

MODELING AND NORMATIVITY: How much revisionism can we tolerate?

Stewart Shapiro

The Ohio State University and The University of St. Andrews

Abstract

The paper concerns the relationship between mathematical logic and its supposed subject matter: correct reasoning. I suggest that a formal logic is a mathematical model of reasoning, in much the same sense, as a system of point masses is a model of moving objects. There are gaps between the model and what it is a model of. The paper explores the possibility of using a logic, so understood, to motivate revisions in reasoning. The main case studies are classical logic, relevance logic, and intuitionism. The paper is an extended commentary on John Corcoran's classic paper "Gaps between logical theory and mathematical practice" (*The methodological unity of science*, edited by M. Bunge, Dordrecht, Holland, D. Reidel, 23-50).

Key words: logic, model, reasoning.

My interest in mathematical logic began, when I attended an National Science Foundation Summer Program in mathematics for high school students, held at Ohio State from 1967-69. The Director, Arnold Ross, taught us a slogan: "think deeply about simple things". Good marching orders for the budding mathematician; crucial for the philosopher. A corollary to this is that we should return again and again to the most basic questions. This legacy from Socrates is one of the most exhilarating aspects of philosophy.

The underlying theme of John Corcoran's [1973] classic article "Gaps between logical theory and mathematical practice" invites us to follow this path. The issue is the purpose of mathematical logic. What is a formal logic supposed to illuminate? How does it accomplish this illumination? Typically, a logic consists of a formal language, a deductive system, and/or a model-theoretic semantics. We know what those things are, but what do they do? What are they for?

A typical elementary logic course begins with the instructor saying that logic is the study of correct reasoning. This opening gambit puts logic within epistemology, since reasoning has something to do with knowledge. Pretty soon, however, the typical student finds himself working with mathematically

defined formal languages, formal deductive systems, and set-theoretic model theory. Our question is this: what do the chicken scratches produced after the first day of the term have to do with correct reasoning, the advertised goal of our subject?

A formal language is a recursively defined collection of strings on a fixed alphabet. This mathematical object somehow corresponds—or is supposed to correspond—to whatever the medium of reasoning is. Well, what is the medium of correct reasoning? Is it a natural language, the realm of propositions, the language of thought? And what do formulas in a formal language have to do with this medium? Our mantra is that validity is a matter of form, and that formulas reflect the logical forms of sentences or propositions. What does this mean? Does each sentence or proposition have a unique form, or can a sentence or proposition have several forms at once?

Corcoran's proposal is that we think of a formal language as a mathematical model of a chunk of natural language (or whatever the medium of reasoning is), in the same sense that a system of point masses is a model of interacting physical objects. As Corcoran sees things, a deductive system is a model of the process of showing an argument to be valid. Formal deductions correspond to discourses in which a person comes to know that a given argument is valid—that its premises guarantee its conclusion. In like manner, a model-theoretic semantics models the ways that someone can come to know that an argument is *not* valid. Intuitively, to establish that an argument in natural language is invalid, one gives another argument in the same form with true premises and a false conclusion. Since formulas in the formal language correspond to logical forms, the various interpretations of the formal language provide a stock, or the stock, of arguments that share a given form. The logic-as-model approach was developed in detail in the recent dissertation of my student Roy Cook.

The [1973] Corcoran paper concerns mismatches between mathematical logic and mathematical deductive reasoning. It is the nature of the enterprise that there is a gap between a model and what it is a model of. Some gaps are desirable. A good model is simple and idealized, and so easier to study than the phenomena being modeled. Consider, for example, the point mass model of physical objects. It would not be appropriate to complain about the gap concerning the infinite density of point masses. This idealization is part of the underlying purpose of this model.

There are two rather different reasons to focus attention on gaps between model and modeled. A typical use of a mathematical model is to infer something or explain something about the modeled phenomena. Suppose that we notice a fact about the model, and want to draw a conclusion about the underlying modeled phenomena. We have to be reasonably sure that the inference does not turn on a gap, a feature of the model that incorrectly matches up with the modeled phenomena. For example, we can use a point

mass model to correctly predict the orbits of planets, since the distances and sizes are reasonable, and friction is negligible. But we should not use the point mass model to predict how long it will take a feather to fall to the earth.

