

Logic Naturalized

John Woods¹

Contents

I- The Mathematical Turn in Logic	3
II- The Naturalistic Turn in Logic.....	12
Notes.....	23

I- The Mathematical Turn in Logic

It is no secret that classical logic and its mainstream variants aren't much good for human inference as it actually plays out in the conditions of real life - in life on the ground, so to speak. It isn't surprising. Human reasoning is not what the modern orthodox logics were meant for. The logics of Frege and Whitehead & Russell were purpose-built for the pacification of philosophical perturbation in the foundations of mathematics, notably but not limited to the troubles occasioned by the paradox of sets in their application to transfinite arithmetic. Logic from Aristotle to then had been differently conceived of, and would be decked out to serve different ends. The Western founder of systematic logic wanted his account of syllogisms to be the theoretical core of a general theory of everyday face-to-face argument in the courts and councils of Athens, and more broadly in the *agora*. Aristotle understood that in contexts such as these premiss-conclusion reasoning² is an essential component of competent case-making. He thinks that when a conclusion is correctly derived from a set of premisses there exists between it and them a truth-preserving relation of *consequence*. This is a distinctively Greek idea, and one that has resonated from then to now. It is the idea that even *everyday* case-making argument is deductively structured when good.³ Deductivism is with us still, albeit often in rather watered-down ways. Even so, the nonmonotonic consequence relations of the 20th and 21st centuries, virtually all of them, are variations of variations of classical consequence. They are, so to say, classical consequence twice-removed.⁴

Whatever we might make of this lingering fondness for deductivism, the logic of premiss-conclusion inference assigns a pair of important tasks. One is to describe the conditions under which premisses *have* the consequences they do. Call this the logician's "consequence-having" task. The second is related but different. It requires the logician to specify the conditions under which a consequence of a premiss-set is also a consequence that a human reasoner should actually *draw*. Call this the "consequence-drawing" task. The dichotomy between having and drawing is deep and significant. Consequence-having occurs in logical space. Consequence-drawing occurs in an inferer's head and, when vocalized, mainly in public space.

This gives us an efficient way of capturing a distinctive feature of modern mainstream logics. They readily take on the consequence-having task, but they respond ambivalently to its consequence-drawing counterpart. This ambivalence plays out in two main ways. I'll call these "rejectionism" and "idealization" respectively. In the first, the consequence-drawing task is refused outright as an unsuitable encumbrance on logic.⁵ Such gaps as there may be between consequence-having and consequence-drawing are refused a hearing in rejectionist logics. However, according to the second, the consequence-having problem not only receives a hearing in logic but derives from it a positive solution. Logic would rule that any solution of the consequence-having problem would *eo ipso* be a solution to the consequence-drawing problem. The desired correspondence would be brought about by *fiat*, by the stipulation that the "ideally rational" consequence-drawer will find that his rules of inference are wholly provided

for by the truth conditions on consequence itself. By a further stipulation, the conditions on inference-making would be declared to be *normatively binding* on human inference-making on the ground.⁶ There is a sense, then, in which an idealized logic closes the gap between having and drawing. Even so, it is clear that on the idealization model the gap that actually does close is not the gap between consequence-having and consequence-drawing on the ground, but rather is the gap between having and idealized drawing. In that regard, the idealization model is in its own right a kind of quasi-rejectionism. All it says about on-the-ground consequence-drawing is that the rules of idealized drawing are normative for it, notwithstanding its routine non-compliance with them. Beyond that, inference on the ground falls outside logic's remit. It has no lawful domicile in the province of logic.

Aristotle is differently positioned. On a fair reading, what he seeks is a new purpose-built relation of deductive consequence-having - *sylogistic* consequence - whose satisfaction conditions would coincide with the rules of consequence-drawing not under idealized conditions but rather those actually in play when human beings reason about things. Accordingly, Aristotle's is a genuinely gap-closing logic, but without the artifice of idealization. The nub of it all is that Aristotle's constraints bite so deeply that for any arbitrarily set of premisses the likelihood that there would be *any* syllogistic consequences is virtually nil; and yet when premisses do have syllogistic consequences, they are at most *two*.⁷

This is a considerable insight. Implicitly or otherwise, Aristotle sees that the way to close the gap between having and on-the-ground drawing is by reconstructing the relation of consequence-having, that is, by making consequence-having *itself* more inference-friendly. To modern eyes it is quite striking as to how Aristotle brought this about. He did it by taking a generic notion of consequence (he called it "necessitation") and imposing additional conditions on it that would effect the desired transformation. This would produce - in my words, not his - the new relation of syllogistic consequence - a proper subrelation of necessitation - whose defining conditions would make it nonmonotonic and paraconsistent, and at least some adumbration of relevant and intuitionist in the modern senses of those terms.⁸ It is well to note that these inference-friendly improvements derive entirely from readjustments to consequence-having, and they put to no definitional work any considerations definitive of face-to-face argumental engagement. In other words, although inference happens in the head, Aristotle's provisions for inference-*friendliness* take hold in logical space.

When we turn from Aristotle's to modern-day efforts to improve logic's inference-friendliness, we continue to see similarities and differences. As with syllogistic consequence, the newer consequence relations trend strongly to the nonmonotonic, and many of them are in one way or another relevant and paraconsistent as well. Others still are overtly intuitionist. The differences are even more notable. I have already said that, unlike his theory of face-to-face argument-making, Aristotle's syllogistic consequence is wholly provided for without the *definitional*⁹ employment of considerations about inference-making or of the beings who bring them off. In the logic of syllogisms there is no role for agents, information flow, actions, times or

resources. In contrast, modern attempts at inference-friendliness give all these parameters an official seat at the definitional table of consequence-having. Consequence-having is now defined for consequence relations expressly connected to agents, information flow, actions, times and resources. There is yet a further difference to respect. It is that although these modern logics give official admittance to agents, actions and the rest, they are admitted as idealizations, rather than as they are on the ground.

In our own day, a case in point is Hintikka's agent-centred logics of belief and knowledge, in ground-breaking work of 1962.¹⁰ Hintikka's epistemic logic is an agent-centred adaptation of Lewis' modal system S4, in which the box-operator for necessity is replaced with the epistemic operator for knowledge, relativized to agents a . The distinguishing axiom of S4 is $\neg \Box A \rightarrow \Box \Box A$. Its epistemic counterpart is $\neg Ka A \rightarrow Ka Ka A$, where "Ka" is read as "It is known by agent a that ...". We have it straightaway that the epistemicized S4 endorses the KK-hypothesis, according to which it is strictly impossible to know something without realizing you do. Of course, this is miles away, and more, from the epistemic situation of real-life human agents; so we are left to conclude that Hintikka's agents are idealizations of *us*. It is a gap-closing arrangement only in the sense that the behaviour of Hintikka's people is advanced as normatively binding on us.

Hintikka has an interesting idea about how to mitigate this alienation, and to make his logic more groundedly inference-friendly after all. Like Aristotle, Hintikka decides to make gap-closing adjustments to orthodox consequence-having, not just by way of specific constraints on it, but also by way of provisions that make definitional use of what agents say. That is, Hintikka decides to fit his consequence relation for greater inference-friendliness not just by imposition of additional semantic constraints, but also by application of *pragmatic* ones as well.

It is a radical departure. It effects the pragmaticization of consequence-having. I regard this as a turning point for most of the agent-based logics ever since. Logics of nonmonotonic, defeasible, autoepistemic and default reasoning also pragmatize their consequence relations.¹¹ Still, radical or not, it shouldn't be all that surprising a departure. How could it be? What would be the point of inviting even idealized agents into one's logic if there were nothing for them to do there?

