Why do informal proofs conform to formal norms?

Jody Azzouni

*Tufts University*

jodyazzouni@mindspring.com

*Abstract*: Kant discovered a philosophical problem with mathematical proof. Despite being *a priori*, its methodology involves more than analytic truth. But what else is involved? This problem is widely taken to have been solved by Frege's extension of logic beyond its restricted (and largely Aristotelian) form. Nevertheless, a successor problem remains: both traditional and contemporary (classical) mathematical proofs, although conforming to the norms of contemporary (classical) logic, never were, and still aren't, executed by mathematicians in a way that transparently reveals *why* these proofs—written in the vernacular to this very day—succeed in conforming to those norms.

Keywords: *mathematical proof*, *reasoning*, *logic*, *meaning*, *natural language.*

Jody Azzouni publishes in ontology, philosophy of logic, mathematics, and science. He has four books: *Metaphysical myths, mathematical practice: The ontology and epistemology of the exact sciences* (1994), with Cambridge University Press, *Knowledge and reference in empirical science* (2000), with Routledge, *Deflating existential consequence: A case for nominalism* (2004), and *Tracking reason: Proof, consequence, and truth* (2006), both with Oxford University Press.

1. *Kant's problem.* Analytic statements are to be recognized by an analysis of the concepts involved in those statements. It can be proven in Euclidean geometry that the interior angles of any triangle sum to 180 degrees. The proof proceeds, however, not by an analysis of a free-standing triangle (and its conceptually-given parts), but by embedding it within a larger figure. This is a ubiquitous and crucial aspect of mathematical proof: proofs are routinely facilitated (and often *must* be facilitated) by introducing concepts and methodology that—strictly speaking—go beyond (sometimes *far* beyond) the conceptual givens of what's to be shown. What makes this a problem for Kant is that (for him and for his contemporaries) analyticity exhausts the resources of *logic*.

It's remarkable that this problem was first discovered by Kant, who wasn't a professional mathematician (although, of course, he was familiar with mathematics). It's perhaps no surprise that *some* of his philosophical predecessors, Locke and Hume, failed to see that their characterizations of analytic truth weren't compatible with the epistemology of mathematical truth—with how, in fact, mathematical truths are (by professionals) recognized as such. It's more remarkable that this "explanatory gap" wasn't seen by Leibniz, a first-rate mathematician as well as philosopher, who—otherwise—thought deeply and far-sightedly about the nature of mathematical proof (see Hacking (1973)).

The problem may be posed this way: a medium for proof is required that bridges the gap between the concepts involved in the statement of a theorem, and those that arise in a proof of that theorem. Kant, notoriously, hoped that *intuition* (of space and time) would fill that gap. This solution loses plausibility once the scope of mathematical proof is recognized to extend well beyond the traditional domains of geometry and number. Strictly speaking (although this is largely a twentieth-century insight) concepts from a subject matter of *any sort at all* may—in principle—be treated in a (purely) mathematical fashion; they needn't arise or be derived from space and time. Four examples of such are these: Turing machines, rigid-body dynamics, "spaces" of functionals, and formal languages. Mathematical practice is indiscriminate with regard to what it may be practiced on, requiring only *proof*. (For further details, see Azzouni (2000a).)

2. *Artificial languages as a medium for proof*? Kant's problem was taken to be solved by Frege's artificial language, one governed by the rules of a logic, and accompanied subsequently by the emergence of set theory (as its ultimate subject matter).[1] Indeed, the subsequent *Principia* of Whitehead and Russell should be seen for what it was: a program of *translating* mathematical theorems—and their proofs—into *formal artificial languages*. This program was successful in this sense: the *results* of mathematics survived the translation. The important question is: what about ordinary mathematical proofs? Did *they* survive translation?

That proofs in artificial languages are key to solving Kant's problem is supported by the fact that such artificial languages can contain vocabulary from ordinary language, and from the specialized subject-matters of mathematics. (The idea is that it's by virtue of such formal proofs that ordinary informal mathematical proofs are understood and recognized to be successful.) However, the meaning-properties of informal-mathematical terms cannot play a role once those

[1] Frege had important predecessors, of course. But it was he—pretty much single-handedly—who both put the full logic together, and conceived the project of its serving as the "foundation" for mathematical proof. I should stress, to avoid misunderstanding, that I'm not speaking of the *other* foundational project of reducing mathematics to set theory.

terms are relocated to one or another formal language *except insofar as such aspects of meaning have been codified formally, i.e., syntactically*.[2]

Indeed, only slowly did the following important properties of Frege's artificial language, and its descendents, emerge. Because the (intuitively) intended interpretations or meanings of the non-logical terms play no role in the reasoning licensed in such languages, those terms are open to *reinterpretation*—compatibly with the axioms and the background logic, of course. Intuitively, the meanings of the terms of ordinary language—and this includes mathematical terms as they are ordinarily understood—are presumed to enable those terms instead to have *specific intended objects* as their referents.

Thus, in contrast to this impression of the meanings of the terms of ordinary languages, the "meanings" of the non-logical terms of formal languages are idle. Their content is instead provided by conditions (axioms) together with the explicitly given background logical rules. As a result, first-order languages have non-standard models; higher-order languages can't fix the specific objects "intended": their models are open to permutations.

Another property is that formal proofs in such artificial languages are mechanically recognizable on the basis of axioms and explicitly given background logical rules. That a

[2] Rav (2007, 303) seems to deny that antecedently meaningful terms can thus appear in a formal setting, and to also deny that syntactic codification of their meanings can (at least in principle) fully reprise their roles in informal mathematical proof. The oddity is that—at least at this point in his paper—he seems to imply that these claims *follow from* the standard technical definitions for "formal proof" or "derivation." They don't.

sentence follows from earlier sentences is determined by the rules applicable to those earlier sentences.[3]

An early, natural, and persistent model for ordinary mathematical proofs is that they are actually proof-*sketches*: many steps are missing, many essential concepts aren't explicit. It would be nice to think that the *Principia* project, and subsequent projects to formalize ordinary mathematical proof, vindicate this view of mathematical reasoning by showing what a proof-sketch in the vernacular becomes once all the *lacunae* are filled in.

There are problems, however, with taking this perspective. Here are *some*. Frege didn't simply invent (or discover) modern logic; as just noted, he invented (or discovered) artificial (or "formal") languages. The logical tools, the connectives and the quantifiers, by which reasoning occurs in these languages aren't *in* natural languages. Frege's artificial language, and the ones that it inspired, are all far more (theoretically *and* proof-theoretically) tractable than the natural product:[4] how reasoning occurs in them is transparently indicated by their terminology.[5]

---

[3] *A refinement*: In first-order (and in some other) artificial languages, proof-theoretic rules exhaust the accompanying notion of validity (the proof-theoretic methods are complete). Otherwise, the *recognizable* reasoning is characterized by such rules, although they are not complete with respect to the accompanying notion of validity.