A second reason to pay attention to gaps is so that we can construct better, more realistic models. This is what happens when physicists develop models of physical systems whose objects have finite volume and which friction. Notice, however, that for some purposes the more “realistic” models are not as useful. The more accurate models are more cumbersome, less tractable. Of course, the fact that simple models are tractable is itself a gap. Reality is rarely tractable. If it were, we would not need the models.

Getting back to the case at hand, can we use mathematical logic to *improve* reasoning? One would think so since, after all, logic is the study of reasoning. A deductive system is supposed to model inference patterns implicit in the practice of an individual or a community. Suppose we come up with a deductive system and realize that in certain instances, a subject or subjects systematically fail to reason in accordance with it. We might think we have discovered a mistake in the practice. But why call this a mistake, and not a gap between theory and practice?

Let us start with a toy example. It is a commonplace that in studying deduction, we introduce the distinction between bound variables and free variables (or parameters). We uncover the exact conditions for certain inferences, such as universal introduction and existential elimination. Then we are in position to be critical of a mathematician who violates these conditions (which they do sometimes).

On the other hand, suppose that a mathematician violates the conditions on free and bound variables, but still insists that her practice contains no mistakes. She claims that she has made no false conclusions, and we have to admit that we cannot find any in her practice. The logician points out that by ignoring the rules for parameters or free variables which the mathematician seems to be doing, she might reason from true premises to false conclusion. The mathematician retorts that the pedantic logician has *mis-described* the practice, and that she is not following the flawed rule he has attributed to her, nor is she incorrectly following what the logician says is the correct rule. The logician then asks her to say what rule she is following. She declines the invitation, claiming to have no patience for logic, but she insists that she knows what she is doing —suggesting that we butt out. The logician might feel like Berkeley complaining about the infinitesimal calculus. And the mathematician might feel like the analysts who felt safe in ignoring Berkeley’s critique, for a time anyway.

I presume that most of us sympathize with the logician in the above dispute. But the perspective of logic-as-model-building makes the logician’s critique problematic. Why think that the *practice* is in error? Would it not be more natural to think that it is our model is not correct, as the imaginary

mathematician claims? In general, when do we blame a supposed gap on the practice, favoring the model over the practice, and when do we blame the gap on the model?

Notice that in other cases where mathematical models are used, it is patently absurd to use the model to criticize what it is supposed to be a model of. Suppose, for example, that a biologist concocts a model of the growth of certain bacteria in Lake Erie. He then sees that in some places, the lake consistently fails to fit the model. Can he coherently maintain that his model is correct, and be critical of the lake, or the bacteria in it, for failing to conform to the correct model?

Some perspective is provided by the first round of John Burgess's [1992] attack on revisionism in his "Proofs about proofs: a defense of classical logic". Burgess compares mathematical logic to mathematical physics and mathematical economics, a simile much in the spirit of Corcoran's "Gaps between logical theory and mathematical practice". As we noted above, Burgess notes that gaps are to be expected:

...like mathematical physics, [logic] deals with an idealization of reality. Its artificial languages are conspicuously simpler in grammar than natural languages, for example. No more than mathematical physics is mathematical logic to be condemned for being unrealistic or idealized. (p. 12)

He makes a distinction between different approaches to logic:

Whenever a community has a practice, the project of developing a theory of it suggests itself. When the practice is one of evaluation, a distinction is made between descriptive and prescriptive theories thereof. The former aims to describe explicitly what the community's implicit standards have been: the theory is itself evaluated by how well it agrees with the facts of the community's practice. The latter presumes to prescribe what the community's standards ought to be: the community's practice is evaluated by how well it agrees with the norms of the theory. (p. 12)

Things are potentially confusing here, since our subject—logic—is itself inherently normative. We do not deal with human reasoning as such, but with *correct* reasoning. As Burgess notes, the target practice is one of *evaluation*. On the first day of a logic class, the typical instructor notes the difference between logic (or epistemology) and cognitive, empirical psychology. Supposedly, psychology is (only) interested in human thought processes, with little or no concern for when those processes are correct. Logic is concerned with how reasoning is supposed to go, and so it is normative from the start. So Burgess proposes a distinction between a descriptive and a normative theory of a normative enterprise.

The absurdity of the above imaginary biologist's attitude towards the bacteria in Lake Erie indicates that one is hard put to come up with a normative mathematical biology. Similarly, I guess we can dismiss a

normative approach to mathematical physics. I am not so sure about Burgess's other example, mathematical economics.