Consider for example, the Hintikkian treatment of logical truth. In the orthodox approaches a wff A is a truth of logic if and only if there is no model of any interpretation in which it fails to hold. In Hintikka's pragmaticized logic, A is a truth of logic if and only if either it holds in every model of every interpretation or its negation would be a self-defeating statement for any agent to *utter*. Similarly, B is a consequence of A just when an agent's joint affirmation of A and denial of B would be another self-defeating thing for him to *say*. The same provisions extend to Hintikkian belief logics. Not only do people (in the model) believe all logical truths, but they close their beliefs under consequence. There are no stronger idealizations than these in any of the agent-free orthodox predecessor-logics.

Closure under consequence is especially problematic. In the usual run of mainline approaches, there exist for any given A transfinitely many consequences of it. Think here of the chain $A \vDash B$, $A \vDash B \vee C$, $A \vDash B \vee C \vee D$, and so on - summing to aleph null many in all. Take any population of living and breathing humans. Let Sarah be the person who has inferred from some reasonably manageable premiss-set the largest number of its consequences, and let Harry be the person to have inferred from those same premisses the fewest; let's say exactly one. Then although Sarah considerably outdraws Harry, she is not a whit closer to the number of consequences-had than Harry is. They both fall short of the ideal inferrer's count equally badly. Neither of them approaches or approximates to that ideal in any finite degree. Now that's what I'd call a gap, a breach that is transfinitely wide. It is also an instructive gap. It tells us that giving (the formal representations of) agents, actions, etc. some load-bearing work to do under a pragmaticized relation of consequence-having is far from sufficient to close the gap between behaviour in the logic and behaviour on the ground.

Still, it is important to see what Hintikka had it in mind to do. The core idea was that, starting with some basic but gap-producing logic, the way to close it or anyhow narrow it to real advantage, is to do what Aristotle himself did to the everyday notion of necessitation. You would restructure your own base notion of consequence by subjecting it to additional requirements. In each case, gap-closure would be sought by making the base notion of consequence a more complex relation, as complex as may be needed for the objectives at hand. In other words,

The turn towards complexification: The complexification of consequence is the route of choice towards gap-closure and inference-friendliness.

One can see in retrospect that Hintikka's complexifications were too slight.¹² There is, even so, an important methodological difference between Aristotle's complexification and those of the present day. Aristotle's constraints are worked out in everyday language. Syllogistic consequence would just be ordinary necessitation, except that premisses would be (1) non-redundant, (2) more than only one and (probably) no more than two, (3) none repeated as conclusion or immediately equivalent to any other that does, (4) internally and jointly consistent, and (5) supportive of single conclusions only. These and others that derive from them would provide that the conclusion of any syllogism is either one that should obviously be drawn or is subject to brief, reliable step-by-step measures to make its drawability obvious. This is got by way of the "perfectability" proof of the *Prior Analytics*. (Even it is set out in everyday Greek supplemented by some modest stipulation of technical meanings for ordinary words).¹³

Modern gap-closers have quite different procedural sensibilities. They are the heirs of Frege and Russell, who could hardly in turn could be called heirs of Aristotle. Frege and Russell were renegades. They sought a wholesale restructuring of logic, of what it would be for, and how it would be done. Those objectives and their attendant procedural sensibilities are mother's milk for modern logicians. Logic pursues its objectives by way of mathematically expressible formal representations, subject in turn to the

expositional and case-making discipline characteristic of theoretical mathematics. There flows from this a novel understanding of complexification. In the modern way, complexifications are best achieved by beefing up the mathematical formalizations of a base mathematical logic. Let's give this a name. Let's say that today's preferred route to gap-closure is the building of more mathematically complex technical machinery. In briefer words, inference-friendly logics are *heavy-equipment logics*.

Johan van Benthem has recently written of an idea that gripped him in the late 1980s:

The idea had many sources, but what it amounted to was this: make actions of language use and inference first-class citizens of logical theory, instead of studying just their products and data, such as sentences or proofs. My programme then became to explore the systematic repercussions of this 'dynamic turn'.¹⁴

In the ensuing thirty years, van Benthem and his colleagues have constructed a complex technology for the execution of this dynamic turn. It is an impressive instrument, an artful synthesis of many moving parts. Here is a close paraphrase of its principal author's summary remarks: With the aid of categorical grammars and relational algebra we can develop a conception of natural language as a kind of cognitive programming language for transforming information. This could be linked in turn to modal logic and the dynamic logic of programs, prompting insights into process invariances and definability, dynamic inference and computational complexity logics. In further variations, logical dynamics would become a general theory of agents that produce, transform and convey information in contexts both social and solo. The result is a *dynamic epistemic logic* (DEL), which gives a unified theoretical framework for knowledge-update, inference, questions, belief revision, preference change and "complex social scenarios over time, such as games." The creator of DEL also

would see argumentation with different players as a key notion of logic, with proof just a single-agent projection. This stance is a radical break with current habits, and I hope that it will gradually grow on the reader, the way it did on me. (p. ix)

Indeed, I would be happy if the viewpoints and techniques offered here would change received ideas about the scope of logic, and in particular, *revitalize its interface with philosophy.*" (p. x; emphasis added)

Van Benthem notes with approval the suggestion in Gabbay's and my 2002 paper, "Formal approaches to practical reasoning: A survey,"¹⁵ that the interface with argument may be the last frontier where modern logic finds its proper generality and impact on human reasoning. Again I paraphrase: Over the last decade this insight has developed into a paradigm of attack-and-defend-networks (ADNs) - from unconscious neural nets, to variations that adapt to several kinds of conscious reasoning. This, too, is a highly complex technology, a fusion of several moving parts. As provided for by Barringer, Gabbay and me,¹⁶ the ADN paradigm unifies across several fields, from logic programs to dynamical systems. AD-networks have some interesting technical capacities. They give an equational algebraic analysis of connection strength, where stable states can be found by way of

Brouwer's fixed-point result. When network activity is made responsive to time, logic re-enters the picture, including the development of quite novel modal and temporal languages. "Clearly", says van Benthem, "this is an immense intellectual space to consider." He adds that he "totally agrees" with the ADN "vision, and am happy to support it." (p. 84)

Here, then, are just two of a great many heavy-equipment technologies, specifically adapted to the requirements of argument. They are unifications of partner elements, some of their authors' own contrivance, but in the main having an already established and well understood methodological presence in the several research communities from which they have been borrowed. Both the DEL and ADN approaches carry the same presupposition for the logic of argument, and underlying it the logic of inference too. It is that argument and inference won't yield the mysteries of their deep structures unless excavated by heavy-equipment regimes capable of mathematically precise formulation and implementation. It is here that the fissure between modern logic and Aristotle's is deepest and most intensely felt, at least by me.¹⁷

Why, then, it might well be asked, my *own* complicity in constructing ADN logics and the formal models on display in earlier work?¹⁸ I am not opposed to heavy equipment methodologies as such. I am perfectly happy to see formally contrived new ideas added to our conceptual inventories, for whatever good may be in them in due course, apart from their beauty as intellectual artifacts. I also concede the necessity of idealizations, formally wrought or not - even those that are transfinitely untrue to what happens on the ground - that are indispensable to the *descriptive* success of the empirical sciences; not excluding those of them that investigate empirically realized and normatively assessable human goings-on *in terra firma*. I also welcome the fact that thinking of things in ways they couldn't possibly be sometimes gets us to thinking of things, even perhaps of other things, in ways they *do* turn out to be.¹⁹ In this present section, I've been trying to set my course for the developments that lie ahead in section II. Part of what I want to say is how much I distrust our present compulsion to mathematicize everything in sight. Compulsions aren't good for intellectual health. They are a drag on the market and a pathological impediment to open-minded enquiry.

