore 1992, p. 9), or we do not grasp it at all; either we ground it all at once, or we do not ground it at all. Furthermore, if to explain a total theory is, to a large extent, to significantly tie together its various significant constituents (e.g., show how its physical and mathematical constituents are related to each other), then, if it has no epistemically significant constituents, such an explanation is impossible. We may

liken our system of knowledge, under this conception, to a huge atom or blob: no inner structure, no differentiation, no interrelations, nothing to work with. The very idea of a system of knowledge becomes meaningless: “truth”, “hypothesis”, “observation”, “inference”—all notions applicable to smaller epistemic units—lose their significance; experiment—an activity that takes a specific segment of our total theory as epistemically significant and subjects it to a test—is ruled out; the connection between knowledge and rationality—a connection that has to do mostly with units smaller than our entire system, is largely undermined. And so on.

Quine’s second conception of holism is more promising: Our system of knowledge is a network of highly interconnected elements whose interconnections assume a multitude of forms. Such a system is, in principle, explainable, learnable, comprehensible, and otherwise accessible, *provided* its interconnections are manageably structured rather than chaotic. The viability of Quine’s holistic model thus requires the imposition of a manageable structure on the interconnected web of elements of knowledge.

4.5 Disconnection from Reality

By itself, a holistic network of interconnected elements (structured or unstructured) need not be connected to reality. Any consistent collection of interconnected elements would pay the bill, be it factual or fictional. To turn his system of knowledge into a system of knowledge of reality Quine has to supply it with a factual anchor or grounding, and this holism by itself cannot do.

The solution to both problems comes in the form of a new thesis, the Center-Periphery thesis, or CP. CP structures our system of knowledge as, figuratively, a circle with two distinguished zones: *center* and *periphery*. Intensionally, the center and the periphery represent two distinct dualities. On the one hand, the periphery is the area where our system of knowledge is directly connected to reality, the center—the area least connected to reality; on the other hand, the center is the center of interconnections (the locus of generality), the periphery—the region of least interconnected units (locus of particularity). Extensionally, logic, mathematics and philosophy occupy the center; experimental science the periphery. Other areas of knowledge—in particular theoretical science—occupy the intermediate zones.

In making its crucial contributions to the Quinean model, however, CP threatens to undermine NAS. There is no longer one, homogeneous domain of knowledge, governed by both veridical and pragmatic norms but, as in traditional empiricist models, a center governed by pragmatic norms and a periphery governed by veridical, experiential norms (the intermediate zone being governed by a combination of the two). It is true that statements in the center are connected to the periphery and conflicts in the periphery can be resolved in the center. But conflicts with reality occur only in the periphery, and changes in the center are purely pragmatic. A logical statement is not accepted (or rejected) because it itself agrees (or disagrees) with reality; it is accepted (or rejected) because its acceptance (rejection) helps other statements, namely statements in the periphery, to agree or overcome conflicts with reality. Logic, in this model, never lies in the periphery, nor does experiential science ever lie in the center. In this way the analytic-synthetic

bifurcation, or its epistemic analog, re-enters through the back door, so to speak, in the guise of a center-periphery duality.

4.6 Solution to the Conflict: From a Static and Absolutist CP to a Dynamic and Contextual CP

Some commentators, e.g., Dummett (1973a, b), react to the fundamental tension between CP and NAS by eliminating the latter. And eliminating the former is, of course, also an option. Our analysis suggests, however, that NAS offers a better strategy than the traditional methodology for achieving a good measure of friction and freedom, while something like CP is needed to complement it. My solution to the conflict between CP and NAS consists, therefore, in a reconfiguration of CP. Center and periphery are reconfigured as job descriptions rather than permanent locations of statements and theories, and in principle both logic and physics have a job in the periphery and a job in the center. Physics' "peripheral" job is to square our system of knowledge with the *physical* (material, empirical) features of reality, logic's "peripheral" job is to square our system with the *formal* features of reality. During periods of changes in physics logic provides the glue that holds our system together (i.e., logic is located in the center), during periods of changes in logic, (experimental) physics provides the constant element in our system (by moving to the center). When veridical norms apply to logic, logic is in the periphery; when pragmatic norms apply to physics, physics is in the center. I will elaborate on the way logic is located in the periphery (logic is grounded in reality) in the next section, but structurally we can see that the model is dynamic and contextual. Metaphorically, disciplines move freely from the center to the periphery and the other way around, and this movement is determined by task and context. During periods of what Kuhn (1970) calls "normal" science experimental physics lies in the periphery, logic and mathematics in the center, and theoretical physics in the intermediary zone. But in the course of a "scientific revolution" their positions may change. In the course of Einstein's revolution geometry moved to the periphery; in the course of the quantum revolution some have considered the possibility of logic moving to the periphery. When our focus is on discovery in physics, logic is held "fixed" in the background (center); when our focus is on (factual) discovery in logic or meta-logic, physics is held fixed in the background (as something that logical discovery should not, or should not lightly, interfere with), while logic, or meta-logic, confronts the facts in the periphery (For an example of logic, or meta-logic, confronting the "facts", think of Gödel's discovery of the incompleteness of 1st-order arithmetic). In this way, the model provides a link between our entire system of knowledge and reality, and the link it provides is holistic (in the structural sense).