Returning to logic, I propose that Burgess's distinction goes something like this. There are norms of reasoning that are implicit in, or somehow underlie, ordinary mathematical practice. A descriptive account is an attempt to uncover or describe those very norms, by constructing idealized mathematical models of them. In contrast, a normative logic is an attempt to say what the norms should be, never mind what the norms are. The normative logician claims that the norms that are implicit in practice are flawed, and he proposes better ones. The descriptive logician makes no such claims concerning logical competence. He attempts to recapitulate the norms as they are. If he is out to improve something, it is logical performance.

To provide further perspective on this issue, let us contrast the modeling approach to logic with Michael Resnik's adaption of the program of "wide reflective equilibrium", formulated by Nelson Goodman [1955] and used by John Rawls [1971] for an account of justice. Resnik states that wide reflective equilibrium is the "method logicians use when constructing systems for codifying correct reasoning or notions of logical necessity or possibility" (Resnik [1997, 159]). His 1997 book contains a succinct account of the method:

One starts with one's own intuitions concerning logical correctness (or logical necessity). These usually take the form of... test cases... that one takes to be logically necessary, inconsistent, or equivalent ... One then tries to build a logical theory whose pronouncements accord with one's initial considered judgements. It is unlikely that initial attempts will produce an exact fit between the theory and the 'data... Sometimes... one will yield one's logical intuitions to powerful or elegant systematic considerations. In short, 'theory' will lead one to reject the 'data'. Moreover, in deciding what must give, not only should one consider the merits of the logical theory *per se*... but one should also consider how the theory and one's intuitions cohere with one's other beliefs and commitments, including philosophical ones. When the theory rejects no example one is determined to preserve and countenances none one is determined to reject, then the theory and its terminal set of considered judgements are in . . . *wide reflective equilibrium*. (Resnik [1997, 159])

So Resnik explicitly allows logic to be normative—on broadly Quinean grounds. The logician is free to propose changes in his or others' inferential practice, in order to bring the total web of belief into reflective equilibrium. If an established practice of making and refusing inferences conflicts with the theorist's overall beliefs, she may try adjusting the inferential practice to achieve equilibrium. Never mind that one of the things up for grabs here is what exactly counts as conflict and accord.

Let us see what to make of a the "gap" in the toy example where the logician develops conditions for handling free and bound variables. Assume that these conditions are regularly violated in practice. Surprisingly, the mathematicians never, or very rarely, deduce false conclusions from true premises. From Resnik's perspective, the logician is proposing a change in the

practice, to help achieve reflective equilibrium. The logician proposes that upon reflection, the mathematician should modify or slightly remold her “intuitions”, for the sake of a more unified and coherent logical theory.

On the logic-as-model approach to logic, the situation is different. The logician first claims that the mathematician violates a rule that is implicit in her own practice. The logician does not merely claim that the mathematician *ought* to be bound by this rule; the logician claims that the mathematician *is* bound by it, just because she is a member of the community of reasoners. The mathematician retorts that there is no such norm, claiming that the logician has mis-described the practice.

So far, the logician is in purely descriptive mode. He and the mathematician disagree on what the practice *is*. What exactly are the norms that underlie mathematical reasoning? When put this way, the dispute seems to be an empirical one concerning what is and what is not implicit in practice. To resolve the dispute, then, we must examine the practice itself. Of course, this practice contains occasional errors, but we are not concerned with those. We must examine *correct* practice. But who decides what is correct? Which aspects of the performance of reasoners are errors? If the logician is to claim that the norms are implicit in the practice, the practice itself must determine what is correct and what is not. So to adjudicate the dispute, we must examine patterns of praise, censure, and condemnation within the community of mathematical reasoners. Presumably, our adjudication of this dispute should conform to the standards of ordinary empirical investigation: obtaining large, stratified and random samples, and invoking standard deviations and statistical significance.

So far as I know, such studies have not been done. Even if there are such studies, they are not used to resolve gaps between logical theory and mathematical practice. But let us push on for a bit. Suppose that the empirical data shows that mathematicians do not censure each other for failing to pay attention to the logician’s supposed rules for free and bound variables, even after prompting, like: “How about this step? Does it seem OK to you?” Or suppose that we show some mathematician subjects pieces of argumentation where false conclusions are thereby obtained from true premises, obtained by violating what we take to be the rules. The subjects consistently find fault with the reasoning, but never at the place where the logician says they should (when new variables are introduced). So we find no empirical support for the logician’s claim that the rule is implicit in practice.