A further reservation concerns the groundlessness of the pretensions of the heavy-equipment methodologies to a normative authority over human cognitive performance in London, Vancouver and Guangzhou. The two most noticeable explanations of the normative authority of ideal models are the *reflective equilibrium* defence, and what I'll call the *mathematico-analytic* defence. According to the first, the correct procedures for action are those implicitly in play in the relevant community of agents. The trouble with this is the impossibility of finding credible candidates to qualify as relevant communities. If it is the human community on the ground - that is, all of us in general - then there is between how we perform and what the orthodox models require us to perform no equilibrium at all. On the other hand, if the authoritative community is the ideal-modelling research community, there will indeed be a nice concurrence between what their

models demand and what they say should be demanded. Which prompts a good, if somewhat informal, question: “Who made these guys king of the normativity castle?” Besides, why would we think that *saying* is a salient consideration? What the experts say at the office is one thing. In all other respects, they are just like the rest of us. Don’t we all put on our trousers in the same way? Don’t we all - routinely so - violate the norms of trouser-manipulation at some or other idealized juncture of perfection?

The mathematico-analytic defence is even more of a muddle. In one version of it, an idealized norm is binding on the ground when it arises in the theory as a theorem. In another, the norm’s authority arises from the fact that it is analytic in the model (i.e. made true there by stipulation). The general idea is this: The proposition “ $2 + 3 = 5$ ” is a theorem of number theory; and some people think that it is true by meanings alone. Its normative authority is straightforwardly clear. If someone in London, Vancouver or Guangzhou wants to add 2 and 3 in the way authorized by number theory, he should not identify their sum as any number that’s not 5. The same would be true for belief-closure. If someone on the ground wanted to close a belief in the way authorized by idealized closing, he would fire away transfinitely. Of course, this is absurd. No one on earth, except for the odd decision-theorist when at the office, has ever heard of the idealized closure-conditions, never mind aspiring to their fulfillment.²⁰

From the very beginnings and most of the time thereafter, the logician had to be two things at once. He would be the setter of the targets for his theory, and he would be the creator of the tools to enable him to meet them. If we were speaking of cars rather than inferences, we could see this duality nicely captured by a quite common division of labour in Detroit. Cars would be sold by the sales staff, but they would be built by the engineers. Things are different in the logic business. Not many of Ford’s sales people know anything much of how cars are built, and engineers are notorious for their poor salesmanship. But in logic, it falls to the logician to build what he sells. He must be his own engineer. It is not at all surprising that Ford’s top salesman might know nothing of engineering. But the same thing in logic would be quite astonishing.

There is a further difference between the car business and logic. Logic’s modern machinery is put together in a quite particular way. Originally designed for expressly mathematical purposes, its creators, then and now, bring a generalized mathematical sensibility to their creative work. In due course it would become apparent that the technical objects of their machinery are themselves possessed of a mathematical character and are eligible for mathematical investigation in their own right. In the car business the work of the engineering division and the work of the sales division is harmonized by the biting discipline of the bottom-line. No engineer will thrive in Dearborn if the company’s cars don’t sell, even if he’s more interested in new equipment than he is in new cars. Logic is different. By and large, the work of logicians is free of commercial expectation.²¹

When we put these two points together, we can see a quite considerable alienation of the mathematical study of logic’s machinery and the attainment of what the equipment might be good for. The factor of good-for recedes

into the background, and technological self-study becomes *sui generis*, and withal the route to the upper echelons of academic achievement and repute.²² The heavy equipment logics of the day have put themselves in an awkward position. They *say* that their technical complexifications are wanted for the inference-friendliness. But they construct their complexifications in ways that discourage if not outright preclude the accomplishment of those ends.

With complexification comes complexity, which is a well-known inhibitor of on-the-ground implementability.²³ This necessitates the reinstatement of idealizations, for two particular reasons among others. Idealizations would simplify and streamline procedures for theorem-proving; and they would explain the broadening gap between having and drawing occasioned by the idealization process itself. This would be brought about in the same old ways: by closing the gap between having and drawing in the heavy-equipment model, and by normativizing the model's drawings in relation to those that play out in the world. This is seriously problematic. Heavy equipment upgrades yield empirically false accounts of on-the-ground drawing, and do so in ways that exacerbate, not solve, the normativity problem. (As for my own involvements with the heavy-technology industry, I have never supposed that the ADN technology is normatively authoritative for anything apart from its own self-created objects. Similarly, I believe, for my ADN co-conspirators.)

One thing that could be done - and in some cases has been - to mitigate the gap-producing difficulties engendered by ideal models is to deny the *transfinitely* false ones a place at the table, and admit only those falsities for which an approximation relation is either definable or at least plausibly entertainable. With belief still our example, closure under consequence would not be permitted, but a sizeable gap could still remain between drawings in the model and drawings in human life. The difference would be that, where the original gap is transfinitely wide, the new gap is smaller - anyhow smaller enough to qualify the new closure-rule as approaching in some finite degree what actually happens here.

So adjusted, the heavy equipment approach to inference-friendliness could now be roughly summed up this way:

Complexity as gap-closing: The heavier the equipment the less empirically unfaithful the machinery's formal models, to the degree they *approximate* to what happens on the ground.

It is an interesting idea, animating another.

Approximation converges on normativity: The closer the approximation of a theoretical model of premiss-conclusion, the greater its descriptive adequacy; and the greater too, its presumptions to normative sway.

As far as I can tell, nothing in the heavy equipment literature puts things in just this way, or even close to it. And a good thing, too, readers may be thinking! Isn't everyone still cringing at *le scandale* created by poor Mill's gaffe in proposing in *Utilitarianism*, chapter 4, that "the sole evidence ... that anything is desirable, is that people do actually desire it"? A not uncommon complaint can be found in Charles Hamblin's observation that "[i]t was given to J.S. Mill to make the greatest of modern contributions to this Fallacy [= the 'naturalistic' fallacy] by perpetrating a serious example of it himself."²⁴

The mockery is misplaced. It is little more than name-calling, occasioned by the critics' misconception that Mill is saying that "The desirable (F) is what's normally desired by us all (G)" is true by meanings, supplemented by the further assumption that believing that something is F entails believing that it is G. The first of these assumptions is implausible on its face. The second owes its truth (if true it be) to the falsehood that believing that this thing *a* is F requires that there be some distinct term "G" that the believer in question believes to be semantically equivalent to "F". Notwithstanding the stout resistance it provokes, the converge of approximation on normativity is an extremely engaging idea, whatever its prior origins. It carries a suggestion of the first importance for the logic of consequence-drawing. It is that the normative authority of a logic converges on its descriptive adequacy. Should this prove to be so, it deserves acknowledgement as a foundational insight for a naturalized logic of inference.

II- The Naturalistic Turn in Logic

It is easy to see that the idea that a theory of premiss-conclusion reasoning's normative presumptions vary in some nontrivial way with its descriptive adequacy bears no *intrinsic* tie to theories contrived in the heavy-equipment way. So it would be a mistake to think that a theory's approximation to empirical fidelity is owed to its mathematical complexity. It lies instead in the fact (if fact it be) that its still-considerable falsities remain eligible for the duties of real-life approximation. It is too early to declare the heavy-equipment industries a dead loss for inference-drawing on the ground. But it surely can't hurt to start looking for alternatives.

I have an idea about where the search might pay off. It might pay off in a logic reconstructed in the way that traditional epistemology was adjusted by Quine and others in 1969 and following. Just as Quine proposed the *naturalization* of epistemology, so I propose the naturalization of *logic*.²⁵ The pivotal point of it all is this. What is the use of admitting to one's logic agents, actions and the like, if we don't admit them as they actually are on the ground - warts and all? Of course the last thing that Quine would ever agree to for logic is to do what he himself did for epistemology. For Quine, logic was first order classical quantification theory and nothing else. Quine wanted no truck with people in his logic, formally idealized or in the flesh. He clung to this conservatism until late in his career, when he grudgingly allowed that physics might require a nonclassical quantum logic, and constructivism in mathematics an intuitionist one.²⁶ The idea of naturalizing logic does not originate with me. In the modern era alone, it is actively proposed by Dewey and sympathetically entertained by Toulmin and Finocchiaro.²⁷ In this second section, I'll give a sketch of how the naturalization of logic might go, with some indication of the good of it. To this end, I'll call upon the reflections of Nat, an imaginary naturalizer.