[4] These are truisms as far as natural-language quantifiers are concerned. Even the more complicated (but still artificial) generalized quantifiers are seen only as illuminating *models* of the complex semantic phenomena found in natural languages. There is controversy, for example, about whether truth-functional connectives are in natural languages: the ordinary language terms to be pressed into service as "logical" all seem to have additional content that goes beyond truth-functionality. The dispute is over the semantic status of that content. See Carston (2002, chapter

So, it seems that whatever the mathematician had (or has) in mind as a way of filling out the missing steps, and the making explicit of tacit concepts in an ordinary mathematical proof, it couldn't have been (couldn't be) *anything* like proof in one or another such artificial language. First, before Frege, there were no such languages. Second, ordinary mathematical proof, and the experience of the mathematician convinced by such, seems quite different from a formal proof, and a mathematician's experience of that. (a) Understanding an ordinary proof seems to proceed by the grasping of mathematical concepts, grasping the nature of the objects those concepts are about, and recognizing what follows from these—not by the recognized application of (explicitly formulated) rules. (b) Related to this, ordinary mathematical proof seems to be topic-specific, its transitions governed by insights into the subject-matter of the proofs.

Recall, furthermore, that Kant's problem is *epistemological*. A medium through which (and by which) the mathematician reasons successfully *in the vernacular* is to be provided. But, in general, (c) when an ordinary proof (in the vernacular) is translated into one or another formal language, comprehension and certitude decline. One—almost always—understands (and is convinced by) the ordinary mathematical proof much better than (by) formalized cousins.[6]

---

3) for discussion of the most straightforward-looking case, "and," and also for references.

[5] Again, modulo the issues noted in footnote 3.

[6] (a), (b) and (c) are points made repeatedly in the literature. I too have stressed them in earlier work, e.g., in Azzouni (2004), (2005), (2006). For another discussion stressing these points about ordinary mathematical proof, along with (some) earlier citations, see Rav (2007). It's also worth noting that the range of formalizations for a particular informal proof isn't sharply defined—options in formalization are always possible. For the sake of argument, I'm not pressing this point.

This suggests that proofs in invented artificial languages *haven't* solved Kant's problem. Kant was concerned with the epistemology of ordinary mathematical proof. But the foregoing points about formal proofs seem to make it impossible for our recognition of the success of informal proofs to arise from our recognition of the properties of the formal proofs that they are (supposed to) correspond to.

The formalist approach to Kant's problem isn't without resources to try to meet these objections. Consider, for example, this response to (c). To be explained is why the mathematician's comprehension of (and certainty in) an informal proof declines when it's formalized. The explanation the formalist can give is that this decline isn't *total*: it's not that one experiences the resulting formal proof as *utterly* incomprehensible. Rather, it's that one is overwhelmed by *details*. One can see that each step correctly follows from earlier ones; one experiences the correctness of each step. What one doesn't experience is a sense of the correctness of the *whole*. (One's experience of correctness fails to be transitive.)

Thus (so this suggestion may go) one or another formal analogue of an informal proof *is* what gives one the impression of correctness as it arises during the understanding of the latter. One abstracts or codifies a series of detailed steps in one's mind, and (somehow) one's sense of the correctness of the whole shines forth. One reason this formalist response isn't convincing is that the reaction to the detailed formal proof isn't an "ah-ha!" reaction: So *that's* why the informal proof works. One already experiences knowledge of that, with respect to (perspicuous) informal proofs, and nothing further seems added—psychologically—by filling in all the details. One's sense of the correctness of the proof—one's understanding of it—is *not* enhanced by studying a formal analogue of an informal mathematical proof. This seems, therefore, an inadequate response to (c), apart from the other problems that a formalist solution to Kant's problem still seems to face.

3. *Some points about the phenomenology of the experience of successful mathematical proof.* I've been speaking in the foregoing of the perception of the correctness of the steps of an informal proof. Three possible misperceptions about the nature of this experience should be gotten out of the way, especially because two of them invite cheap solutions to Kant's problem.

First, this recognition of the correctness of an informal mathematical proof is accompanied by a certain striking *experiential* quality. This is the sense of *necessity* that accompanies successful proof: that the steps in a proof *must* follow from earlier steps.[7] Notably (and here's the first misconception to avoid), this sense of necessity isn't *ontological*; it isn't due to the subject matter. There is, of course, a venerable tradition of trying to explain the sense of necessity as a reflection of the timeless and unchanging (necessary) nature of mathematical *abstracta*. But apart from the fact, noted earlier, that notions of any sort—including quite temporal ones—may be treated mathematically, there's no reason why the topic of one's reasoning being timeless and necessary ought to affect one's experience of one's *reasoning*. The motivation for this line of thought seems to turn on a confusion in the application of modalities (and perhaps even on a use/mention error): opposing it is that one can reason about the unchanging and timeless (the ontologically necessary) without that reasoning *itself* striking one

[7] I've elsewhere described this as a *sense* of compulsion. See Azzouni (forthcoming). I don't mean this to be a metaphor. Notably, one doesn't experience the "compulsion" in question while inferring a conclusion. It arises when one tries to avoid believing that conclusion (after having seen that it follows). The sense of compulsion is especially visible when one tries to live within the confines of an alternative logic, and *really disbelieve classical implications*.

as necessary (or compelling) in any way. The cheap solution to Kant's problem that this invites, of course, is an explanation of one's recognition of successful proofs in terms of one's intuitive appreciation of the objects the proof is about. For a contemporary version, see Katz (1998). For criticisms, see Azzouni (2000b).

Next, this sense of necessity shouldn't be confused with the (mere) recognition of the *correctness* of an inference being its conformity with certain rules. One can recognize of rules *of any sort at all* that they have been executed correctly. But this sense of "necessarily follows" needn't be accompanied by the recognition that rules are even *relevant*; this is often the case with informal reasoning. We recognize that T must follow from S even though we don't recognize *any* rules as involved. We reasoned for centuries—for *millennia*, actually—before it was recognized that there might be rules we were reasoning with respect *to*.

It should be stressed that the experience of step-by-step "necessarily following" of a (fully explicit) proof occurs only in *classical* settings, or with respect to the application of logical rules (in a nonclassical setting) that overlap with applications of rules of classical logic. If the logic has rules that license inferences disallowed classically, then when one applies such a rule, one experiences no sensation of anything "necessarily following." One experiences only the sheer mechanical nature of the steps: that a sentence follows from others by virtue of the rules. One experiences this even if the rules—in question—are to codify the "meanings" of the (non-classical) connectives. One might be tempted to the view that the item missing from the analysis here is the notion of *truth*. Only the classical inference rules are "truth-preserving," not the alternatives. Unfortunately for this line of thought, the notion of truth is promiscuous: it can be grafted onto any logical system *whatsoever*, and a corresponding notion of validity can be defined relative to that logic. (I discuss this further in Azzouni (forthcoming).)