If the logician is to stay in descriptive mode, then, it seems that she must concede. This is indeed a gap between logical theory and *correct* mathematical practice—assuming again that have the community of mathematicians is the final arbiter of what is correct mathematical practice. And who else gets to decide that?

This is an instance of a general issue concerning naturalized epistemology. W. V. O. Quine [1981, 72] characterizes naturalism as “the abandonment of

the goal of first philosophy” and “the recognition that it is within science itself ... that reality is to be identified and described”. The naturalist sees epistemology as a branch of psychology, perhaps ignoring Frege’s rejection of psychologism. The naturalist need not abandon the distinction between correct and incorrect practice, but she must look to a practice itself for cues as to what the correct practice is. There is no perspective outside the practice of mathematics from which to evaluate it. The philosopher and the logician alike are sailors on the ship of Neurath—so long as the logician stays in descriptive mode, and limits his target to norms that are implicit in practice.

This perspective leads to some uncomfortable counterfactuals. In a similar context, Penelope Maddy [1997, 198] writes:

... To take a wild example, suppose mathematicians decided to reject the old maxim against inconsistency—so that both “ $2+2=4$ ” and “ $2+2=5$ ” could be accepted... If mathematicians themselves insisted that... they were pursuing a legitimate mathematical goal,... I find nothing in the mathematical naturalism presented here that provides grounds for protest.

Maybe the best conclusion to draw here is that the reaction against first philosophy has gone too far, and perhaps the same goes for the reaction against a normative perspective in logic. Presumably, with any norm, or any proposed norm, there is a telos. We can always ask what the point of the norm is. Someone can make a proposal concerning the underlying purpose of a certain practice, and then propose norms that help us toward that goal, whether or not those proposed norms are already implicit in practice. If that is revisionism, and a normative perspective on logic, then so be it.

Our logician can argue, quite plausibly, that an essential purpose of deductive reasoning is to preserve truth—to make sure we only infer true conclusions from true premises. The restrictions on free and bound variables pursue this goal. The logician shows that if the rule is ignored willy nilly, then one can infer a false conclusion from true premises. The mathematician might have some instincts that prevent this disaster without conforming to the rule, or perhaps she has just been lucky, but adopting the rule provides insurance that truth is preserved.

The logician’s proposal is similar to a parent’s insistence that everyone wash their hands before eating. There is an agreed telos—avoiding falsehood in the one case, avoiding illness in the other—and the rule is proposed to aid in the pursuit of this avoidance. There need be no claim that the rule is already implicit in the practice. If the dispute continues after this, it will turn on whether the proposed rule does in fact pursue the goal—truth preservation or bodily hygiene. The point here is that our theories and proposals are, or may be, revisionist in spirit. If the logician locates a flaw in the practice, he switches to a normative mode, saying what the rules ought to be, never mind what the rules in fact are.

There are philosophers and logicians who are far more critical of mathematical practice than the above logician who complains about variables. Intuitionism comes readily to mind, as does quantum logic and some of the literature on relevance logic. Again, most logicians are resistant to such revisionism. We are out to model *mathematics*, not some supposedly improved substitute for mathematics. But what are the grounds for this anti-revisionism, and how do we distance it from more mundane uses of logic to improve reasoning (like the toy example)? If we admit that the practice is, or may be flawed, then why are we so quick to dismiss intuitionism, quantum logic, and relevance logic?

With this framework and our toy example in place, let us turn to relevance logic, and then later to intuitionism. The relevance logician claims that a conclusion does not follow from a set of premises unless there a relevant connection between them. Relevance logicians disagree among themselves concerning what this relevance comes to, but they all agree in rejecting the validity of the inference *ex falso quodlibet*: From Φ , and $\neg\Phi$, infer Ψ . The following argument is typical:

Either we will eat carrots or we will eat peas.
We will not eat carrots.
We will not eat peas
Therefore, we will eat broccoli.

Broccoli? A quick look at truth tables shows that this argument is classically valid, but the relevance logicians reject it.