Nat holds a naturalized logic to an adequacy condition of "empirical sensitivity". To bring it off, the logician would familiarize himself with the data that the cognitive sciences seek to account for. He would develop an informed acquaintance with the findings of the empirically best-confirmed of those theories. On pain of abandoning his logic, the naturalizer would offer a satisfactory account of its empirical disconformities with the data and findings of the partner sciences. Nat was also quick to appreciate a fact openly on view on the ground if we take the trouble to look. It is that the human animal is a knowledge-seeking organism and that premiss-conclusion reasoning is an important facilitator of its achievement. Accordingly, in addition to its empirical sensitivity, a naturalized logic of reasoning should display an *epistemic* sensitivity as well. This is already shaping up to be a coherent methodology:

Respect for data: Logic should develop a healthy respect for the data and should respond to them with empirical and epistemic sensitivity. And correlatively it should respect the *difficulty* of paying them proper respect.

Consider in this regard Gerd Gigerenzer's alarm about the "data-bending" he sees in the application to cognitive processes of methods developed by theories of statistical computation and theory testing in the

1940s and '50s. Under the influence of such assumptions, theories of cognition in experimental psychology

were cleansed of terms such as restructuring and insight, and the new mind has come to be portrayed as drawing random samples from nervous fibers, computing probabilities, calculating analysis of variance, setting decision criteria, and performing utility analyses.²⁸

Taken together, these assumptions conceptualize the human reasoner as an “intuitive statistician”. This “radically changed the kind of phenomena reported, the kind of explanation looked for, *and even the kind of data that were generated.*”²⁹ What is more, researchers who adopted the methods of inferential statistics were unaware of this change, since these methods had become canonical in psychology.³⁰

Data-bending is misconstrual of how the data actually are, owing to the wrong sort of data-loading assumptions. It is conceiving of facts on the ground in ways that facilitate pre-conceived theoretical outcomes. This is nothing to make light of. We have known at least since Bacon that data aren't self-announcing, and yet that they can't be grasped at all without some prior or concurrent conceptualization of them. This generates a nasty problem for the experimental theorist. He can't proceed without conceptualizing his data, and yet the line between conceptualization and misconceptualizations is easily transgressed. There are no algorithms for the avoidance of data-bending. But there are some useful lessons to be learned. One is to be careful. Another is to respect these difficulties.

Nat attaches a singular importance to these lessons. He also views with suspicion the most prominent justification for sticking with an empirically false theory. According to that view, it is not the aim of such theories to be descriptively adequate; the goal is to establish rules that are normatively authoritative for human practice on the ground. Fundamental to Nat's project is his rejection of this assumption, at least until such time as it might come to have a convincing independent defence. Like the rest of us, Nat knows full well that science frequently adopts false idealizations, such as the infinite cardinality of populations in population genetics. These are, he says “virtuous distortions”³¹, and are so when they are compensated for by the confirmation of the theory's observational predictions. Normatively idealized theories of empirically realized and normatively assessable human performance can't meet that test - consider here the perfect information assumption for ideally rational decision systems. So the theory's only payment alternative would appear to be its success at the *normative* checkout. But Nat's view is that there is little in the way of a settled consensus about what doing well there actually amounts to. In the absence of a sounder grasp of normative legitimacy, Nat thinks that the better option is to try to handle the problem of empirical infidelity in a different way.

Nat's own response is twofold. He stresses the priority of attending with care to the actual details of human reasoning on the ground, without preconception and well in advance of assessments of goodness or badness. His worry about preconceptions is an ecumenical one. He dislikes the preconceptions of orthodox logic. But he dislikes no less the preconceptions of orthodox psychology.³²

Nat has a working idea of how he thinks the naturalization process should go. To this end, he has availed himself of a fable, a kind of thought-experiment about the ways of naturalization. Nat imagines that for some years now a visiting team of cognitive anthropologists from a distant and unknown extraterrestrial place has been hard at work in our midst. Its earthly mission is the examination of human cognition. These cosmonauts from afar are well-equipped for their work. Themselves organic beings, they have been able to acclimate to the particularities of planet Earth. Themselves cognitive agents, they have some acquaintance with how cognition works under the ecological constraints of habitat. They are also accomplished field linguists and intelligent problem-solvers. Nat has adapted this fiction of the visiting scientist from van Fraassen's notion of an epistemic marriage, which envisages an epistemic partnership between dolphins, extraterrestrials and us.³³ Just as Wittgenstein would wonder whether if a lion could talk, we would be able to understand him, the visitors were interested in whether the information-processing activities of us rose to the bar of cognition in a manner that would be discernible to their own methods of enquiry.³⁴

The visiting anthropologists are here to do some descriptive epistemology, to take the earthly cognitive pulse as best they can with the exploratory resources available to them. They will run their investigations in compliance with their own understanding of best scientific practice. They will leave themselves free to consult the published record of *our own* conception of best scientific practice, but without prior commitment to defer to ours when it conflicts with theirs. Neither will they consult our philosophers, not even the ones who do logic; at least not until their own work is finished. Why would the visitors bother with philosophers? They are not themselves philosophers; they are scientists conducting a naturalistic examination of the naturally occurring phenomena of cognitive behaviour in their subject population. This exclusion is neither ideological nor hostile. It simply reflects the plain fact that scientists do their thing and philosophers do their largely different thing. Philosophers and scientists can share a common subject-matter, but they each tend to cut their respective cakes largely in independence of the other. A sensible course for each would be to arrive at a decision about what happens now. There might be some advantage for each in exposing one another's views of the matters they have in common to a reciprocal scrutiny. But the visitors' point, and Nat's too, is that this scrutiny is better reserved until the scientists have done their business.

In this same spirit, the visitors decided that it was not part of their mission to discern how on-the-ground human inference-makers *defined* inference - or knowledge or belief or whatever else. Nor would they circulate questionnaires asking their subjects to tell them whether they thought that this, that or the other thing is *bona fide* knowledge or good reasoning. The visitors have long been aware that, among their own kind, people do quite well at knowing and reasoning without being very good at saying what these things are. Until they learned otherwise, they would assume the same for us. From this came another foundational insight:

Being good at and knowing what: Being good at knowing things and reasoning well does not require knowers and reasoners to know how to define what knowledge and good reasoning are, or to specify the conditions that bring these things about.

Substitution here for “reasoning” of “remembering”, “imagining”, “seeing”, “high-fly ball catching” and a plethora of others generates a host of sentences most of us would consider too obviously true for words. Consider, for example, our knowledge by looking that it has snowed overnight. Everyone in town with eyes to see knows this to be so. But hardly any of them has much of an idea about how the mechanics of visual cognition actually play out. Why wouldn’t this also be true for “knowing” and “reasoning”?

Perhaps it will be protested that if we are so good at knowing things why are we so bad at knowing what knowing is? It is an excellent question but a feeble protest. Does the questioner really think that his question carries the force of rebuttal? If so, he would do well to explain why. When, in the late stages of their enquiry, the visitors relaxed their exclusion of earth-bound philosophy, they were astonished to learn - disapprovingly so - that some of the first of our great philosophers were in thrall to the idea that a person’s knowledge requires a prior or concurrent grasp of a real definition of it, that a concept can’t be instantiated in human behaviour in the absence of the behavior’s “analytic grasp” of it. This, the visitors thought, amounted to a scepticism so corrosive as to make the attainment of knowledge a generally impossible target. When some of our own earthbound professors of epistemology protested in turn the excessiveness of the visitors’ alarm, they answered in unison: Socrates and *les autres* could not have thought this way had they paid scrupulous attention to the ups and downs of what really goes on in human life; especially in the *agora*.