5 Logic's Embeddedness in the World

One of the distinctive characteristics of the neo-Quinean model is its treatment of logic. By placing logic (in certain respects and during certain periods in the development of our system of knowledge) in the periphery, the new model subjects

it to the first requirement of friction, namely, groundedness in reality. But *need* logic, from its own point of view, be grounded in reality? *Can* it be so grounded? And *how* is it to be grounded? In the present section I will offer a general answer to these questions, and in the next (and last) section I will respond to some objections and offer a few clarifications.

The two main reasons for grounding logic in the world, in my view, are:

1. Logical theory (like physical theory) has to “work” in the world.
2. Logical theory (like any other theory) is “immanent” or “world-oriented”.

5.1 Logical Theory has to “Work” in the World

It is a straightforward observation that in the same way that using a defective physical principle can cause a system (dependent on it) to malfunction, so using a defective logical principle can result in a system malfunctioning. This is not to say that we have no latitude in selecting our logical (or, for that matter, our physical) laws. But there is a very real sense in which logical theory, like physical theory, either works or doesn’t work in the world.

Take, for example, the system of drag and lift in an airplane. It is clear that adopting a flawed lift principle, e.g., “*to achieve a lift effect, set the flaps at a large downward angle*”, could cause an airplane to crash. But adopting a flawed identity principle, e.g., “*If a is not identical to b and a has property P, then b does not have property P*”, can also cause an airplane to crash (for example by skewing the calculations used to determine whether a certain action leads to a lift effect). In both cases we affirm a non-existent correlation between properties of objects, and both errors can lead to disaster. To prevent an error in our physical laws we ground them in reality, and to prevent an error in our logical laws we need to ground them, too, in reality. The grounding of a law of identity may require different methods from the grounding of a law of lift, but if, in reality, *non-identity* is *not* correlated with disjointedness of properties, then (in the absence of appropriate correcting measures) our logical theory should not contain a law that says it does.

5.2 Logical Theory is “Immanent” or “World-Oriented”

A more theoretical argument for logic’s embeddedness in reality is based on the observation, due to Quine, that “[l]ogical theory... is... world-oriented rather than language-oriented; and the truth predicate makes it so” (Quine 1970, p. 97). The idea is that truth depends on the way the world is, and since logic purports to produce *true* logical laws, logic is world-oriented. One way to explain this idea is based on what I have elsewhere (Sher 2004) called the “Immanent” conception of truth (I do not suggest that this is the conception Quine actually had in mind, but this is a conception that explains and justifies his claim). Truth, according to this conception, lies at the juncture of three basic modes of human thought:

1. the immanent mode (lower-case ‘i’),
2. the transcendent mode, and

3. the normative or critical mode.