It is fair to assume that most reasoners—even most competent reasoners— would balk at this argument, at least until they have taken a logic course. I balked at it myself once upon a time, and I presume that most of you did as well. After taking a logic course, most students end up accepting the validity of this argument. The question is whether they are rationally convinced or just beaten into conformity.

The soundness of *ex falso quodlibet* is sometimes established by a short argument attributed to C. I. Lewis, but apparently known to medieval logicians:

- | | | |
|-----|------------------|-----------------------------|
| (1) | Φ | premise |
| (2) | $\neg\Phi$ | premise |
| (3) | $\Phi \vee \Psi$ | 1, disjunction introduction |
| (4) | Ψ | 2,3 disjunctive syllogism |

In one way or another, relevance logicians reject this reasoning. The most common route is to deny the validity of disjunctive syllogism, but others reject the universal applicability of disjunction introduction, and at least one of them accepts both inferences, but denies the transitivity of deduction.

Using Resnik's reflective-equilibrium framework, the debate can be framed as follows. The classical logician concedes that *ex falso quodlibet* does

not accord with our initial intuitions concerning logical validity. Pre-theoretically, we are not inclined to accept arguments like the one about broccoli. But this intuition is modified in light of theory. We see that classical logic provides a useful and powerful theory, and the datum of rejecting the broccoli argument is given up. We are initially inclined toward disjunction introduction, disjunctive syllogism, and the transitivity of deduction, but we are inclined against *ex falso quodlibet*. We thus notice a conflict in our intuitions—thanks to C. I. Lewis. For the sake of reflective equilibrium, the classical logician *decides* to give up the disinclination toward *ex falso quodlibet*, maintaining the other intuitions.

The relevance logician makes different decisions when confronted with the same dilemma concerning initial intuitions. She insists on the invalidity of *ex falso quodlibet*, and compensates elsewhere. Resnik allows for the possibility that the classical logician and the relevance logician might both achieve reflective equilibrium, each by his own lights. That is each of them might get to the point where his or her own logical “theory rejects no example one is determined to preserve and countenances none one is determined to reject”.

How does the debate play itself out on the logic-as-model approach? The first thing to note is that some relevance logicians, perhaps most of them, advocate teaching classical logic in introductory courses, for roughly the same reason that we teach the mechanics of point masses in beginning physics classes. That is, advocates of non-classical logic hold that classical logic is a good model of correct inference. This model is simple, easy to present and study, and more or less tractable. Just what we want in a model. However, the relevance logician then claims that *ex falso quodlibet* is a *gap* between logical theory and inferential practice. She claims that a relevance logic is a more accurate, but perhaps less tractable, model of correct inference.

From the logic as model perspective, this is just the place where the classical logician disagrees. He claims that *ex falso quodlibet* is implicit in the practice of mathematics. The validity of this inference in the model is not a gap between logical theory and mathematical practice.

But what of the reactions against inferences like the broccoli argument? The classicist might argue that our negative reaction to this argument is due to Gricean pragmatic considerations. If we were given the premises about eating carrots or peas, and not eating carrots and not eating peas, we would have no earthly reason to draw *that* conclusion, even though it would be correct to do so. Notice, however, that the distinction between implicature and logical implication is highly theoretical, and begs the present question.

So how do we adjudicate this dispute? How do we determine just what is and what is not a gap between logical theory and mathematical practice? Our two theorists agree on a model, but they disagree over whether a certain inference pattern is a gap between theory and practice. As with the toy case of free and bound variables, the dispute seems to be an empirical one. To see

if *ex falso quodlibet* is in fact implicit in practice, we must have a look at actual mathematical practice, to see if the inference is ever used, and, if it is, whether this use is regarded as correct, or is rejected as a mistake in inference—as opposed to rejecting it as foolish even if technically correct—by competent reasoners. And again, the envisioned study should conform to the standards of normal empirical inquiry, using stratified and random samples, calculating statistical significance, etc.

As far as I know, there have been no such studies. However, several years ago, there was an interesting exchange Burgess and Stephen Read over this issue. Burgess claimed that *ex falso quodlibet* is regularly used in mathematics. A typical instance occurs in a proof-by-cases. Suppose that a mathematician is trying to establish a conclusion Ψ . She shows that there are three cases to be considered, that Ψ follows from two of those, and that the third case leads to contradiction. So she concludes that Ψ holds in any case.