Perhaps the visitors were a trifle hasty. They didn’t have time to immerse themselves in the history of earthly epistemology. But they’d had enough exposure to it to have tasked themselves with the four questions. One is whether having a real definition - or conceptual analysis - of it is a condition on knowing what knowledge is. A second is whether knowing what knowledge is is a condition of there being any. The third is whether the ancients were inclined to favour affirmative answers. The fourth is whether this favoritism has a discernible presence in modern-day analytic approaches to epistemology. Not having the time or inclination for extended consideration of these matters, they came together on three summary positions; *first*, that any notion that a scientific knowledge of knowledge is to be got by a conceptual analysis of “knows” is a misbegotten idea; *second* that the only room for big-box scepticism in the science of human knowledge is by way of the default rule that big-box scepticism in the science of human knowledge should not be so much as entertained, never mind rebutted, except for weighty cause; and *third* that the earthly epistemological tradition betrays too little heed of these constraints; not without exception, but dominantly so. They thought that a thoughtful examination of the cognitive routines of human life make it clear that just about the last thing that could be true of it is radical scepticism of any broad kind. On the contrary, they thought, it was empirically evident that human beings are good at knowing things - not perfect but good; that they have lots

and lots of it about lots and lots of different things; that the human cognitive harvest is both abundant and diverse.

A further shock was administered when, shortly before heading for home, the group began to look into what earthbound *logicians* had been saying about these things. They were taken aback to discover the confidence and wide-spreadness of the dogma that premiss-conclusion reasoning is no good when it fails to be deductively valid or at least inductively strong in the technical sense familiar to the statistico-experimental sciences. What surprised them most was the utter lack of behavioural recognition of this would-be fact in the subject population. Most of the reasoning that passes muster there - and is evidently accepted as good - fails both these standards. This left the visitors with two choices. They could condemn their human subjects as across-the-board losers in the reasoning game. Or they could reject the validity-or-inductive strength condition as a general requirement for good premiss-conclusion drawing. Of course, they chose the latter. How could they have not? They were natural scientists who faithfully respected the necessity of respecting the data on the ground. I hardly need say that Nat was quick to sign on, and in short order would see that one of the essential tasks for a naturalized logic of inference is the discovery of the conditions under which this “third way reasoning” would be properly achieved; and in the process to see that the resulting third way logic would be the natural home for virtually all the nonmonotonic logics currently on offer, once such compensating adjustments as might be required were worked out.³⁵

It was equally clear to the visiting team that there is another abundance for the naturalized logician to pay attention to in the cognitive ecologies of the subject population. Knowledge exists there in abundance. Error is another of its abundances. The human animal makes lots and lots of errors about lots and lots of different things. A logic of reasoning must bring these abundances into a benign harmony, in response to yet another empirically discernible feature of human cognition.

Enough-enough: Human beings know enough about enough of the right things enough of the time for survival and prosperity and, from time to time, for the erection of civilizations of dignity and lingering worth.

There is no question here of our prosperity *entailing* an abundance of enabling knowledge. Cognitive abundance is our visitors’ working hypothesis. It is not ruled out that there is none better. It is not ruled out that they might come to think that a more plausible working hypothesis would be that the beliefs that serve us well are mainly false, in which case, *right* belief would separate away from *true* belief. But, if anything was clear to the anthropologists, it was that beings like us are awash in what we might call alethic beliefs, which are believings-to-be-true. This gives the new working hypothesis a bit of a twist, providing that the best way of achieving prosperity is believing-to-be true propositions that are false. This handed the visitors a chuckle and the idea was dropped like a hot potato. Where, they asked, is it empirically discernible in the belief-prosperity data that this story would hold in the general case?

Nat followed the visitors in thinking that an attentive naturalizer would lodge his curiosity about good and bad reasoning in a default principle for the logical theorist:

NN-convergence: Take it as given in the absence of particular reasons to the contrary that humans reason well when they reason in the ways that humans normally reason in the conditions of real life.

Nat was careful (and well-advised) to add an important caveat. The NN-convergence principle is *not* a safe default for all aspects of human cognition; its application here is reserved for premiss-conclusion reasoning. It tells us to judge premiss-conclusion reasoning in roughly the same sort of way that we'd check the subject's pulmonary behaviour. What this amounts to is that, for the most part, a human agent's premiss-conclusion reasoning is the right way to reason when his conclusion-drawing mechanisms are in good working order, and at present working in the right way, engaging good information in the absence of hostile externalities.

Breathing is like that too. By and large, the goodness of good breathing depends on the ship-shapeness of the pulmonary equipment. At bottom, as we might say, good breathing is not down to us; it is down to our equipment. Sometimes, however, it *is* down to us. We can't sing grand opera, or achieve a place on the Vancouver Canucks, if we don't breathe in the required ways. We can't bring this off until we learn to do this. Almost always we'll have to be taught by an expert, and oftener than not it won't work. Hardly anyone has the breath for *Tosca* or for the seventh and deciding game of the Stanley Cup playoffs.³⁶

These same passivities and activities are discernible in reasoning. Most of it happens passively and largely out of sight of the mind's eye. Most of it is down to good machinery. Indeed, it is *always* down to good machinery, but sometimes it is also down to us - to our disciplined, patient and skilled application of the expert routines that knowledge sometimes requires. Think here of the incompleteness of formal arithmetic.

We've already said that our visitors aren't philosophers or logicians of the orthodox sort. They are cognitive anthropologists. This is not Nat's situation. Nat is a philosopher whose dissent is launched from within the very orthodoxies he finds fault with. It might be appropriate for the visitors not to take a philosophical position on knowledge. But Nat can hardly proceed to completion without allowing his epistemic reach to carry *epistemological* implications as well. Nat knows that philosophical theories of knowledge broadly partition into two paradigms, one of them more historically dominant than the other. This first, he calls the Command and Control Model, and the other the Causal Response Model. We see in this division a nice correspondence with the already noted distinction between active-case knowledge, whose attainment is significantly down to us, and passive-case knowledge, whose attainment is largely down to our equipment. Having noticed the statistical dominance of the passive over the active, together with the passive underlay of even the active, Nat concluded that the right base epistemology for a naturalized logic is the CR model, provided that, where indicated, it accommodates the CC model as a proper subtheory for various kinds of knowledge acquisition "up above", in which the knowing agent has an active and self-aware role to play. Think again of the second incompleteness theorem.

The idea that a naturalized *logic* needs a base *epistemology* might strike some readers as implausible - anyhow puzzling. Why would I assume it? The answer is that in the human world reasoning is transacted by cognitive

beings, and one of its principal functions is the facilitation of knowledge-seeking and the attainment of epistemic goals. Nat is a philosopher who wants a philosophically tenable account of what makes such reasoning in pursuits of such ends the right way to reason or, as the case may be, the wrong. A logic so designed cannot be judged in the absence of a philosophically convincing understanding of what it is that is facilitated or attained when these goals are in play and properly handled.³⁷

Given Nat's naturalistic leanings, this seems much the right choice. Right or not, it is a fateful one. It effects a considerable scrambling of the once-pacific distinction between reasons and causes, and it makes of knowledge a much more causal phenomenon than an intellectually wrought achievement. The same holds for reasoning. By a statistically large measure, our premiss-conclusion reasoning is right when the conclusions we draw are causally induced by belief-producing devices when working as they should. In shorter words still, in the general case you don't have to be smart to reason well; you have to be healthy.