Lastly, this sense of necessity should not be characterized as the *certainty* that one experiences when comprehending a valid proof.[8] The cheap solution to Kant's problem that *this* conflation invites is a dissolution of it: There is nothing special about one's perception of mathematical *proof*; it's just like one's perception of anything else that one is (pretty) sure of. However, it's widely recognized (or ought to be) that one can be utterly certain that something follows from something else, and yet be *wrong*. What one experiences (although fallibly) is the sensation that a step (in a proof) *must hold* if earlier ones do. What one experiences may also be put this way: a proof, if successful, is "valid." The important question (apparently still) facing us is Kant's problem: the source of one's experience of its "validity."[9]

---

[8] *I* have stressed this point in earlier work, *pace* Rav (2007, 317).

[9] There are some terminological issues I've been trying to dance around. It's best to just make them explicit. I'm focused on a certain phenomenological experience during (certain cases of) reasoning, that of recognizing "correctness," or "validity": given certain assumptions, that something must follow from them. But various technical notions have co-opted these ordinary ways of speaking. "Validity" is often used to describe model-theoretic validity, and "correctness," correctness according to certain rules. When I use "correct" or "valid," (and their cognates) as when describing our recognition that a certain piece of informal mathematical reasoning is "correct" or "valid," I'm speaking of intuitive recognition, and not of the technical notions. Similarly, when we (intuitively) say that T "must be true if S is true," (or that a valid piece of reasoning is "truth-preserving"), this is an ordinary use of "truth" that's restricted to (and only applies to) intuitively-recognized correct (classical) reasoning. As noted in this section, the truth idiom is promiscuous, and can be co-opted for use in (non-classical) logical formalisms. When describing the intuitive idea of "consequence" or "implication," I'll speak in one or

Posing "the important question" the way I just have gives the impression that the state of play is our needing to reprise Kant's problem in light of the fact that proofs in artificial languages can't help us solve it. But that's not *quite* where we're at. Frege's invention changed the philosophical landscape forever, and it isn't Kant's problem we now face, but a successor problem. I now turn to motivating (and then stating) this successor problem.

4. *Formalized proofs are norms for informal mathematical proofs*. The first point to observe is that formalized proofs have become the norms of mathematical practice. And that is to say: should it become clear that the implications (of assumptions to conclusion) of an informal proof cannot be replicated by a formal analogue, the status of that informal proof as a successful proof will be rejected.[10]

Rav (2007, 306-7) denies this, but for what can easily be seen to be weak reasons. He suggests that should a computer check of a derivation produce the claim "not valid," it will be presumed that there is a flaw in the formalization or a bug in the computer program. It may be presumed indeed. But this observation misses the point. The norm is this: *There is* a formal analogue of a purported informal mathematical proof or else the latter fails to be a proof. It's

---

another of these ordinary ways (just as we ordinarily do). I'll sometimes use quotation marks to indicate that I'm speaking intuitively, and I'll often add technical adjectives (e.g., "*model-theoretic* validity") when I want to allude to technical notions instead. The co-opting of ordinary terminology by logicians has been so thorough that it has become impossible to allude to the ordinary experience of reasoning *without* using notions that have been so co-opted. Thus the need for footnotes like this one to avoid misunderstandings.

[10] There is flexibility in the application of this norm insofar as tacit assumptions (and concepts) can be (and often are) taken as presumed present in an informal proof.

another question entirely whether one—in particular circumstances—can show (or be confident) that there *isn't* a formal analogue.

A similar point applies to Kreisel's (1969, 27) remark, quoted by Rav (2007, 308) with approval (italics in the original):

> What use is it to the mathematician to be told that his arguments *can* be formalized and thereby justified … if (i) he does not formalize them, yet (ii) is convinced of the conclusion? The evidence of *this* conviction must be elsewhere.

Maybe so, but none of this touches (or is even relevant to) the evident fact of the presence of the norm of formalized proof with respect to ordinary mathematical practice. The nonexistence of a formal analogue *does* rule out the status of a purported ordinary mathematical proof being a successful proof, regardless of the felt convictions of mathematicians about this.

5. *A matter of certainty vs. a matter of justification*. We should distinguish two (ordinary) ways of speaking of justification. A mathematician may be described as "justified" in his belief that an informal proof is indeed successful. And this will turn on issues like: when does a mathematician *really know* (in practice) that an informal mathematical proof is a good one? When he really knows, he really is justified in his claim that such-and-such is a good proof. He has a right to his *certainty* about this. This use of "justification" needn't be (directly) related to *what it is about the proof* that makes it a good proof. For example, a mathematician may be entirely justified in believing a certain ordinary informal proof is a good one (is "valid") on the grounds of professional testimony: he can trust his (very competent) colleagues about this.

On the other hand, there is a (quite distinct) use of "justification" that requires alluding to the norms involved—to whatever it is that (officially) *makes* an ordinary informal proof a good one. That is, that there is a formal-language analogue of the proof.

Why not claim that the norm in question is the (model-theoretic) validity of the informal proof, and not its corresponding to a formal analogue? (And, correspondingly, why not claim that the failure of there being a formal analogue of a proof is taken as reflecting the failure of the informal proof to be valid?) In the classical setting, in the first-order case, these (extensionally) amount to the same thing. In classical higher-order logics, these *still* (in practice) amount to the same thing since any higher-order language and deductive system corresponds to a many-sorted first-order system. That is to say, for any informal mathematical proof, taken as (implicitly) higher-order, *there is* some formal higher-order language and deductive system wherein a corresponding formal proof *exists*, and so there is some first-order system wherein a corresponding formal proof exists as well. Apart from the cases where model-theoretic validity can be recognized proof-theoretically, model-theoretic validity can have no normative role to play because it's noneffective. I dwell on this important point in section 8.

6. *The successor problem to Kant's problem*. That said, there is an important issue in the *neighborhood* of all this. Any norm has to be practically implementable. Therefore, if a practice *is* in accord with a norm, we need (at least) *some* idea of how that's being managed. The striking point about the *Principia* program is this: implementing it didn't expose a crisis in ordinary mathematical practice (i.e., the revelation of numerous informal mathematical proofs failing to actually show what they were purported to show). The recent computerization of formal analogues of informal mathematical proofs, and the checking of such reveals exactly the same thing: there has been no widespread rejection of mathematical results. Ordinary mathematical

practice already *is* in accord with formalization norms, and apparently has been for its very long history.[11]

MacKenzie (2001, 323) writes:

> Many rigorous arguments in ordinary mathematics have been replaced successfully by formal proofs, using automated theorem provers and proof checkers, especially the AUTOMATH and MIZAR systems … What is most remarkable about these many replacements of the rigorous arguments of mathematics with formal, mechanized proofs … [is that] it is a conservative process. Applied to programs, hardware designs, and system designs, efforts at formal, mechanical proof frequently find faults and deficiencies that have not been detected by other means … Applied to rigorous arguments within mathematics, however, efforts at mechanized proof nearly always suggest at most the need to remedy matters that a mathematician would regard as basically trivial, such as typographic errors or failures to state the full range of conditions necessary for a theory to hold.

