Australasian Journal of Philosophy
Vol. 86, No. 3, pp. 481–497; September 2008

KITCHER, MATHEMATICS, AND NATURALISM

Jeffrey W. Roland

This paper argues that Philip Kitcher's epistemology of mathematics, codified in his Naturalistic Constructivism, is not naturalistic on Kitcher's own conception of naturalism. Kitcher's conception of naturalism is committed to (i) explaining the correctness of belief-regulating norms and (ii) a realist notion of truth. Naturalistic Constructivism is unable to simultaneously meet both of these commitments.

Introduction

Prima facie, an epistemology of mathematics should at least answer the question.

(ME) What justifies us in believing propositions of mathematics?[1]

where justification is taken to be whatever makes the difference between believing truly and knowing.[2] Philip Kitcher has proposed a naturalistic answer to (ME) in the form of his Naturalistic Constructivism [1984; 1988]. This invites the question of how we should understand naturalism. Though naturalistic tendencies are widespread in contemporary philosophical discussions, there is (notoriously) no consensus on just what naturalism is or consists in. I will not attempt to rectify this situation here.[3] Rather, I will assess Kitcher's view in the light of his own conception of naturalism, which he has discussed in (among other places) Kitcher [1984; 1988; 1993] and detailed most thoroughly in Kitcher [1992].

For present purposes, the most important features of Kitcher's naturalism are its commitments to

(KC_1) explaining the correctness of belief-regulating norms (i.e., norms governing belief formation and retention)

[1]Those who have qualms about the metaphysical status of propositions can take 'proposition' and its cognates in this paper as shorthand for contents of beliefs, whatever one thinks such contents are.
[2]This conception of justification sometimes goes by 'warrant' [Plantinga 1993b; Plantinga 1993a; Kitcher 1984].
[3]I have given considerable attention to this issue elsewhere [manuscript a].

Australasian Journal of Philosophy
ISSN 0004-8402 print/ISSN 1471-6828 online © 2008 Australasian Association of Philosophy
http://www.tandf.co.uk/journals DOI: 10.1080/00048400802001947

(KC_2) a realist notion of truth, in the sense that what makes a statement (belief, etc.) true or false is independent of us and our cognitive activities (modulo statements about us and our activities, of course).

Call these *Kitcher's commitments*. I will argue that Kitcher's Naturalistic Constructivism has trouble simultaneously honouring both of these commitments. If my arguments are correct, then, by his own lights, Kitcher has not provided a naturalistic answer to (ME). Before turning to the question of whether Naturalistic Constructivism honours (KC_1) and (KC_2), I briefly motivate including these commitments among those of Kitcher's naturalism.

I. On Kitcher's Commitments

Consider (KC_1). According to Kitcher, '[t]he central problem of epistemology' by naturalism's lights 'is to understand the epistemic quality of human cognitive performance, and to specify strategies through whose use human beings can improve their cognitive states' [Kitcher 1992: 74–5]. Cognitive improvement is cognitive progress—i.e., progress toward 'impersonal epistemic goals' [Kitcher 1993: 93]. Kitcher writes that '[s]ome types of [cognitive] processes are conducive to cognitive progress; others are not' [Kitcher 1993: 186]. What separates the former from the latter? In a word: reliability.

> Focus for the moment on the simplest kind of cognitive progress, that of finding a true answer to a significant question. Imagine that various subjects have all the information needed to generate belief in the correct answer: there is an inferential process that could lead any of them from items in declarative memory to a state of belief in the correct answer and all of them have the propensities required to undergo this process. Some of them activate the right propensities and achieve the true answer. Others activate propensities that are very unlikely to generate true answers (for example, suppose that they lexicographically order the alternatives and choose the eleventh), and they come to believe incorrect answers. There is a distinction to be drawn here. Some undergo *processes that reliably generate true beliefs*, while others undergo *processes that have a very small chance of yielding true beliefs*.
>
> [Kitcher 1993: 185, emphasis added]

So cognitive processes are belief-forming and sustaining processes and, according to Kitcher, we are to address the central problem of epistemology 'by describing [such] processes that are reliable' [Kitcher 1992: 75–6], where reliability is a matter of generating true beliefs. Thus, Kitcher's brand of naturalism concerns itself with the reliability (truth-conduciveness) of belief-forming and sustaining processes. That Goldman [1976; 1986a; 1986b] and Kornblith [1980] are approvingly cited by Kitcher [1984: 18–19; 1992: 65; 1993: 184] lends credence to this reading.