An interesting case in point, and one to which the visiting team paid a lot of attention, is the utter widespreadness of the phenomenon of being got to know things by being told them by others. Both models recognize the phenomenon, but they give it quite different theoretical treatments. The central point of contention is whether *justification* is a general control point for knowledge-transmission, with the CC-theorist voting aye and the CR-theorist nay; that is, nay as a general condition on transmission. The CR-reservation came down to this: If we pay close attention to what happens on the ground, the presence of justificatory involvement is *markedly* less discernible than the levels of recognizable cognitive satisfaction among transmitters and recipients alike. As the visitors came to appreciate, the CC-crowd has spared no effort to attribute the workings of justification even in the absence of its recognizable behavioural presence in the general case. Nat has noticed this too. The visitors didn't quite know what to make of it. But Nat knew. What he made of it was that the CC-crowd weren't paying sufficient heed to the respect-for-data rule.³⁸

The on-ground data also disclose how much of this causal belief-inducing activity proceeds out of reach of the mind's eye, without notice or attentiveness, without deliberation or overt case-making. As Nat puts it, most of human inference is inference "down below". If this is so, it plays straight into the normativity question. If we wish to remain faithful to the NN-convergence thesis, we will need to derive a naturalistic account of errors of reasoning; for it is to precisely these that the badness of bad reasoning is owed, is it not? Such accounts don't lie idly about. They are not numerously available or free on board. They have to be laboured after. First and foremost perhaps is the reconciliation of our two abundance theses - the abundance of knowledge and the abundance of error. Trailing along is a perfectly natural puzzlement about how, if we're so good at reasoning, why are we so bad at avoiding errors.

Nat has an answer to this which (very sketchily) comes to this: A standing liability for the human knower, no less than the human high-jumper, is that there is only so much he can do. There is only so much that it

makes sense for him to want to do - or have the slightest interest in doing. He must learn to live and to set his sights within his cognitive means. As with any design-constrained and resource-limited activity, cognitive success calls for economic alignment of capacity with resource-availability. Nat noted that human individuals are graced with very efficient feedback mechanisms. It is a welcome advantage, enabling a better record at error detection after the fact than before. He concluded from this that it is more economical for a resource-bound individual to correct mistakes after the fact than to avoid them before commission. He also observed that, in case upon case, the human animal is more adept at conclusion-drawing than he is at premiss-selection. Think of the former as errors *of* reasoning, and of the latter as errors *in* reasoning. The dominant fault of premiss-selection is *misinformation*; and it is markedly easier to be misinformed about something than to draw from it the wrong conclusion. On thinking it over, Nat came to the view that

Misinformation and misinference: There are significant global variations in peoples' well-informedness - matching kindred variations in region-to-region levels of ignorance - and yet comparatively uniform performance-levels in human conclusion-drawing.

Nat's naturalizing focus is on premiss-conclusion reasoning, chiefly on the drawing side. If, as it appears, this is something we're uniformly good at, then errors *of* reasoning, when committed at all, must arise from external hostilities or equipment failure. Generally speaking, when we make mistakes of conclusion-drawing, we are dog-tired or strung out, or leveled by a stroke; or the conclusion-drawing equipment isn't - like the nineteen year old Ford - in quite good enough working order; or we are awash in information-overload. If this is right, something else is also bound to be right. It is that

The thickness of error: Error is a natural phenomenon, with a dominantly descriptive character, but also imbued with what Bernard Williams calls *thickness*. Roughly speaking, this means that to find error in a person's performance is to see it as an utterly natural phenomenon but one that is out of joint with the how things are supposed to go.³⁹

Since his first course in logic, Nat has known that logic from its very beginnings sought for the attainment of a decent theoretical command of what Aristotle calls *fallacious* reasoning. Since his second course in logic, Nat has also known that the fallacies programme is nowhere in sight in any of the going mainstream logics.⁴⁰ It is not hard to see why. We've already said that the modern orthodoxies were built for the relief they promised for metaphysical and epistemological anxieties in the foundations of mathematics. They weren't built for human reasoning, even for when it is transacted fallaciously. This, of course, can't be Nat's own position. Nat wants a logic for real-life premiss-conclusion inference. He wants his logic to solve the normativity problem. He thinks that the correct account of errors of reasoning will be the key to its solution. Surely, one would think, no account of reasoning errors could be complete if it didn't revive and make some headway with the fallacies project.

Nat knows that on the traditional approach fallacies are errors of reasoning having certain distinguishing features. One is that they are attractive, hence inapparent. Another is that they are universal, in the sense that virtually all of us are disposed to commit them with a frequency higher than our error-makings in general. Yet another is their incorrigibility; that is, even after detection and correction, rates of post-diagnostic recidivism are extremely high. Like the rest of us, Nat is also familiar with the traditional list of the fallacies - hasty generalization, *argumentum ad verecundiam* (argument from authority), *argumentum ad ignorantiam* (arguments from ignorance), and so on. It wasn't long before Nat made a discovery that genuinely surprised him. He saw that a proper regard for the respect for data principle discloses that virtually all the items on the traditional *list* of fallacies have no discernible presence in the instantiation-class of the traditional *concept* of fallacy. Accordingly,

Concept-list misalignment: Virtually none of the fallacies in the traditional list lies in the extension of the predicate "is a fallacy" as traditionally interpreted. Either they aren't errors, or they are not attractively inapparent, or not universal or incorrigible.

Here is an example that especially impressed Nat. Nat is on his first trip to Brazil, a country he knows little of, but enough to know that animals called ocelots are resident there. One day, Nat and his Brazilian host are tramping the countryside. "Look", exclaims Luis, "an ocelot!" Nat is surprised. "Good heavens, Luis, I had always imagined ocelots as two-legged, not four." Nat has come to see that, on the basis of a single encounter, ocelots are four-legged. On thinking it over a bit, he also realized that the true generalization "Ocelots are four-legged" is not falsified by the plain and subsequently discovered fact that this other ocelot, Ozzie, a beloved resident of the zoo of his friend Luis' hometown, is three-legged. And in no time at all it became apparent to him that this kind of hasty generalization is, in the human species, as common as dirt (as the saying goes), and that to a quite marked degree the generalizations hastily drawn are actually right, not wrong. Whereupon we have it that

Hasty generalization: Hasty generalization is not a fallacy in the traditional sense. Indeed, comparatively speaking, it is hardly ever wrong when actually performed.

The obvious question now is whether *anything* instantiates the traditional concept. It is, as I write, an open question in Nat's logic.⁴¹

The naturalistic turn pulls logic and cognitive science in the opposite direction from the mathematical turn. For three decades and more, the mathematically shaped enquiry has searched out an affiliation with a closer attachment to agent-centred, goal-directed, resource-based, time and action systems. These enrichments are a considerable complication for, whose management more basic formal equipment must be upgraded with new machinery of correspondingly greater complexity, sometimes problematically so. Nat's worry, like my own, about these heavy-equipment upgrades is that the more complex they are to handle, the likelier they are to invite the solace of a new batch of simplifying purpose-built performance norms. Nat isn't opposed to the enlargement of capital assets as such. His

reservation about heavy-equipment upgrades is that they leave the normativity problem undealt with. His present inclination is to enrich the logic of premiss-conclusion reasoning with naturalistic assets, especially those of them that improve our grasp of the on-the-ground management of error - its avoidance, its commission, its detection and repair. For it is here that he sees promise of a principled solution to the normativity problem. The empirical turn is an attempt to reshape this enlargement of capital assets, by lightening up on the notion that theory-building is intrinsically theorem-proving. Theorem-proving is not demanded for population biology. Why, asks Nat, should it be demanded here?

What Nat proposes for logic is what Quine proposed for epistemology. Not everyone thinks that Quine's is a tenable project or, so far, a well executed one. But no one should think that these days naturalized epistemology is anything but a well-accepted part of mainstream epistemology. It is vanishingly unlikely that a friend of naturalized epistemology would dislike Nat's proposal for logic because he distrusts the role of naturalism in philosophy. But he might well dislike it because it lies in the nature of logic not to take well to naturalistic intrusion. This, after all, was Quine's own position. At the core of it all is modern logic's deeply dug-in loathing of psychologism. Epistemology leaves lots of room for psychology, and logic leaves none. Naturalized epistemology may now have found a place in the big leagues, but this is the last thing that one could say of logic. This makes Nat's proposal a radical one for logic if not any longer for epistemology. It also makes Nat's proposal a contentious departure from the still well-favoured normative idealization approach to the social sciences. In a nutshell, what Nat wants is logic's reinstatement of psychologism.