He continues this way:

---

[11] Given the nonexistence of formal proofs and the languages of such previous to Frege, it's a subtle historical question how formalized proofs so easily became norms of ordinary mathematical proof. My (brief) answer is that this is due—in some part—to a confusion. The earlier vision of treating ordinary mathematical proofs as missing steps, and utilizing tacit assumptions was taken to be vindicated by the invention of formalized languages in which mechanically-recognizable proofs (in accord with classical logic) could be constructed.

> Research for this book has been unable to find a case in which the application of mechanized proof threw doubt upon an established mathematical theorem, and only one case in which it showed the need significantly to modify an accepted rigorous-argument proof.

I should add that failing to recognize the full range of conditions necessary for a theorem to hold *can* turn out to be very significant. The point is that *in these cases* it has not.

What needs to be explained, therefore, is *how* the mathematician's apparent attention and focus on the surface features of ordinary mathematical proof enables such widespread and nearly universal compliance with the formalization norm.[12]

---

[12] Rav (2007, 317) makes a suggestion that might be seen as addressing this question. He writes:

> Like all human knowledge, mathematical knowledge is fallible …. Due to the high precision with which mathematical assertions are stated, with technical terms always precisely defined, errors are minimized and members of the mathematical community normally weed out flaws in proofs through reciprocal criticisms. … [P]rofessional mathematicians spot mistakes in a proof due to their special aptitude, sharpened by the process of lengthy training and experience.

Maybe so, but is it really believable that those professionals, engaged in programs, hardware designs, and system designs, are so much sloppier than mathematicians? (Rhetorical question.)

This is the successor problem to Kant's problem. Any explanation of how mathematicians recognize (and historically, recognized) successful ordinary informal proofs must explain how that recognition conforms to formalist norms, in particular, to the classical logic embodied in the earliest formal systems.[13]

7. *Can mathematicians directly sense validity without a (proof-theoretic) medium to sense it through (or via)*? The failure of formalist approaches to solve Kant's problem invites the thought that we have been trying to locate the recognition of the correctness of informal mathematical proofs in the wrong place. Such recognition must be enabled by the surface properties of such things. In this sense, the old traditional view that filling in details enhances one's sense of the correctness of proofs is just *wrong*. But what surface qualities of an informal mathematical proof can one's recognition of its correctness be due to? Well, pretty obviously, all that's available on the surface are the meanings of the terms that appear in the proof: the various mathematical concepts so ingeniously employed by mathematicians. Since syntactic manipulations are pretty much absent,[14] it must be the mathematician's appreciation of the

[13] Things change in the twentieth century: mathematics breaks free from classical logic. In large measure this is because so breaking free only becomes possible when logic has finally been made explicit. (We can now *see* that there are rules we can break, if we want to.) See Azzouni (2006, chapter 6) for further details on this.

[14] This may be a serious exaggeration. Computation—which usually *does* amount to step-by-step mechanical manipulation—is ubiquitous in mathematical proof. But the professional mathematician's response to its appearance is telling, and supportive of this line of thought. It isn't that he looks over such a computation, and just "sees" that it's valid; rather, he either checks (some of) it, or he takes (some of) it on *trust*.

meanings of mathematical terms coupled with his recognition of truth and validity that does the job.[15]

This may look atavistic. After all, wasn't it Kant who (long ago) showed us that reliance on *meanings* won't provide an explanation of the nature of (informal) mathematical proof? Well, a lot has changed with respect to theoretical approaches to meaning since Kant's day. Back then, the notion was restricted, pretty much, to an analysis of concepts almost along the model of Lego blocks. A concept—corresponding to a word—was taken to be composed of parts, and recognizable as so made up. This was why Kant's examples of a priori synthetic truths, such as "7 + 5 = 12," were successful *counterexamples*. It was presumed recognizable that the concept of "+," for example, isn't part of the notion of "12." The contemporary notion of meaning that Rav alludes to, however, has additional resources. This is indicated by the fact that its formal counterpart is Tarski's model theory, with its accompanying (and defined) notion of validity. Rav (2007, 312) writes (italics his): "… the implication relation … with meanings (relative to a theory) to justify claimed consequences, can be considered as the *informal* counterpart of the semantic model-theoretic perspective."

8. *Reasons to doubt that mathematicians directly perceive the validity of informal proofs via the meanings of the terms in those proofs*. It's clear that an approach—to the mathematician's recognition that an ordinary mathematical proof is a good one—that takes this to occur by his perception of the validity of such a proof (based on the meanings of the terms occurring in it) *is* responsive to the successor problem. This is because the meanings of the terms occurring in an ordinary mathematical proof, and the perception of validity in the steps of that informal proof can be taken to license the existence of a formal proof in one or another artificial language (by

---

[15] See Rav (1999), (2007).

the metalogical considerations raised at the end of section 5). The problem is more basic: it's that the notions invoked require *magical insight* on the part of the professional mathematician. Validity is *noneffective*. This is most visible in the model-theoretic semantics that Rav (2007, 312) alludes to. *Criterial* of an inference, S ⇒ T, being valid is that for *any* model, if S is true in that model then T is true in that model as well. To mildly understate the point, this is rarely recognizable *model-theoretically* (by an examination of, say, *every* model). Under certain (rather unusual) circumstances, the models in which S holds can be characterized in such a way that it can be proven model-theoretically that T holds in them as well. But, in general, it's only via proof-theoretic tools that such validities are recognized.[16] Even if the phenomenology of the experience of understanding an informal mathematical proof gives the impression of the direct perception of validities via the meanings of the terms occurring in that proof, *this cannot be how it is actually being done*.

This is a (purely) *epistemological* issue that's being raised: how is the status of an informal mathematical proof recognized? Rav (2007, 313) provides a "typical text-book proof" starting with a lemma that there is a bijection between $\mathbb{R}$ and the open interval (0, 1):

> Let $S$ consist of all binary sequences $(s_n^{(k)})_{n\in\mathbb{N}} = (s_1^{(k)}, \ldots, s_n^{(k)}, \ldots)$, where by definition, for each $k$ and $n\in\mathbb{N}$, $s_n^{(k)} = 0$ or 1. Define now a new sequence $\Delta = (\delta_n)_{n\in\mathbb{N}}$ by setting $\delta_n = s_n^{(n)} + 1 \pmod 2$. Since for each $n\in\mathbb{N}$, $\delta_n \neq s_n^{(n)}$, it follows that the diagonal sequence $\Delta$ is not in $S$. Hence the totality of all binary sequences is non-denumerable. This in turn

[16] Gödel's *completeness* results are *celebrated* too. There's a reason for this.

implies that the set of all real numbers in the interval (0, 1) is non-denumerable, and hence, in view of the lemma, $\mathbb{R}$ is non-denumerable. Q.E.D.