Clearly, the mathematician's argumentation here is correct. The dispute concerns how it should be understood. The classical logician reconstructs the reasoning as follows. The statement about the three cases amounts to an intermediate conclusion in the form $\Phi_1 \vee \Phi_2 \vee \Phi_3$. We are done if we can show that Ψ follows from each disjunct (via disjunction elimination). The mathematician derives the desired conclusion Ψ from Φ_1 and derives Ψ from Φ_2 . That leaves Φ_3 . This case is shown to be contradictory, and so impossible. Typically, the mathematician stops here, satisfied that the proof is complete. In the classical reconstruction, we invoke *ex falso quodlibet* to derive Ψ from the contradiction, and then invoke disjunction elimination to finish things up.

I presume that classical logicians accept this reconstruction as giving the underlying form of the mathematician's reasoning. The mathematician did in fact invoke *ex falso quodlibet* as part of his proof by cases. But the relevance logician disagrees, arguing that classicist has mis-described the practice. The relevance logician agrees that the conclusion Ψ is acceptable—either outright or provisionally on the presupposition that the background mathematical theory is consistent—but she insists that *ex falso quodlibet* was not invoked to get it. Moreover, unlike our imaginary mathematician in the toy scenario concerning free and bound variables, the relevance logician is not without resources here. Read and other relevance logicians provide an alternate reading of the mathematician's reasoning, a reading that does not invoke *ex falso quodlibet*. It turns out that the rules underlying proof by cases are more subtle than the classical model indicates. Or so says the relevance logician.

I am at a loss to see how to continue the debate, let alone resolve it in the favor of the classical logician, much as I would like to resolve it that way. The parties disagree over whether a certain inference pattern is in fact implicit in the practice of mathematics, whether or not pre-theoretic intuitions militate against the inference. Unfortunately, mathematicians do not usually justify their steps in reasoning by citing rules using the terminology of logic books.

“Here we use conjunction elimination; here we use disjunction syllogism; here is *ex falso quodlibet*.” In just about every case, including the straightforward undisputed ones, the logician has to *reconstruct* the reasoning, hopefully to get it to fit his or her favored model. In the present case, we are faced with alternate reconstructions of the practice.

The classical reconstruction is simpler, and uses fewer, less wordy rules of inference. The relevance logician concedes this, but argues that his reconstruction is the more accurate, since it does not invoke the invalid *ex falso quodlibet*. To support this last, the relevance logician cites the initial, pre-theoretic reactions against this inference, and his ability to reconstruct the reasoning. The classical logician surely concedes the data concerning the initial reaction. He argues that, nevertheless, the inference is correct. He might cite the Lewis argument, together with the fact that *ex falso quodlibet* comes up pretty often in his own reconstruction of mathematical reasoning.

As with the toy example above, either party has the option to concede the descriptive issue, and go on to make a normative claim. The classical logician might grudgingly admit that *ex falso quodlibet* is not implicit in practice, but then argue that the inference is nevertheless correct. He might cite a goal of mathematical inference—the preservation of truth—and then argue that *ex falso quodlibet* furthers this goal. It is obvious that one never goes from truth to falsehood with this inference—because the rule never starts with truth.

Analogously, the relevance logician might also wax normative. She might concede that *ex falso quodlibet* is implicit in practice, but argue that it should not be—for reasons analogous to logical hygiene. Disjunctive syllogism and *ex falso quodlibet* is safe enough if one is certain that the background principles and theories are true (as in, say, Peano arithmetic). However, *ex falso quodlibet* is dangerous in situations where one is reasoning from faulty, or potentially faulty premises. Suppose that a subject realizes or suspect that her beliefs are inconsistent, but does not know which of them to reject. So she provisionally accepts each of her beliefs (like we all do). *Ex falso quodlibet* would allow her to infer anything at all, surely an unwise license.

Rather than pursue this normative matter further, let us turn to intuitionism. I characterized the debate over relevance logic as a dispute over whether a certain inference pattern constitutes a gap. Unlike relevance logic, intuitionism is a straightforward revisionism, urging the rejection of the law of excluded middle, and other principles and inferences based on it. I presume that contemporary intuitionists concede that excluded middle, double negation elimination, and the like, are well entrenched in the practice of mathematics. Uses of the axiom or rules in question are not questioned, and theorists are subject to criticism for failing to use the rules. That is, intuitionists are losing the battle for the hearts and minds of the mathematical community.