The *immanent mode* is the mode of attributing properties and relations to objects or saying that things are one way or another; the *transcendent mode* is the mode of transcending a given thought in order to *reflect* upon it; and the *normative* or *critical mode* is the mode of critically evaluating whether a given thought satisfies a given standard or desideratum. *Truth is a standard for a positive answer to the critical transcendent question “Is it so as thought t says it is?”, applicable to any immanent thought t . If t is the immanent thought that object o has the physical property P , then t is true if (in the world) o has the physical property P ; if t is the immanent thought that the collection C of objects satisfying a certain condition p has the logical property of being non-empty (i.e., the existential-quantifier property, \exists), then t is true if (in the world) C is not empty, and so on. The precise conditions for having a physical property P and having a logical property L (e.g., \exists) may vary, but in both cases the truth of the relevant thought depends on the way the world is.*

Thus, take the logical theorem-schema


$$(Cont) \sim \exists x(Px \& \sim Px),$$

where ‘ P ’ is a schematic one-place predicate. *(Cont)* says that there is no property P and object o such that o has both the property P and its complement, not- P . This is a law concerning the formal behavior of objects and properties in the world, a law that, as we shall see below, satisfies an especially strong invariance condition which distinguishes it from other, non-logical, laws.

Indeed, not only is logical truth grounded in reality, *logical consequence*, too, is so grounded. To see how logical consequence is grounded in reality, I will proceed in three steps:

5.3 Consequence Relations in General

Consequence relations in general are relations of transmission of truth from sentences to sentences: Let S_2 be a consequence of S_1 , and let S_1 be true. Then truth is transmitted from S_1 to S_2 . Now, for S_1 to be true the world has to be as S_1 says it is—let us say that situation S_1 has to hold. Also, if the truth of S_1 is to be transmitted to S_2 , the world has to be as S_2 says—say, situation S_2 has to hold. Now, suppose the world is such that S_1 holds but S_2 doesn’t. Then, S_2 is not a (valid) consequence of S_1 . *In order for a consequence relation to work properly (i.e., transmit truth from sentences to sentences) the world has to cooperate.* We may say that the transmission of truth from sentence S_1 to sentence S_2 is grounded in a certain connection between the situation S_1 and the situation S_2 . Graphically:



The truth of S_1 is transmitted to S_2 if, when, and because S_1 is appropriately connected to S_2 : S_1 's holding in the world is sufficient for S_2 to hold as well.

5.4 Strong Consequence Relations

Now, suppose the consequence relation in question is “strong”, i.e., it satisfies some necessity standard (physical necessity, metaphysical necessity, logical necessity, etc.). Then the transmission of truth from S_1 to S_2 is based on an appropriate *necessary* connection between S_1 and S_2 . Graphically:

$$\begin{array}{ccc}
 \text{Language:} & T(S_1) \text{ ----- Nec -----} & \rightarrow T(S_2) \\
 & \Downarrow & \Uparrow \\
 \text{World:} & S_1 \text{ === Nec =====} & \rightarrow S_2 .
 \end{array}$$

5.5 Logical Consequence

To understand how logical consequence is grounded in the world, then, is to understand what kind of necessary connections must hold between S_1 and S_2 if S_2 is to be a logical consequence of S_1 . I will offer an elaborate account of these connections in Sect. 6 below, but briefly, my view is that logical consequence is grounded in *formal laws* governing (and relating) structures of objects (properties, relations), where the idea of a “formal law” is very similar to that of “mathematical law” in structuralist accounts of mathematics (e.g., Resnik 1997; Shapiro 1997). On this view, what distinguishes formal laws from other kinds of law is the strong invariance conditions they satisfy. While physical laws distinguish between isomorphic structures whose objects differ (e.g., between structures whose objects are physical and isomorphic structures whose objects are mathematical), formal laws do not distinguish between isomorphic structures no matter how different their objects are from one another. Logic is characterized by its invariance under isomorphisms of structures, and it is this property that distinguishes its operators, its truths, and its consequences, from the operators, truths, and consequences of other disciplines.

According to this view, logic is grounded in a specific facet of reality, namely, the formal facet, and to satisfy the friction requirement logic has to get this facet right. If, to return to our earlier example, logic wrongly assumes that non-identical objects must differ in all their properties, its rules of inference for identity statements will be skewed. This would result, practically, in an increased likelihood of airplanes crushing, cars stalling, criminals going free, innocents being incarcerated, etc., and theoretically, in a violation of the first friction requirement, namely, groundedness of knowledge in reality, or more specifically, groundedness of logic in the formal laws governing reality.

6 Objections and Clarifications

The above discussion leaves many issues to be further elaborated. I will conclude by offering two sets of clarifications in response to two objections:²

6.1 First Objection

The neo-Quinean model of knowledge construes logic as “grounded in reality” and in principle revisable as our “web” of knowledge is confronted by evidence. One standard objection is: If our system of knowledge is “a web of interconnected elements”, what constitutes the interconnections? If they are *logical* connections, then how can logic itself be subject to revisions? When it’s revised, won’t the interconnections be revised as well? And if so, how can the web remain intelligible after revision?