Rav (2007, 313) claims, among other things, that the "part of the proof claiming that Δ is not in *S* is not a 'proof sketch' that indicates a derivation,"[17] that the structure of this proof "does not depend even implicitly on a deductive calculus," and that "it depends on an understanding of the terminology, i.e., of the meaning of the terms used in that claim and on the background knowledge that $1 + 1 = 0 \pmod 2$)." In one sense, he's absolutely right about the phenomenology: the professional experiences his understanding of this proof as due to the meanings of the terms appearing in the proof, and to the truth of what the sentences say as a result.[18]

There is an important aspect of meaning—as it's ordinarily understood—that may seem relevant to this. Knowing, say, what a rose is, and that (therefore) the object in front of me is a rose, I may be able to draw (not deduce!) the conclusion that there is one red rose. I *see* that at

[17] My uses of "indication" (in "derivation-indicator" and the like in Azzouni (2004), (2006)) may have been unfortunate in the impression that they can give that the "indicating" of a derivation by an informal proof is supposed to be phenomenologically visible to the working mathematician. This was never my intention, as the discussions in Azzouni (2004), (2006) make clear.

[18] Although Rav officially directs this example of a typical informal mathematical proof against what he takes to be my views about the matter, one will find in Azzouni (2005, 19) an example of (part of) a proof from MacLane et al. (1967) that's invoked to illustrate these same points.

least one red rose exists, and I can assert this (in part) by virtue of what "rose" and "red" mean, and thus (in part) by my ability to apply these terms correctly to things.[19]

Given *abstracta* (plus a—broadly speaking—perceptualist epistemology to describe our access to those *abstracta*), we *could* characterize the understanding of the proof given above in these terms: one *sees* that $\Delta$ isn't in $S$. The terms used in the proof refer to objects, and given certain properties of those objects—ones that the reader can be brought to see—he recognizes certain truths as a result. If this construal is ruled out—on the grounds that even if *abstracta* exist, no such epistemological story is acceptable—then the relevance of application-aspects of our understanding of (ordinary-language) meaning is ruled out too. That leaves meaning only the role of providing implicit constraints that hold of what the meanings are of.

9. *Subdoxastic recognition of indicated derivations?* Let's, therefore, set aside an approach to the recognition of successful ordinary informal mathematical proofs that operates only via the surface properties of those proofs. Perhaps what's needed is a revitalization of the formalist approach by letting it borrow a methodological page or two from linguistics. Notably, the view *there* is that the subdoxastic grammar of a natural language contains all sort of apparatus and categories that are invisible to the ordinary speaker-hearer. He can recognize (within certain performance parameters) grammatical sentences, and distinguish them from

[19] This aspect of meaning is to be captured by the analysis of "application conditions," or something like that. Strictly speaking, it's not to be captured (or not very well, anyway) by (neo-Davidsonian) truth-conditional approaches to meaning, ones that handle it not by an analysis, but instead via deference to terms (in the language within which the truth conditions are given) that are taken to have the same truth conditions, meanings, or whatever.

strings of vocabulary items that aren't sentences. But the source of his competence in this is unavailable to him introspectively, and can only be discovered empirically.

In a similar way, we can treat formalization projects as empirical ones: attempts to discover the subdoxastic properties of mathematical proof. This goes a long way towards responding to (a) and (b), as stated in section 2. The subjective experience of the qualities of proof as arising from meanings is similar to the ordinary speaker-hearer's impression that the grammar of his sentences turns on rules of thumb that he's (roughly) aware of. The appearance that his reasoning (during the course of an ordinary informal mathematical proof) is topic-specific is an introspective illusion; in point of fact, as one or another formalist project reveals, all reasoning is topic-neutral in nature. It would be in this (empirical) sense, and in this sense only, that an ordinary informal mathematical proof would "indicate" one or another formal analogue.[20]

I've become convinced that this can't work.[21] The main reason is that the linguistic model turns on a psychological—ultimately neurophysiological—assumption: the real description of the grammar (as opposed to the ordinary speaker-hearer's impression of what it is) is subdoxastically (but nevertheless *really*) psychologically instantiated. It plays a genuine *causal*

---

[20] I'm skipping details. Gödel incompleteness, for example, requires that there not be a single formal language but families of such. And, given the developments of twentieth-century mathematics, what's also required are *isolated* families of such, based on alternative logics (and other differences). See Azzouni (2006). Also see Azzouni (1994), where a version of the "algorithmic-systems" approach to mathematical proof was first given.

[21] See Azzouni (2004), (2005) for details.

*role* in the ordinary speaker-hearer's recognition of grammaticality. In doing so, this approach opens itself to empirical study and test.

Indeed, empirical reasons for believing in such capacities rely on nativistic evidence: the early and automatic acquiring of ordinary language, this capacity being largely modularized (at least with respect to general intelligence), there being a small window of time during which language acquisition abilities are operative, and so on. No similar evidence exists for the corresponding mathematical thesis. Although there are mathematical savants and prodigies, their capacities arise in quite topic-specific ways—often with respect to certain narrow families of concepts (e.g., in one case, *just* prime numbers)—and not with respect to mathematical reasoning generally. Furthermore, such neurophysiological evidence as we are beginning to gather, seems to indicate something similar. Whatever it is that's neurophysiologically embodied in us by way of mathematical capacities, it doesn't seem to be topic-neutral reasoning abilities, but far more specific (and, quite frankly) gerrymandered capacities with small groups of concepts (e.g., "one," "two," and "three") that we can automatically rely on, and that are creatively (and perhaps idiosyncratically) built upon as an individual's abilities in mathematics grow. (For a taste of the on-going research in this area, see Dehaene 1997.)

10. *The inference package approach to mathematical proof.*[22] Focusing on the surface properties of ordinary mathematical proof is the right strategy for attempting to characterize how mathematicians grasp those proofs. But meaning, as noted, isn't a useful tool. This can already be seen by the fact that the concepts (the "meanings") that mathematicians employ aren't typical. That mathematics is so usefully applied everywhere in ordinary life and in the sciences shouldn't

[22] What follows is drawn from Azzouni 2005, where the inference package approach was first discussed.

blind us to the fact that nevertheless the concepts utilized in ordinary mathematical proof are (from the point of view of ordinary language) peculiar and specialized. They seem to refer to *very strange things*. One needs to get a grip on why such and not the terms of ordinary life are (and need to be) employed by mathematicians. Blandly noting that mathematical terms are "meaningful," and that mathematicians (somehow) rely on those meanings in order to recognize good informal proofs isn't helpful.