The traditional intuitionists, L. E. J. Brouwer and Arend Heyting, claim that the logic revision is tied to the mind-dependent nature of mathematics:

The... point of view that there are no non-experienced truths ... has found acceptance with regard to mathematics much later than with regard to practical life and to science. Mathematics rigorously treated from this point of view, including deducing theorems exclusively by means of introspective construction, is called intuitionistic mathematics... [Classical mathematicians believe] in the existence of unknown truths, and in particular [apply] the *principle of excluded third* expressing that every mathematical assertion... either is a truth or cannot be a truth. (Brouwer [1948, 90])

[W]e do not attribute an existence independent of our thought, i.e., a transcendental existence, to... mathematical objects... [M]athematical objects are by their very nature dependent on human thought. Their existence is guaranteed only insofar as they can be determined by thought. They have properties only insofar as these can be discerned in them by thought ... Faith in transcendental... existence must be rejected as a means of mathematical proof... [T]his is the reason for doubting the law of excluded middle. (Heyting [1931, 52-53])

In reply, the defender of classical logic may argue that mathematical objects and/or mathematical propositions are in fact mind independent. The classical counter-claim would be that the metaphysical principle of independence or determinacy underwrites excluded middle. Or else the classicist can argue that excluded middle does not depend on any metaphysical principle. That is, one might try to defend classical logic independent of the metaphysics. After all, the correct logic should not depend on substantial, contentious metaphysics. Logic is supposed to be neutral on such matters, since logical reasoning is involved in metaphysical, or any other, disputes.

Michael Dummett shifts the focus of this debate away from metaphysics to semantics. Starting with Dummett [1973], he argued that any consideration concerning logical principles must ultimately turn on questions of *meaning*. Since language is a public medium, the meanings of the terms in a language are determined by how the terms are correctly used in discourse. This common-sense view of language supports Dummett's *manifestation requirement*, a thesis that anyone who understands the meaning of an expression must be able to demonstrate that understanding through her behavior —through her *use* of the expression. Other pieces of Dummett's argument are a requirement that the logical constants be graspable one at a time, independent of each other (separability) and a requirement that there be harmony in the use of each connective.

Dummett argues that these considerations have ramifications for the proper meaning of the logical terminology. In the prevailing Tarskian semantics, the truth conditions of a complex formula are defined in terms of the truth conditions of its subformulas. Typically, truth conditions can hold (or not) independent of our abilities to know that they do. On Dummett's diagnosis, this is the source of bivalence and excluded middle. Dummett argues that if the language is undecidable, a semantics like this violates the manifestation requirement. On a classical, bivalent interpretation of a mathematical theory,

the central notion is that of truth: a grasp of the meaning of a sentence . . . consists in a knowledge of what it is for that sentence to be true. Since, in general, the sentences of the language will not be ones whose truth-value we are capable of effectively deciding, the condition for the truth of such a sentence will be one which we are not, in general, capable of recognising as obtaining whenever it obtains, or of getting ourselves into a position in which we can so recognise it. (Dummett [1973, 105])

Dummett claims that *verifiability* or assertability should replace truth as the main constituent of a compositional semantics. In mathematics, verification is *proof*. Dummett's proposal thus invokes the central theme of *Heyting semantics* for intuitionistic logic. Instead of providing truth conditions of each formula, we supply proof conditions. Supposedly, this semantics does not sanction excluded middle, double negation elimination, etc.

In response to this charge, the classicist has a few options. First, she can meet the charge head on, and try to undermine Dummett's considerations concerning the learning of language. As it stands, the manifestation requirement is innocent enough, probably beyond dispute, but one can challenge Dummett's articulation of it, along with the requirements of harmony and separability. Or the classicist can concede that classical mathematics does not enjoy the level of justification that Dummett demands, but then argue, or just assert, that mathematics does not need this level of justification. Mathematicians have developed their own standards of justification and coherence, and these are not in need of fixing.