The key to overcoming this objection lies in the *holistic* and *dynamic* structure of the neo-Quinean model. The network of interconnections among different constituents of our system of knowledge consists of connections of numerous kinds. These include, in addition to logical connections, also non-logical principles of reasoning (e.g., inductive principles of inference), general methodological principles (e.g., the principle of unity), pragmatic principles (e.g., economy), psychological principles (e.g., preference of the familiar), physical principles (e.g., the laws of gravity cannot, in the absence of a compelling reason, be violated by revisions outside the theory of gravity), everyday items of knowledge (e.g., “I have two hands”), etc. Generally, it is the principles held “fixed” during a particular revision that play the dominant role in connecting the different parts of our system of knowledge together during that revision. The elements we hold fixed during revision *r* are what “glues” our system together during *r*. The dynamic nature of the model is reflected in the fact that while some elements must be held fixed during every revision, not the same elements have to be held fixed during all revisions. The fixed-unfixed division may change from revision to revision. Kuhn (1970), for example, distinguished two types of revision in empirical science: revision in “normal” science and “scientific revolution”. During revisions of the first type our logical, mathematical, and conceptual principles are held fixed; during revisions of the second type some of these principles become unfixed, and simple observations, common-sense principles, scientific principles not undergoing change, as well as logical, mathematical, and conceptual principles not involved in the change, take their place as fixed elements. In the course of a revision in logic, non-logical principles, as well as logical principles not undergoing revision, are held fixed (partaking in the gluing of our system), while those principles undergoing change are temporarily excluded. When the revision is “factual” (rather than, e.g., pragmatic), this is a manifestation of the movement from *center* to *periphery* characteristic of the neo-Quinean model. During (factual) revision in physics logic lies in the center, physics in the periphery; during (factual) revision in logic, everyday physics occupies the center, logic the periphery. Furthermore: during

² I thank an anonymous reviewer for inviting these clarifications. The two objections are taken almost verbatim from his/her comments.

(factual) revision of special principles in logic, those logical principles occupy the periphery; other logical principles occupy the center, and vice versa.

Furthermore, due to the multi-dimensional nature of the model, even principles undergoing change can play some role in connecting our system together, i.e., can lie in the center on one level (at one time and/or from one perspective), in the periphery on another level (at another time and/or from another perspective). For example, suppose we want to replace our currently dominant bivalent logic by a three-value logic. One way to carry out this revision is to resort to the familiar distinction between *object-theory* and *meta-theory* (*meta-theory* and *meta-meta-theory*, and so on). Suppose, at the outset logic lies in the center. We may proceed as follows. First, we make two-copies of our bivalent logic, leaving one in the center and moving the other to the periphery. Then, using the logic in the center (as well as other resources, e.g., various mathematical resources) as our meta-theory, we turn the “peripheral” logic from a bivalent to a trivalent logic. Finally, we move the revised logic from the periphery back to the center, where it replaces the old, bivalent, logic. Henceforth (i.e., until the next revision in logic), we use our newly revised, trivalent, logic to interlink our system of knowledge (In this scenario, we use the old logic temporarily, as a Wittgensteinian ladder).

In practice, revision in logic will, in all likelihood, involve all the above means of keeping our system connected. Take a dramatic case of revision in twentieth-century logic, namely Russell’s revision of Frege’s logic in light of his discovery of an inner contradiction in it. What background logic did Russell rely on in discovering the paradox? What logic did he use in looking for ways to overcome the paradox, in replacing Frege’s logic by his own type-theoretic logic? It is unreasonable to assume that the logic Russell relied on in discovering the paradox was itself type-theoretic logic, the logic he constructed *subsequent* to this discovery. Nor is it reasonable to assume that he made do with some pre-Fregean (e.g., Aristotelian) background logic. In all likelihood Russell used some combination of Fregean logic, Cantorian set-theory, and various other mathematical and common-sense principles, as his background theory, shifting back and forth between holding Fregean logic fixed and unfixing it, suspending some of its principles but not others. The result was the discovery of “Russell’s” paradox, the devising of a method to avoid it, and the construction (based on this method) of a new logical theory, the “theory of types”. This logic replaced Frege’s logic, at least for a while, in its role of holding our system together and ensuring its intelligibility (a role that, in our holistic conception, it shared with other elements of the system).³

To sum up:

1. The network of interconnections contains numerous types of connections, not just logical connections.
2. The dominant interconnections during a given revision are those held fixed during that revision. Since revision in logic is often partial, it allows some logical laws to be held fixed.