As already indicated, one important clue about what's actually going on is the apparent topic-specificity of mathematical reasoning. But another equally important clue is the surprisingly ubiquitous role of (various kinds of) diagrams that occur in (even contemporary) mathematical proof. What the presence of diagrams points to is that crucial to the understanding of a proof is the employment of various psychological techniques. If one, for example, imagines triangles on a plane, and visualizes moving them around (or expanding them uniformly), and compares those same operations on triangles when they appear on the surface of a sphere or on an ellipsoid, one will *visually experience* how the curvature of such surfaces affects the figures. In doing so, a capacity is exhibited of exploring the *simultaneous effect* of several assumptions embodied in the properties of the surfaces imagined, and seeing how those assumptions affect (in particular) what sorts of triangles are possible on a surface, and how they shift in their properties when moved or deformed in various ways.

This is a perspicuous example of the utilization of an *inference package*—a capacity to recognize the implications of *several* assumptions by means of the representations of objects wherein those several assumptions have been knit together (psychologically). I also claim that the employment of inference packages shows up everywhere in mathematical practice.[23] What an

[23] I speculate that inference packages are "off the shelf" psychological capacities with "cognitive

inference package allows a mathematician to do is to *reason* about a subject matter compatibly with a formal mechanical proof. His reasoning, however, is topic-specific[24] and the various assumptions operative in that reasoning function together in a way that makes them phenomenologically invisible to him.

Inference packages are brought to bear on mathematical proofs in such a way that it can appear that they operate via the meanings of the terms appearing in those proofs. But the meaning-properties of mathematical terms will not illuminate what's happening here. To see this, notice that mathematical practice treats mathematical concepts as quite open-ended in their contents. One reason for this is that a set of mathematical concepts M aren't fixed in the inference packages that can be brought to bear on them. And what *this* corresponds to is that other concepts C can be brought to bear on the implications taken to follow from M and its truths. But such supplementations of a set of concepts (and truths governing them) with other

---

homes" elsewhere, in the various capacities we can bring to visualization, as well as to our other sensory experiences. Such inference packages are *grafted*—often creatively—into mathematical contexts by mathematicians. Diagrams are one tool for priming the use of inference packages during attempts to master an informal mathematical proof.

[24] Topic-specificity arises, to begin with, because inference packages seem directed towards specific subject-matters. For example, our capacity to visualize shapes seems intrinsically geometrical. In practice, the mathematician can reinterpret those inference packages similarly to how geometrical concepts can be reinterpreted to apply to something non-geometrical. Although these reinterpretation techniques, in effect, amount to treating inference packages as topic-neutral, that isn't the experience of the employment of them (even when they have been reinterpreted for applications).

such concepts (and truths governing them) although required to preserve what can be shown to follow from the concepts of M alone, are not conservative. Additional new results become possible.

It's terminologically reasonable to call this the "holism" of mathematical concepts. One example is the widely-noted non-conservative effect of the truth predicate, when it's allowed to occur within the induction schema in Peano arithmetic. Another is the non-conservative effect of the introduction of negation into the positive propositional calculus (e.g., with the resulting provability of Peirce's law—see Church (1956, 161)). But—from the mathematician's point of view—strict non-conservativity isn't the issue. Instead, ease of provability is (the resulting shortness and/or perspicuity of an informal proof). Expanding, for example, real transcendental functions uniquely into their complex analytic extensions yields a wealth of simple and easily applied tools for integrating the real functions one started with; this is valuable regardless of whether the resulting introduction of new concepts is conservative or not.

11. *The successor problem again*. This isn't the place to provide a full defense of the inference-package picture of informal mathematical proof, especially since I've attempted that in Azzouni (2005). What's crucial to notice now is that this approach to informal mathematical proof seems to face the successor problem to Kant's problem. Inference packages, operating as they do so near the surface of informal mathematical proof, aren't explicit about the logic that governs the inferences that they facilitate.[25] How, then, do the proofs that they enable the understanding of so uniformly correspond to formal derivations?

[25] It's a striking fact that the project of uncovering the logic of mathematical reasoning arises *so late* in the mathematical tradition. Logic—as a subject matter—although not *quite* as old as mathematics, is certainly old enough for the inability of syllogistic logic to capture mathematical

12. *A content-containment notion of implication*. The first question to answer is what someone thinks (ordinary) implication comes to when not only is the logic tacit, but even the idea of *there being a logic involved at all* is tacit as well—e.g., how is someone like *Descartes* supposed to think of implication? One answer, of course, is that if the assumptions are true, then the conclusion *must* be true as well. But there's more involved in the phenomenology than *that*; there is also the intuition that this necessity is due to what the sentences in question *say*, their "content." An implication T of a sentence S can't say more than S does, or something different from what S says; it can only say the same or *less*.[26] It's important to realize how frail this notion is—how little work it actually does when we concretely experience one proposition following from others. It *does* work quite well (intuitively) with conjunctions: One senses that "John is running" follows from "John is running and Sally is running" because the latter sentence just says what the former sentence does *plus* something else. But it works badly even with otherwise intuitively compelling applications of modus ponens: "If John is running then Sally is running, and John is running" doesn't (intuitively) *say* that Sally is running. Still, many have the sense

---

reasoning to have been clear in its significance *sooner*. (As noted, this was first seen by Kant.) One explanation for the (historical) lateness of the project is that the inference packages crucial to mathematical reasoning bury so completely (phenomenologically speaking) the assumptions they rely on that it becomes hard to even see that they rely on (logical) assumptions *at all*. Uncovering assumptions—of any sort—in an informal method of proof is, after all, a mathematically creative, and quite difficult, activity.

[26] This idea is in the same family as the modern-philosophy idea of the content of a predicate being contained in the content of the subject term (in the case of analytic sentences)—a notion that was discussed briefly in the opening pages of this paper.

that the first sentence "tacitly" or "implicitly" says that. But thinking in terms of tacit "sayings" won't carry us very far. Such content-containment intuitions are decidedly *not* transitive: even quite short chains of (otherwise recognizably valid) reasoning can obliterate any sense that content-containment is the source of one's sense of "validity"—e.g., after several applications of modus ponens. Note the point: the obliteration of the sense that an inference accords with content-containment isn't an (intuitive) recognition that it isn't content-containment that's behind one's sense of implication. Rather, one thinks that it's still a matter of content-containment—but only insofar as an implication brings out explicitly what the premises only say "tacitly."