Our question here is what to make of this debate? As we saw, Resnik's method of wide-reflective equilibrium allows revisionist elements in logic. The logician is explicitly told to adjust his inferential practice in light of his overall beliefs, including philosophical views. So the intuitionist can be described as busy trying to attain reflective equilibrium. What are we to make of intuitionism on the logic-as-model approach? I submit that it is not straightforward to simply dismiss the whole program, out of hand, as pursuing a different goal than that of other logicians. The toy example indicates that normative, and potentially revisionist elements can arise in a standard, descriptive program of logic.

So what is the position of these intuitionists with respect to the above imaginary logician who proposes restrictions, and thus revisions, in how free and bound variables should be handled? Let us put aside for now the question of whether the intuitionist, or the logician for that matter, is correct. The issue concerns the status of normative, revisionist claims. From the logic-as-model perspective, does the intuitionist even get on the table? I propose that the difference between the intuitionist and the above imaginary logician is merely a difference in degree, not a difference in kind. All of the revisionists argue that the practice is flawed, and they propose ways to put it right. The difference lies in how flawed they think the practice is (and, of course, whether they are correct that the practice is flawed).

One difference, perhaps, is that the logician in the toy example is not asking the mathematician to give up any *theorems*. In the scenario, the logician and the mathematician agree on the correctness of the various *results* of mathematics. The logician may be surprised that the mathematician is not led into falsehood, but the two agree on what falsehood is. The logician is merely proposing changes in the deductive reasoning toward true conclusions—on the grounds of logical hygiene. In contrast, the intuitionist demands that large portions of established mathematics be given up. For example, the intermediate value theorem is not intuitionistically correct. However, there is no guarantee that “mere” logical hygiene will not lead to changes in results. If the logician is correct that the mathematician is being sloppy, then it is more than likely that the latter will eventually make some mistakes, and these can show up in incorrect conclusions. What is the point of so-called hygiene if it ignoring it is completely safe? Moreover, the rigorization of the calculus (whether via the ϵ - δ route or non-standard analysis) shows that logical rigor can lead to revisions in results. For example, by not paying attention to the differences between pointwise and uniform convergence and continuity, famous mathematicians did make mistakes.

Another potential difference between the imaginary logician and the intuitionists is that the former’s claims are purely logical, while the latter’s are explicitly philosophical. But this assumes that we can make out a relevant distinction between internal logical matters, and external philosophical meddling. If the distinction could be maintained, then perhaps we could dismiss the philosophical meddling, out of hand, as not part of our program. However, I submit that there is no distinction to be drawn here. The logician’s claim about free and bound variables is based on a model of how language functions, and, thus, on a model of the truth conditions of natural languages. The question of whether this model is apt is philosophical if anything is. The intuitionist is likely to dispute that very model, and the mathematician who is prone to ignore the logician is free to dispute it as well.

To reiterate, then, I suggest that the difference between the logician proposing care with free and bound variables and the various intuitionists is a distinction in degree, and not a distinction in kind. Both theorists propose changes in existing practice, in order to make it faithful to what the practice is—or should be—attempting to accomplish, perhaps in light of a disputed thesis concerning the goals of the practice. I presume that there are limits to how much revision we can tolerate, and still claim that we are being faithful to an existing practice, or to existing goals. One might dismiss intuitionism, not for being revisionist, but for going too far down the revisionist path. But if we are going to dismiss a revisionist program simply because it goes too far, then we need a principled account of what the limits of revision are, or what the limits should be. I do not know how to do that without begging the question at hand.

This paper is perhaps disappointing in that it has no definitive conclusions. The philosophical topic of mathematical models could use further exploration. Is there a common feature of all models, whether in a physical sciences like chemistry or a social science like economics? What is the relationship between a purported model and what it is a model of? Models are often part of explanations. How does this work, given that strictly speaking, most models are false—and gaps are to be expected?

As a classical logician, I'd love to have a definitive knock-down argument against relevance logic and intuitionism. I have not given up the belief that classical logic is correct, although I am not sure what this correctness comes to. Moreover, I do not see a definitive resolution of the issue on the near horizon. The fruitful logic-as-model perspective provides a general framework for raising these descriptive and normative issues, and for distinguishing descriptive and normative programs, but we do not know enough about how to determine what the norms underlying a practice are, let alone what those norms should be. I'll have to rest content with the hope that I have pointed in an interesting direction, along with the call to follow Arnold Ross's advice once more: think deeply about simple things. Another corollary of Dr. Ross's slogan is that there is work to do.