³ Today, Russell’s system is no longer the prevalent logical system. This change could be explained in similar, though less dramatic, terms, with a greater role possibly assigned to pragmatic considerations than in Russell’s revision of Frege’s logic.

3. Due to its special features, the neo-Quinean model allows even revised laws, including revised logical laws, to play a limited role in interconnecting our system of knowledge:
 - a. The neo-Quinean model is a holistic model, and holistic models can sustain some circularity.
 - b. The neo-Quinean model is a dynamic and perspectival model. As such it allows considerable flexibility in:
 - i. maneuvering the fixity and non-fixity of logical principles,
 - ii. moving logical elements back and forth between center and periphery.

Let us now turn to the second objection, or request for further clarifications.

6.2 Second Objection

In the neo-Quinean model it is not the “physical” reality that provides friction for logic, but rather “formal” reality. But what is “formal” reality? In Sect. 5 it is claimed that since logic concerns the true it is “world-oriented” rather than “language-oriented”, and the “world” toward which logic is “oriented” is the “formal world”, consisting of “the formal behavior of objects and properties in the world”. But what does that mean? And why should we accept that there are such objects and properties in the world? Is the world these objects and properties belong to the physical or formal world? Furthermore, when logic is said to be characterized by its invariance under isomorphisms of structures, which isomorphisms are these? And which structures? What counts as a structure? Going back to the main question, we can put it in provocative terms as follows: For logic to meet the epistemic friction test, it needs a reality to provide friction. What kind of reality can do this job? Not physical reality. So do we simply call the needed reality “formal reality” and say that it is whatever will give logic the needed epistemic friction? But wouldn’t this be another case of “theft over honest toil”? It might help illuminate the explication of formal reality to address the question of whether mathematical reality is grounded in formal reality, or in yet some other kind of reality.

This objection applies to the present paper one of the standards this paper applies to all theories and disciplines: substantiveness. As a piece of philosophy, this paper is required to provide a substantive explanation of one of its main themes, the formality of logic. I have given extensive (and, I hope, substantive) explanations of this issue elsewhere (Sher 1991, 1996, 1999b, 2001, 2008, 2009), and I cannot repeat all these explanations here. But I can and will offer an abbreviated explanation of this issue, an explanation that complements the general outline presented in Sect. 5.

First, I would like to make clear that reality, on the conception proposed here, is *one*. There are no two (or more) separate realities, a platonic reality of formal entities, and a physical reality of entities accessible to sense perception. There is one, multi-faceted reality. This is integral to the present version of holism. There is one, rich and heterogeneous reality, a reality whose different, yet interconnected, facets are studied by different branches of our system of knowledge. Logic is

grounded in, and is constrained by, the formal facets of this one reality, and to say that logic is grounded in formal reality is an abbreviated way of saying that logic is grounded in the formal features, structure, and laws of reality or the world (I use “reality” and “world” as synonyms).

Before proceeding to explain the claim that reality has formal features, let me note that for the purpose of the present discussion reality is thought of in the familiar terms of a domain of objects with properties and relations, where these properties and relations themselves have properties and relations, and so on. I.e., reality constitutes an intuitive hierarchy of objects and properties (relations, functions) of various levels, the first few of which we sufficiently understand (on an intuitive level) to proceed with further explanation. In this hierarchy, individual objects are said to be of level 0, properties and relations of individual objects—level 1, properties and relations of first-level properties and relations—level 2, and so on. Individuals are objects that in the present context we treat as atomic. Properties and relations we treat, for the purpose of the present discussion, as standing on their own, leaving open the question of whether ultimately they are, or are not, objects (Our working assumption is that within certain limits, we can adjust our account to various conceptions of ontological categories).