I prefer to take cognitive processes as widely individuated, in the sense that they are not merely or wholly internal processes. This seems to conform with Kitcher's understanding of cognitive processes. For instance, he writes: 'People can make cognitive mistakes, perceiving badly, inferring hastily, failing to act to obtain inputs from nature that would guide them to improved cognitive states' [Kitcher 1993: 185–6]. So perceiving and acting to acquire information from one's environment count as cognitive processes; barring solipsism, both involve more than just internal, psychological states. Moreover, widely individuating cognitive processes adds to the cogency of the question of the reliability of cognitive processes. The reliability of perceptual faculties, for instance, depends in part on the (external) conditions in which they are exercised or applied. Seeing in good light from not too far away tends to produce true beliefs; seeing in poor light at a great distance does not. The reliability of inferential processes is sensitive to the conditions in which they are applied, as well. Though *modus ponens*, for example, is necessarily truth preserving, it is only contingently reliable. Inferring truths using *modus ponens* depends on reasoning from true premises (or at least a true conditional premise), and whether or not the premises of a given argument are true is very often a contingent, and external, matter. The point is that since the reliability of many faculties which contribute to belief formation and maintenance is sensitive to conditions of application, inquiring after the reliability of a cognitive process involving such a faculty only makes sense if one takes into account the relevant conditions of application. Widely individuating cognitive processes does this by building those conditions into the processes themselves.[4]

Since Kitcher takes enhancement of understanding to be the essence of explanation [1976; 1981; 1985; 1989] and the 'central problem' of naturalistic epistemology involves understanding the reliability of belief-forming and sustaining processes, Kitcher's naturalism is committed to explaining the reliability of these processes. Notice, though, that the norms governing these processes are implicit in the processes themselves. Widely individuated cognitive processes incorporate conditions of application. We (often tacitly) take those conditions into consideration when forming beliefs and 'deciding' whether or not to maintain them, and in doing so we (often tacitly) deploy belief-regulating norms. The correctness of these norms is intimately connected with the reliability of the processes they ratify: correct belief-regulating norms are just those that ratify reliable belief-forming and sustaining processes.[5] Given this, the commitment of Kitcher's naturalism to explain the reliability of belief-forming and sustaining processes is also a commitment to explain the correctness of our belief-regulating norms.

Notice that the concern with explaining the reliability of belief-forming and sustaining processes is neither unusual, even with respect to mathematics, nor necessarily tied to a reliabilist epistemology. For instance, Penelope Maddy *qua* set-theoretic realist writes that '[e]ven if reliabilism turns out not to be the correct analysis of knowledge and justification . . . the Platonist still owes us an explanation of how and why [an expert set

[4]There is evidence that Kitcher widely individuates processes relevant to knowing [2000: especially §III].
[5]Assuming, of course, that we aim to have true beliefs.

theorist's] beliefs about sets are reliable indicators of the truth about sets' [Maddy 1990: 43]. I take it that beliefs about *F*s being reliable indicators of truth about *F*s is a matter of their being reliable beliefs about *F*s, and that this is a matter of the reliability of the processes by which beliefs about *F*s are formed and sustained. At the other end of the ontological spectrum, Hartry Field raises the issue in much the same way. The epistemological challenge to the mathematical realist, which largely motivates Field's mathematical fictionalism, 'depends on the idea that we should view with suspicion any claim to know facts about a certain domain if we believe it impossible in principle to explain the reliability of our beliefs about that domain' [Field 1989: 233].