It's clear: content-containment intuitions don't arise from a systematic (and careful) look at the actual inferences we engage in, and what—if anything—is being "carried along" when we see that one proposition follows from others. Rather, they arise from a kind of general impression that an implication of something said cannot *add* content to that something.[27]

An important piece of evidence that something like content-containment intuitions do play a role in the ordinary sense of what implication comes to is the widespread intuitive response to tautologies and contradictions. These—according to the contemporary model-theoretic notion of implication—are the endpoints of a continuum: any statement implies a tautology, and a contradiction implies any statement. These are the appropriate views—I stress—when implication is construed model-theoretically. But *both* contradictions and tautologies are

[27] Thus, many are uncomfortable with the idea that "John is running or Sally is running" is implied by "John is running." The problem *isn't* that the ordinary language "or" is exclusive, or that there is a violation of a Gricean implicature. It's that we are uncomfortable engaging in an inference that *so blatantly* adds content to something said, and yet is still characterized as an *implication* of that something.

experienced as "failing to say anything," as failing to have "content"; and so they aren't (intuitively) seen as having *any* implication-relations to other statements.[28] This impression, in the case of contradictions, is important because it's the first (obvious) place where informal mathematical proof can be seen to deviate from the dictates of ordinary intuitions with respect to implication. Proof by contradiction doesn't sit very well (as a practice) with the perspective that a contradiction doesn't say anything—as Quine (1953, 5) noted so long ago.

A related piece of evidence for the role of content-containment intuitions in the ordinary notion of implication is the strong (intuitive) importance of relevance to ordinary inference. What a sentence S "says," what it's "about," makes its implications look topic-specific. And therefore, how can a sentence T that is about "other things" nevertheless be an implication of S? (Notice that relevance appears to be a necessary condition for the satisfaction of content-containment.) It's a testimony to the intellectual integrity of ordinary informal mathematical proof that the shock of mathematical reasoning so often violating relevance intuitions never seems to have threatened the practice itself in any way.

It can be asked: what exactly is it about ordinary mathematical proof that allows it to so clearly violate relevance (and consequently, content-containment) intuitions? One thing is that such intuitions turn on the impression of the subject matter the sentences are about; but these intuitions can be violated when proofs involve constructions of types of mathematical objects out of other types of objects that we are prone to think about in very different ways. (It's this that

[28] Actually, I've found intuitions in this area to be a little more complicated than this—although in a way that still supports the importance of content-containment intuitions. Some people experience "John is running and it is not the case that John is running," as saying something that implies (for example) that "John is running," although not anything about (say) *Sally*.

makes the equivalence of the various formulations of the axiom of choice, for example, so surprising.) Another factor is that reasoning via "contradictions" and "tautologies" (logical truths and falsehoods) is indispensable to systematic reasoning. But, simultaneously—at least in the context of the tacit reasoning practices in place in mathematics for over two thousand years—such reasoning facilitates violations of relevance and content-containment.[29]

This is why the recent official acceptance of the bundle of classical logical notions, truth-functional "if … then," "implication" (construed model-theoretically), relevance violating inference rules, and so on, were all (nevertheless) so compatible with already-existing informal mathematical reasoning. The latter was already operating (implicitly) with such notions; it was already violating content-containment intuitions and relevance intuitions, *while nevertheless mathematicians continued—when they could—to rely on such intuitions to guide their understanding of proofs*. Notice, by the way, a certain resulting instability in the content-containment notion of implication. It's okay as a tool in one's recognition of implications—as far as it goes—as long as it remains tacit. It breaks down, however, if it's made explicit. And, as the unappealing eccentricities of the relevance tradition in logic make clear, highly dramatic changes must occur in one's logic if (aspects of) relevance, and content-containment, intuitions are to be retained in one's inferential practices.

13. *How we retain content-containment intuitions in the face of their failing to underwrite much of step-by-step mathematical reasoning*. It's already been noted that in the case of certain intuitively-valid inferential steps, such as uncontestable cases of modus ponens, we intuitively

[29] The same phenomenon, incidentally, is operating in the background of the various popular "slingshot" arguments against facts or, more generally, against a fine-grained individuation of what sentences are "about." See Neale 1995 for a useful discussion of these.

protect content-containment intuitions by taking refuge in the idea that something is tacitly said by premises that's made explicit by inference. But there are common patches of mathematical reasoning in many proofs, for example numerical and algebraic computation, where intuitions about "what follows" play little or no role. Instead, one is intensely aware that one is applying mechanical rules of one sort or another, and that these must be implemented without making a mistake. Furthermore, when one checks for mistakes in these cases, it's not done by recognizing that the result "doesn't follow."

Nevertheless, it's thought that such (mechanical) rules can (or should) be *justified*. And what this is taken to mean is that how such rules satisfy the requirements of inference should be indicated (by proof). Then we know they satisfy the requirements of implication even though we don't experience them as doing so. We *see that* they are "truth-preserving," even though we don't *experience them* (in the moment of going through them) as being so. These aspects of reasoning, coupled with our so often treating inference as bringing out what's tacitly said, together protect our intuitions that inference is a matter of content-containment against refutation by counterexamples from our actual reasoning processes. For this reason, content-containment intuitions play a strong role in our (tacit) views about what logical inference *is* even if those intuitions don't play (much of) a role in how that reasoning itself is recognized (when, i.e., we recognize an inference to be "valid").

14. *The bridge between informal mathematical proof and formal artificial language derivations*. With the role of this intuitive notion of content-containment delineated, we can now see what connects ordinary mathematical proof to formal derivations. First, and most important, there is a practice, already in place in informal mathematics, of externalizing the content of a concept (a term) into axioms. In this way, that content *isn't* treated as part of the inferential

process (e.g., as number-theoretic presumptions are so treated in inferences licensed by ordinary-language numerical quantifiers), but instead as assumptions that one (officially) uses to draw inferences *from*. Second, the classical (and intuitionistic) connectives and quantifiers are generated by content-containment conditions.[30] Notice that what's crucial isn't that, in the very steps recognized as correct steps in an informal mathematical proof, the mathematician *sense* content-containment at work. What's required, rather, is that content-containment be presumed to be a *requirement* on logical inference: that logical inference not *add* content. This suffices for the content-containment conditions on the connectives and quantifiers. Third, in place is the assumption that any proposition can be negated, and that iterations of negations are cancelling. This forces a collapse of all the connectives (and quantifiers) to classical—not intuitionistic—ones. The crucial point is this: *as long as the logic itself remains tacit, there is no need (in the vernacular) for idioms that explicitly embody these logical connectives*. All that's needed is that the inferences allowed be recognized as conforming to these content-containment conditions on logical inferences (and that the inferential tools employed—e.g., computational devices, visual aids, etc.—themselves be in accord with these content-containment conditions). In addition, various terms, "if … then," "and," and so on, can be drafted for use in mathematical reasoning, even if they have additional content, provided that content is ignored so far as (mathematical) inference is concerned. This explains how mathematical practice can conform to a logical

[30] E.g., following Koslow (1992), a conjunction is defined by the standard conjunction elimination rules plus a minimality condition that the conjunction (of two propositions A and B) is the *weakest* proposition that obeys the standard conjunction elimination rules. That is, there is no content in a conjunction *other than that* which licenses the elimination rules. (This is why conjunctions—so defined—obey the standard conjunction introduction rules.)

formalism that's nevertheless literally inexpressible in the vernacular. And this is how a solution to the successor problem is managed.