Now, the present view is that objects and properties in the world, regardless of their ontological status (physical, non-physical), have formal properties and stand in formal relations, and logic is grounded in the laws governing those properties and relations. What leads us to think that objects in the world have formal features? We are led to this view by the observation that no matter what individuals and properties we start with, we find that they have properties that are intuitively (and, as we shall see below, also theoretically) classified as formal. Even if we start with just physical individuals and physical properties (relations), it is hard to deny that they have (some) formal properties and stand in (some) formal relations. To see how we arrive at this observation, take a domain of individuals limited to physical objects, and suppose that it has, among other things, flowers and (flower-) petals. It is hard to deny, as Gödel liked to point out, that it is as much a property of some flowers—a *formal* property of some flowers—they have five petals as that they have, say, red petals. Or, to use a Fregean analysis, it is hard to deny that the property “ x is a petal of a ”, for some flower a (with five petals) has the formal property “ X holds of five things” (or “the cardinality of X is five”). Likewise, it is hard to deny that some flowers stand to other flowers in the formal relation of having *half as many* petals. Similarly, each flower has the formal property of “ x is *identical to itself*”, and each stands in the formal relation “ x is *different from*⁴ y ” to all other individuals in the domain. Also: each flower stands in the *complement*, *intersection*, and *union* of some properties, i.e., each has properties that are obtained by the formal operations of complement, intersection, and union from other properties. Likewise, some relations between flowers have the formal property of being *reflexive*, others of being *symmetric* and/or *transitive*. For example the relation of being larger than is transitive. These are just a few formal properties and relations that individuals and properties in the world, physical or non-physical, commonly possess (stand in).

⁴ not identical to.

What are the distinctive characteristics of formal properties and relations? My proposal is that formal properties are distinguished from other properties by their strong invariance. Specifically, formal properties are invariant under all isomorphisms of structures. What does this mean and what is the motivation for this characterization? I have discussed the intuitive motivation for this characterization, its historical roots, and its relation to the structuralist conception of mathematics in Sher (*op. cit.*), and I will not be able to repeat all these here. But I will briefly explain what it means. The basic idea is that formal properties and relations do not take into account which individuals populate a given domain. For example, they do not distinguish physical individuals from biological individuals, biological individuals from psychological individuals, psychological individuals from mathematical individuals, and so on. If something has a formal property in a structure whose domain consists of physical individuals, then all its images in all isomorphic structures have this property, regardless of whether their domains consist of physical, biological, psychological, ..., or mathematical individuals.

As an example, take a second-level cardinality property C , a property like “exactly n ”, or “finitely many”, or “half (the objects in the domain)”. This is a second-level property, and its arguments are first-level properties, whose extensions, in a given domain A of individuals, are subsets B of A . Now, this property has the distinctive feature that it does not distinguish between isomorphic structures of the form $\langle A, B \rangle$, where A and B are as above. What this means is, roughly, that if B satisfies C in the structure $\langle A, B \rangle$, then for any structure $\langle A', B' \rangle$ which is isomorphic to $\langle A, B \rangle$, B' satisfies C in it. Similarly, if B does not satisfy C in the structure $\langle A, B \rangle$, then for any structure $\langle A', B' \rangle$ isomorphic to it, B' does not satisfy C in it. Going back to our examples, it is easy to verify that indeed, if the structures $\langle A, B \rangle$ and $\langle A', B' \rangle$ are isomorphic, then B has exactly n elements (in A) iff B' has exactly n elements (in A'), B has finitely many elements (in A) iff B' has finitely many elements (in A'), B contains exactly half the elements in A iff B' contains exactly half the elements in A' , and so on. Properties like exactly n , finitely many, and half are preserved under all isomorphisms of structures (of the relevant form). Such properties take into account only the structural features of given arguments, but not their other features. They do not distinguish between situations that are structurally alike. We call such properties, *formal* properties. Formal properties do not have to be of the second level. Take, for example, the identity relation, $=$, of level 1. Given any pair of individuals, a and b , in any domain A , if $\langle A, a, b \rangle$ is isomorphic to $\langle A', a', b' \rangle$ then a is identical to b (in A) iff a' is identical to b' (in A') (Here the relativity to A and A' is redundant, but in other cases it may not be). In contrast, physical and biological properties are not preserved under isomorphisms. If P is a property of physical objects but not of numbers, then it is not preserved under isomorphisms that hold between physical and, say, mathematical structures (i.e., structures whose domains consist of physical objects and structures whose domains consist of numbers). Similarly, if P is a property of animate physical objects but not of inanimate physical objects, it is not preserved under isomorphisms holding between structures of animate and inanimate objects. The relation “is a father of” holds of Barack Obama (b) and Malia Obama (m) in the domain $\{b, m\}$, but it does not hold of New York State (n) and California (c) in the domain $\{n, c\}$, or