Nat fully acknowledges the efforts of the newer developments in heavy-equipment logics to do better on the score of on-the-ground inference-friendliness. His chief reservations are two. The heavy-equipment technologies don't solve the normativity problem; and bulking up the formal machinery hasn't closed the gap between the logic's theorems and settled practice on the ground. As far as psychology goes, there are weighty autoepistemic considerations to take respectful notice of. If the heavy-equipment crowd thought that there was room for psychology in their projects, they'd have put some in. But they haven't. so they don't.

If the case against the normative presumptions of heavy equipment logics could be made to stand, it carries like consequences for the normative presumptions of the ideal models approach to the social sciences generally. This is getting to be quite a bit of nay-saying. If acquiesced to, all of normatively presumptive science would be put on hold until the normative authority problem is properly sorted out. This separates Nat from virtually all the going traffic in agent-based logics, including logics of probabilistic reasoning, belief-change and decision, epistemic and justification logics, fallacy theory, discourse analysis and normative psychology. Of course, enquiries of every kind are needful of starting points and of assumptions that frame their conceptual spaces. These are assumptions that lend enquiry

its procedural and organizational shape. This produces a pair of important consequences for the would-be dissident.

One is that by the very nature of received opinion, there is not much pent up enthusiasm for paradigm-overthrow. The other is that, for want of practice, the orthodoxies' disciples aren't much good at defending them. There is a story making the rounds in which the dialethic logician Richard Routley once challenged the logically more strait-laced David Lewis to prove the classically interpreted law of noncontradiction.⁴² There was point to his challenge. Routley thought that, when interpreted the right way, the law of noncontradiction did *not* preclude the truth of some select contradictions. So the challenge to Lewis was to show that this couldn't be so. When it came down to it, Lewis didn't bite; he refused to be drawn. He told Routley, in effect, to grow up and stop horsing around. Lewis' was a telling response. It was an outright and unconsidered dismissal. The point of this little *tableau* is dialectical. Orthodox assumptions carry and are protected by high levels of dialectical *inertia*. So Nat would have been foolish not to have anticipated that his own dissensions might receive scant attention in the high courts of received opinion.⁴³ On the other hand, Quine and others prevailed to good effect in the aftermath of 1969, against the grain of stiff resistance. So who knows? Perhaps Nat's naturalistic prospects will have brightened forty or so years hence⁴⁴

Acknowledgements: For comments on earlier drafts or discussion of closely related matters, I thank most warmly Johan van Benthem, Franz Berto, Dov Gabbay, John Greco, Maurice Finocchiaro, David Hitchcock, Jaakko Hintikka, Frank Hong, Ralph Johnson, Lorenzo Magnani, Christopher Mole, Adam Morton, Ahti-Veikko Pietarinen, Shahid Rahman, Harvey Siegel, Robert Thomas and Yi Zhao.

Notes

1 The Abductive Systems Group, Department of Philosophy, University of British Columbia, john.woods@ubc.ca, www.johnwoods.ca.

2 Readers might wonder about the concentration on premiss-conclusion reasoning. Doesn't logic also investigate higher order cognitive practices, such as decision-making, belief-change and game theoretic enterprises of all kinds? There are two reasons for the concentration. One is that the investigation of the consequences relations that underly premiss-conclusion reasoning is what logicians are best at. The other is that, in any event, most of higher level inference involves in some way or other premiss-conclusion reasoning.

3 Although Aristotle is a deductivist about *sylogistic* reasoning, it is manifestly not his view – or a Greek one either – that all good inference is deductive. See here M.F. Burnyeat, “The origins of non-deductive inference”, in Jonathan Barnes, J. Brunschewig and M.F. Burnyeat, editor, *Science and Speculation: Studies in Hellenistic Theory and Practice*, pages 193-238, Cambridge: Cambridge University Press, 1982. Reprinted in M.F. Burnyeat, *Explorations in Ancient and Modern Philosophy*, volume 1, pages 112-151, New York: Cambridge University Press, 2012.

4 These variations on variations are more robust and complex than straightforward set theoretic restrictions. See here David Makinson's *Bridging From Classical to Nonmonotonic Logic*, volume 5 of Topics in Computing, London: College Publications, 2005. See also John Woods, *Errors of Reasoning: Naturalizing the Logic of Inference*, volume 45 of Studies in Logic, London: College Publications, 2013; chapters 7 and 8. An exception is autoepistemic consequence. Autoepistemic reasoning is discussed in chapter 10 of *Errors*.

5 See Peirce; “My proposition is that logic, in the strict sense of the term, has nothing to do with how you think”, p. 143 of Charles S. Peirce, *Reasoning and the Logic of Things: The Cambridge Conference Lectures of 1898*, Kenneth Laine Keiner, editor, Cambridge, MA: Harvard University Press, 1992, and Gilbert Harman, “Induction”, in Marshall Swain, editor, *Induction, Acceptance and Rational Belief*, Dordrecht: Reidel, 1970, *Change in View*, Cambridge, MA: MIT Press, 1986; chapter 1. See also in this same vein Jaakko Hintikka's distinction between “definatory” (roughly, consequence-having) rules and “strategic” (roughly, consequence-drawing) rules in his “The role of logic in argumentation”, *The Monist*, 72 (1989), 3-24, and Hintikka, Ilpo Halonen and Arto Mutanen, “Interrogative logic as a general theory of reasoning”, in Dov Gabbay, Ralph Johnson, Hans Jürgen Ohlbach and John Woods, editors, *Handbook of the Logic of Argument and Inference: The Turn Towards the Practical*, volume one of Studies in Logic and Practical Reasoning, pages 295-337, Amsterdam: North-Holland, 2002.

6 In a variation, normatively binding idealizations are achieved not by the logician's mere sayso, but in some putatively *à priori* sort of way, such as a conceptual analysis of “rational agent”.

7 John Woods *Aristotle's Earlier Logic*, 2nd edition, London: College Publications, 2014. My thanks to David Hitchcock for helpful instruction on this last clause.

8 *Aristotle's Earlier Logic*, chapters 3 and 5.

9 As opposed to “motivational”, which is another story entirely.

10 Jaakko Hintikka, *Knowledge and Belief: An Introduction to the Logic of Two Notions*, Ithaca, NY: Cornell University Press, 1962.

11 A not untypical example is any system that requires the reasoning agent to impose – if only provisionally – the closed world assumption.

12 Hintikka is not indifferent to this difficulty, and seeks out some relief from it by modifying the interpretation of the K-operator. He allows it as a representation of knowledge both express and *tacit*. There is something good about this and also something less so. Cognition on the ground routinely operates tacitly and inarticulately. So it is good that an agent-based logic would take some notice of this. Not so good is the theoretical cost of the measure. Hintikka now owes us what the 1962 monograph doesn't begin to deliver; and that's a cost. Also questionable is whether Hintikka's tacit knowledge actually does effect gap-closure. Take the second incompleteness theorem as an example. We all have a

Hintikkian tacit knowledge of it just because, in principle, we can all be got to see the self-defeating character of its denial without having to draw on any new empirical information. The trouble lies with the “in principle”-clause. “In principle” is justly infamous as a gap-widener, not closer. My thanks to Frank Hong for discussion of this point.