One worry may be this. If content-containment conditions govern the (implicit) logical apparatus—the reasoning that's in accord with conjunction, disjunction, quantifiers, and so on—how can this possibly be compatible with the clear *violations* of content-containment constraints on implications that nevertheless occur in ordinary mathematical proof? The answer, briefly, is that it's one thing for the logical items by which reasoning occurs to respect content-containment in the sense that those items themselves are defined in such a way that *they* don't contribute anything "extra" to the reasoning process; it's quite another for that reasoning process itself to conform to such content-containment intuitions. The former can be satisfied while the latter isn't.[31]

And with this characterization of logical inference, an answer to the successor problem has been given. On the one hand, content-containment intuitions regarding the logical transitions (in informal mathematical reasoning) directly correspond to the logical operators in classical first-order languages. On the other, inference packages are recognized as satisfying these conditions on logical transitions with respect to background assumptions that—when made explicit—are treated axiomatically.

---

[31] An example. Following Koslow (1992), a negation of a proposition A is the *weakest* item that together with A implies an arbitrary proposition B. Here, negation respects content containment—*it* contributes nothing to the reasoning process over and above its licensing an inference to B—while at the same time an inference from A and its negation to B violates intuitions about the implications of A (and its negation) respecting content-containment.

From the mathematical-practice point of view, it is *luck* that we have the rich and flexible set of inference packages that we have (one particularly *trustworthy* set of such is largely geometric—and it's applied by increasingly sophisticated reinterpretations to ever-new kinds of abstracta). Nothing prevents—in principle—an inference package from being misleading (and useless for mathematical practice). A clear example are the kinds of (subdoxastic) methods that generate our statistical intuitions—intuitions that are, pretty much, untrustworthy. (As those who wisely invest in gambling casinos are well aware.)

15. *One loose end*. For those who see ordinary informal mathematical proof as occurring via the deployment of concepts, proof seems historical or "time-dependent" (Rav, 2007, 294). In mathematics, new concepts are invented regularly (incessantly, one could say); and if these are seen as part of the *proof-procedures*, as part of the inferential apparatus, then clearly the nature of "proof" changes with each new concept introduced into mathematical practice. If, on the other hand, mathematical concepts are just codings for various implicit assumptions, then proof won't be seen as "time-dependent." In past work, I've spoken of traditional informal mathematical proof as atemporal, and for just this reason: the content of such concepts is treated *by ordinary (and traditional) mathematical practice* as externalizable into postulates or assumptions or axioms, or whatever, *as soon as that content ceases to be tacit*. Ordinary mathematical practice, that is, doesn't treat such content as *inferential content* when that content has become explicit—as soon, that is, as mathematicians become aware of it, and endeavor to characterize it.

16. *Last remarks*. It's unfortunate, I think, that the ideas developed in Azzouni (2005) occurred to me well after Azzouni (2006) was in press. During my writing of Azzouni (2006), I continued to struggle with the successor problem, and that's why I continued to use the phrase "derivation-indicator view of mathematical proof" to characterize my view of the relationship of

ordinary mathematical proof to that of formal derivations. This nomenclature acknowledged the normative status of formal derivations; but it also indicated my thinking that ordinary proof being in accord with those norms couldn't be a mere coincidence—that there had to be *something* occurring in ordinary mathematical proof that *explained* its conformity with formal derivational norms. Unfortunately, during the writing of Azzouni (2006), I only had the resources of syntactic pattern-recognition, and the like, to draw on, and so I kept falling (against my will) into a view that mathematicians had to be engaged in *something* like sophisticated syntactic pattern-recognition while perusing informal mathematical proofs, so that they would be sensitive (without realizing it) to a background of *nonexistent* formal derivations. Topic-specific inference packages and content-containment conditions on the logical tools used for inference were responses to the successor problem that only occurred to me during the writing of Azzouni (2005), and so are absent from explicit discussion in Azzouni (2006), except for an inserted footnote.

Acknowledgements

My thanks to Jeremy Avigad, and to a referee, for helpful comments on an earlier version (consequently modified) of this paper. My thanks also to Ian Dove for creating the occasion for this paper's existence.

Bibliography


Azzouni, Jody: 1994, Metaphysical myths, mathematical practice: The ontology and epistemology of the exact sciences. Cambridge University Press, Cambridge, England.

Azzouni, Jody: 2000a, Applying mathematics: An attempt to design a philosophical problem. The Monist 83: pp. 209-227.

Azzouni, Jody: 2000b, Stipulation, logic, and ontological independence. Philosophia Mathematica (3) 8: pp. 225-43.

Azzouni, Jody: 2004, The derivation-indicator view of mathematical practice. Philosophia Mathematica (3) 12: pp. 81-105.

Azzouni, Jody: 2005, Is there still a sense in which mathematics can have foundations? In Giandomenico Sica (ed): Essays on the foundations of mathematics and logic, Polimetrica S.a.s., Milan, Italy, pp. 9-47.

Azzouni, Jody: 2006, Tracking reason: Proof, consequence, and truth. Oxford University Press, Oxford.

Azzouni, Jody: forthcoming, The compulsion to believe: Logical inference and normativity. Protosociology.

Carston, Robyn: 2002, Thought and utterances: The pragmatics of explicit communication. Blackwell Publishing, Oxford.

Church, Alonzo: 1956, Introduction to mathematical logic, Volume 1. Princeton University Press, Princeton, New Jersey.

Dehaene, S. 1997. The number sense. Oxford University Press, Oxford.

Hacking, Ian: 1973, Leibniz and Descartes: Proof and eternal truths. In Historical ontology. Harvard University Press, 2002, Cambridge, Massachusetts, pp. 200-13.

Katz, Jerrold J.: 1998, Realistic rationalism. The MIT Press, Cambridge, Massachusetts.

Koslow, Arnold: 1992, A structuralist theory of logic. Cambridge University Press, Cambridge, England.

Kreisel, G.: 1969, The formalist-positivist doctrine of mathematical precision in the light of experience. L'Âge de la Science 3: pp. 17-46.

MacKenzie, Donald: 2001, Mechanizing proof: Computing, risk, and trust. Cambridge University Press, Cambridge, Massachusetts.

MacLane, S., and G. Birkhoff: 1967, Algebra. The Macmillan Company, New York.

Neale, Stephen: 1995, The philosophical significance of Gödel's slingshot. Mind 104: pp. 763-825.

Quine, W.V.: 1953, On what there is. In From a logical point of view, Harvard University Press, Cambridge, Massachusetts, pp. 1-19.

Rav, Yehuda: 1999, Why do we prove theorems? Philosophia Mathematica (3) 7: pp. 5-41

Rav, Yehuda: 2007, A critique of a formalist-mechanist version of the justification of arguments in mathematicians' proof practices. Philosophia Mathematica (3) 15: pp. 291-320.