of the numbers 1 and 2 in the domain $\{1,2\}$, (or, even, of m and b , in that order, in the domain $\{b,m\}$), in spite of the fact that the structures $\langle\{b,m\}, b,m\rangle$, $\langle\{n,c\}, n,c\rangle$, and $\langle\{1,2\}, 1,2\rangle$ (as well as $\langle\{b,m\}, m,b\rangle$) are all isomorphic to each other. It is in this sense that the properties “exactly n ”, “finitely many”, “half”, and the relation “=”, are all *formal*, whereas the relation “is a father of” and other physical and biological relations are not.⁵

Our next topic is formal laws. How do we arrive at such laws?—We reason that if objects, properties, and relations in the world have formal features, then these features, like, e.g., their physical features, are potentially governed by laws and regularities. What discipline studies the laws governing these features? The most natural candidate for such a theory is pure mathematics, or some branches thereof. Not all pure mathematics need to fall under this category—after all, mathematics is a broad discipline, with many goals and interests. But some of its parts—e.g., arithmetic, set theory, and/or some other theories—do.

How do we explain the fact that these parts of mathematics have an ontology that goes far beyond the physical ontology we have started with?—We reason as follows: Formal laws, like all other laws, have a certain modal force, a modal force that goes beyond their actual applications. Indeed, due to the strong invariance characteristic of formal properties, the laws governing them have an especially strong modal force.⁶ To account for the full force of these laws (as well as to satisfy various methodological constraints and desiderata) the relevant mathematical theories are formulated in terms of an enormous ontology. Whether we eventually treat this ontology as part of the ultimate “furnishing” of our world is immaterial (from the present perspective). What matters is that real objects (real properties) have formal features, and the laws governing these features are real.

Why and how does logic enter into the picture?—Logic enters into the picture in two ways: a practical way, and a theoretical way. Theoretically, logical inference is an interesting subject matter in its own right. Being a logically-valid inference is a property of some n -tuples of sentences and not others, a property with unique and interesting features, hence something we wish to know and understand. It is an abstract property, part of the most abstract facets of reality, yet something we are capable of studying and understanding. Practically, logical inference is an effective means for extending knowledge. Due to our cognitive limitations and relatively short life-span, we cannot hope to establish all our knowledge directly but have to resort to *inference* to obtain a considerable portion of our knowledge. Inference in general is transmission of truth, and transmission of truth, in the case of expansion of knowledge (where premises are in fact true), is based on relations that hold between the situations that make the premises true and a situation that (in successful inference) makes the conclusion true. A method of inference that is both highly reliable, universal (topic neutral), and stable (largely insensitive to changes in knowledge in most specific areas), is especially advantageous for knowledge, and such a method is supplied by logic. Logic provides a method of inference based on knowledge of basic laws governing the formal “behavior” of objects and properties

⁵ For a more general and precise account see, e.g., Sher (1991, 1996, 1999b).

⁶ I discuss this point in detail in Sher (1999b, 2008, 2009).

in all areas of knowledge. Given that formal features of objects are constantly referred to in all discourse—one cannot talk about anything without saying that certain objects are in the *complement* or *intersection* of certain properties, that certain properties are *non-empty*, or *universal*, or have κ objects falling under them, etc.—we can use our knowledge of these features to develop a wholesale method for expanding our knowledge. Logic, on this conception, utilizes our knowledge of the formal behavior of objects (given by basic mathematics) to formulate rules of inference that sanction our movement from what we know to what (prior to this movement) we did not know. This it does by holding “fixed” certain constants referring to formal properties and functions (“is” of identity, “all”, “some”, “not”, “and”, “or”, and possibly others), and developing a method of inference based on the laws governing the denotations of these constants. Due to the strong modal force of formal laws, the logical method provides an especially strong guarantee of transmission of truth and, based on it, expansion of knowledge (For further elaborations see Sher 1991, 1996, 1999b, 2001, 2008, 2009).