13 See Aristotle’s *Earlier Logic*, Appendix on Categorical Syllogisms.

14 Johan van Benthem, *Logic and Dynamics of Information and Interaction*, New York: Cambridge University Press, 2011; p. ix. I have drawn this paragraph and the one that follows from my “Advice on the logic of argument”, *Revista de Humanidades de Valparaíso*, 1 (2013), online version at <http://www.revistafilosofiauv.cl/>

15 In Gabbay *et al*, editors, *Handbook of the Logic of Argument and Inference*, 449-481.

16 Dov M. Gabbay, “Equational approach to argument networks”, *Argument and Computation*, 3 (2012), 87-142; Howard Barringer, Dov M. Gabbay and John Woods, “Temporal argumentation networks”, *Argument and Computation*, 2-3 (2012), 143-202; and Barringer, Gabbay and Woods, “Modal argumentation networks”, *Argument and Computation*, 2-3 (2012), 203-227. See also Artur d’Avila Garcez, Dov M. Gabbay, Olivier Ray and John Woods, “Abductive reasoning in neural-symbolic systems”, *Topoi*, 26 (2007), 37-49, and Artur d’Avila Garcez, Howard Barringer, Dov M. Gabbay and John Woods, *Neuro-fuzzy Argumentation Networks*, to appear.

17 John Burgess reports the standard view that “formal logic” is a pleonasm and “informal logic” an oxymoron. See his *Philosophical Logic*, Princeton: Princeton University Press, 2009, p. 2. In the late 1970s, informal logicians used to fret about the suitability of the adjective “informal” in apposition to so noble a noun as “logic”. “Informal”, they feared, bespoke a kind of casualness or, heaven forbid, sloppiness. One day Michael Scriven stated his own view of the matter. “Informal logic” was indeed the wrong name for their enterprise. “But what should we call it?”, he was asked. “Call it what it is”, said Scriven. “Call it ‘logic’”.

18 See, for example, Dov M. Gabbay and John Woods, *Agenda Relevance: A Study in Formal Pragmatics* and *The Reach of Abduction: Insight and Trial*, volumes 1 and 2 of *A Practical Logic of Cognitive Systems*, Amsterdam: North-Holland, 2003 and 2005.

19 Consider, for example, population genetics, in which the idealization that populations are infinitely large plays an indispensable role. No one thinks that population genetics tells us anything at all about the cardinality of empirically realizable populations. But everyone knows that it tells us a great deal about natural selection on the ground.

20 For a more detailed consideration of these points readers could consult my “Epistemology mathematicized”, *Informal Logic*, 33 (2013), 292-331, and *Errors of Reasoning*, chapter 2, section 3.

21 This is less so in departments of computer science. It is an interesting story to tell, but longer than there is time for here.

22 It would be rude to speak here of heavy-equipment autoeroticism, but we would know what was meant.

23 There is a significant ambiguity between something that can be implemented *in* a human agent, for example, a model that is realized in a neural net, and something that can be implemented *by* a human agent, as for example, a rule of decision manoeuvre that is simple enough for a person to deliberate upon and follow. Complexity is much less a problem for the first than the second. Nevertheless, the purport of heavy equipment gap-closers is to facilitate implementability-*by*. For implementability-*in*, see again Gabbay *et al.*, *Neuro-fuzzy Argumentation Networks*. On the other hand, *Errors of Reasoning* offers an alternative: For implementability-*by*, it is advisable to drop the idea that real-life reasoning competence is intrinsically, or even generally, a matter of following *rules*.

24 C.L. Hamblin, *Fallacies*, London: Methuen, 1970.

25 “Epistemology naturalized”, in *Ontological Relativity and Other Essays*, pages 114-138, New York: Columbia University Press, 1969.

26 I am just one of the many who think Quine’s contributions have done epistemology nothing but good. This makes me hopeful that a like success might be had with logic.

27 John Dewey, *The Later Works*, 17 volumes, Carbondale: Southern Illinois Press, 1981-1991; volume 12, 27; Stephen Toulmin, *The Uses of Argument*, Cambridge: Cambridge University Press, 1958, 257; and Maurice Finocchiaro, *Arguments About Arguments*, New York: Cambridge University Press, 2005, 6-7. I first became aware of Finocchiaro's naturalistic leanings when viewing a poster for his 1987 lecture at the University of Groningen, entitled "Empirical logic".

28 Gerd Gigerenzer, "From tools to theories", in Carl Graumann and Kenneth J. Gergen, editors, *Historical Dimensions of Psychology Discourse*, pages 336-359, Cambridge: Cambridge University Press, 1996; 339.

29 Idem; emphasis added.

30 Further details may be found in Dov Gabbay's and my "Filtration structures and the cut down problem for abduction", in Kent A. Peacock and Andrew D. Irvine, editors, *Mistakes of Reason: Essays in Honour of John Woods*, pages 398-417, Toronto: University of Toronto Press, 2005; 411-414.

31 John Woods and Alirio Rosales, "Virtuous distortion in model-based science", in Lorenzo Magnani and Walter Carnelli, editors, *Model-Based Reasoning in Science and Technology: Abduction, Logic and Computational Discovery*, pages 3-30, Berlin: Springer, 2010.

32 Nat has heard and liked an amusing bit of mischief about physicists. It is said that physicists find fault on two counts with biologists. One is that biologists aren't very good at data-analysis. Another is that they aren't particularly adept at model-building. Of course, it's only a story, a kind of jape really. But it strikes Nat that, given the marked difference in the respective complexities of the motivating and confirmatory data of these two disciplines, perhaps the story has something of the ring of truth. It also strikes him that if this were really so, how less it could be so for the *social sciences*?

33 Bas van Fraassen, "The day of the dolphins: Puzzling over epistemic partnership", in Peacock and Irvine, *Mistakes of Reason*, 111-133.

34 The literature of lionspeak and related matters is large, and amply cited in Dorit Bar-On's "Expressive communication and continuity scepticism", *Journal of Philosophy*, 60 (2013), 293-330.

35 Still, in so saying, Nat was not giving up his renunciation of the orthodox approaches to reasoning even where deductive validity or inductive strength is the appropriate assessment-standard. This calls to mind an amusing story recounted by Johan van Benthem in a talk at UBC in 2012. When a proof theorist says his piece to a roomful of working mathematicians they simply can't believe how little he knows about the real business of proving mathematical theorems.

36 A similar approach is taken, with a good deal more detail and sophistication, in Ernest Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, volume 1, Oxford: Oxford University Press, 2007, and volume 2, 2009. See also John Greco, *Achieving Knowledge*, New York: Cambridge University Press, 2010, and *Errors of Reasoning*, chapter 3, notwithstanding modest demurrals in chapter 2, p. 53 n. 13.

37 I am grateful to John Greco for having pressed me on this.

38 Told-knowledge is discussed at length in chapter 10 of *Errors of Reasoning*. Knowledge occasioned by not being told its opposite is dealt with in chapter 11.

39 *Ethics and the Limits of Philosophy*, Cambridge, MA: Harvard University Press, 1985.

40 Informal logic is a different story, but it is a story without readers in the mainstream. Recall Burgess' quip that according to the elites "informal logic" is a contradiction in terms. When have we seen a fallacies paper in the *Journal of Symbolic Logic*? The answer is: *Never*.

41 Notwithstanding some answers ventured in *Errors of Reasoning*, chapter 15.

42 See my "Wrestling with (and without) dialetheism", *The Australasian Journal of Philosophy*, 83 (2005), 87-102.

43 In fact, Lewis was not quite true to his word. Instead of just walking away, he asserted that LNC is a truth that neither requires nor admits of proof. This was a mistake, apparently not noticed by Routley. Consider the following:

1. LNC is a true proposition that neither requires nor admits of proof. (Lewis)
2. LNC is true. (from 1, De Morgan and \wedge -elimination)
3. LNC can't be proved. (ditto)
4. But this proves that LNC is true and that it can't be proved. (from 2 and 3, \wedge -introduction)

44 Notwithstanding our differences of methodological perspective, I say this in the spirit of van Benthem and his colleagues in the closing lines of their Introduction to *Logic in Action*: “And with this much of the five logician actors out of the way, we draw the curtain for this little book – and invite you to enter our world.” The actors are Johan van Benthem, Paul Dekker, Jan van Eijck, Maarten de Rijke and Yde Venema, *Logic in Action*, Amsterdam: Institute for Logic, Language and Computation, 2001; p. 5, emphasis added.