Acknowledgments I would like to thank the Lady Davis Foundation for a fellowship supporting the writing of this paper, the Hebrew University in Jerusalem for its hospitality, and Carl Posy for sponsoring my fellowship. I have greatly benefitted from discussions with Carl Posy, Oron Shagrir, Mark Steiner, Ygal Kwart, and many other colleagues in Israel. An earlier version of the paper was presented to the Rationality Group at the Hebrew University, the department of philosophy at Ben Gurion University, and the philosophy of logic workshop at the University of North Carolina Chapel Hill, and I am very grateful to the audiences in these places. Special thanks go to John MacFarlane, my commentator at the Chapel Hill workshop, Peter Sher, and an anonymous reviewer of this journal whose comments led to significant improvements.

Open Access This article is distributed under the terms of the Creative Commons Attribution Non-commercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Alston, W. (1976a). Has foundationalism been refuted? *Philosophical Studies*, 29, 287–305.
- Alston, W. (1976b). Two types of foundationalism. *Journal of Philosophy*, 73, 165–185.
- Alston, W. (1983). What’s wrong with immediate knowledge. *Synthese*, 55, 73–95.
- Belnap, N. D. (1962). Tonk, plonk and plink. *Analysis*, 22, 130–134.
- Dummett, M. (1973a). *Frege: Philosophy of language*. New York: Harper & Row.
- Dummett, M. (1973b). The significance of Quine’s indeterminacy thesis. *Truth and Other Enigmas*. Cambridge: Harvard University Press, 1978: 375–416.
- Fodor, J., & Lepore, E. (1992). *Holism*. Oxford: Blackwell.
- Glymour, C. (1980). *Theory & evidence*. Princeton: Princeton University Press.
- Gödel, K. (1953–1959). Is mathematics syntax of language? *Collected Works III*. Oxford University Press, 1995: 334–362.
- Kant, I. (1781/1787). *Critique of Pure Reason* (1st and 2nd eds.), Tr. Kemp Smith. Macmillan, 1933.
- Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.). Chicago: University of Chicago Press.
- McDowell, J. (1994). *Mind and world*. Cambridge: Harvard University Press.
- Prior, A. N. (1960). The runaway inference-ticket. *Analysis*, 21, 38–39.
- Putnam, H. (1962). The analytic and the synthetic. *Mind, Language and Reality*. Cambridge: Cambridge University Press, 1975: 33–69.

- Quine, W. V. (1935). Truth by convention. *The Ways of Paradox and Other Essays* (Revised ed.). Cambridge: Harvard University Press, 1976: 77–106.
- Quine, W. V. (1951). Two dogmas of empiricism. *From a Logical Point of View* (2nd ed.). Cambridge: Harvard University Press, 1980: 20–46.
- Quine, W. V. (1954). Carnap and logical truth. *The Philosophy of Rudolf Carnap*. Schilpp (ed.). Illinois: Open Court, 1963: 385–406.
- Quine, W. V. (1970). *Philosophy of logic*. Cambridge: Harvard University Press, 1986.
- Resnik, M. D. (1997). *Mathematics as a science of patterns*. Oxford: Oxford University Press.
- Russell, B. (1919). *Introduction to mathematical philosophy*. New York: Simon & Schuster. (1971).
- Shapiro, S. (1997). *Philosophy of mathematics: Structure & ontology*. Oxford: Oxford University Press.
- Sher, G. (1991). *The bounds of logic*. Cambridge: MIT Press.
- Sher, G. (1996). Did Tarski commit 'Tarski's fallacy'? *Journal of Symbolic Logic*, 61, 653–686.
- Sher, G. (1999a). Is there a place for philosophy in Quine's theory? *Journal of Philosophy*, 96, 491–524.
- Sher, G. (1999b). Is logic a theory of the obvious? *European Review of Philosophy*, 4, 207–238.
- Sher, G. (2001). The formal-structural view of logical consequence. *Philosophical Review*, 110, 241–261.
- Sher, G. (2004). In search of a substantive theory of truth. *Journal of Philosophy*, 101, 5–36.
- Sher, G. (2008). Tarski's thesis. In D. Patterson (Ed.), *New essays on Tarski and philosophy* (pp. 300–339). Oxford: Oxford University Press.
- Sher, G. (2009). Is logic in the mind or in the world? *Synthese* (forthcoming).