That Kitcher's naturalism is committed to (KC_2) can be argued in at least two ways. The first relies on the way in which Kitcher sees the relationship between the epistemology of mathematics and that of science. In short, Kitcher takes the epistemology of mathematics to be part of the epistemology of science generally. For instance, he claims that 'it is possible to argue for a reduction of the epistemology of mathematics to the epistemology of science' [1988: 301] and that 'the problems of epistemology of mathematics reduce to questions in the philosophy of science' [1988: 317]. Moreover, Kitcher holds that as a consequence of his 'views on the nature of mathematics ... mathematical knowledge is similar to other parts of scientific knowledge, and there is no basis for a methodological division between mathematics and the natural sciences' [1989: 423]. What we have on Kitcher's view, then, is mathematics and science generally fitting into a common epistemological framework. So what goes for the epistemology of one goes for the epistemology of the other; after all, they're not actually distinct in kind.

Add to this that Kitcher is committed to a realist conception of truth in the context of epistemology of science. In connection with the question of his commitment to a correspondence theory of truth in his account of scientific knowledge in [1993], he writes:

> [T]here seem to be three possibilities: (A) my account of scientific progress presupposes the correspondence theory, and, since the correspondence theory is unsustainable, that account of progress is wrong; (B) my account of scientific progress presupposes the correspondence theory of truth, and, since the correspondence theory can be sustained, this is unproblematic; (C) my account of scientific progress can be combined with global realist (or not anti-realist) positions that do not involve a correspondence theory of truth. In [*The Advancement of Science*] I hoped to remain agnostic between (B) and (C). However there are clearly places in the book where I presuppose (B), rather than (C).
>
> [Kitcher 1995: 662–3][6]

As recently as [2002], Kitcher advocates a correspondence theory of truth in defending scientific realism. So there is reason to believe that Kitcher has a preference for (B), hence for a robust realist conception of truth in

[6]Kitcher defends a correspondence theory of truth in [1993: 128–33].

epistemology of science. But since, for Kitcher, epistemology of mathematics is a part of epistemology of science, this means that Kitcher's epistemology of mathematics also deploys a realist conception of truth.

The second way of seeing that Kitcher's epistemology of mathematics is committed to a realist notion of truth relies on Kitcher's affinity for reliabilism, especially as espoused by Alvin Goldman [1976; 1986a; 1986b; 1992]. Kitcher's general reliabilist tendencies came out in the discussion of (KC_1) above, as did his endorsement of Goldman-style reliabilism, in particular. Though these occur in the context of Kitcher's epistemology of science, Kitcher expresses the same tendencies and endorsement in the context of his epistemology of mathematics.

Kitcher [1984] begins with some epistemic stage setting. There Kitcher posits, as a general analysis of knowing, that S knows that p if and only if S believes that p, p is true, and S's 'belief that p was produced by a process which is a warrant for it' [17]. Obviously, the understanding of 'warrant' is of central importance here. Kitcher declines to specify how he understands 'warrant', but he identifies Goldman's reliabilism, as presented in [1976; 1986a], as the 'best available account of warrants' [1984: 18, n. 6], an opinion he subsequently maintains [1993: 162, n. 46]. Moreover, Kitcher is quite explicit that an account of mathematical warrant should fit into a general account of warrant: 'By considering mathematical knowledge from a psychologistic perspective, I hope to amass new data which a general account of warrants should accommodate' [1984: 18]. It's reasonable to conclude that, since on Kitcher's view a general account of warrant should apply to mathematical warrant and Goldman's account of warrant is the best going, Kitcher endorses a Goldman-style account of mathematical warrant.

Goldman [1986b] vigorously defends a realist conception of truth for use in his epistemology. There he quite clearly says that it is crucial to his epistemological view that 'truth ... not be an epistemic matter' in the sense that 'what makes a [true (or false) statement] true (or false) is independent of our knowledge or verification' [Goldman 1986b: 143]. In other words, Goldman endorses what he calls *verification-transcendent truth*, according to which the truth value of a statement is independent of 'our knowledge, or verification, of it (or even our *ability* to verify it)' [Goldman 1986b: 143, original emphasis]. This verification-transcendent conception of truth is precisely the realist conception of truth in (KC_2). Moreover, Goldman doesn't simply endorse a realist conception of truth; he argues that it is essential to his epistemological view, i.e., to his account of warrant. Thus, given that Kitcher endorses a Goldman-style account of mathematical warrant, it's reasonable to conclude that Kitcher's epistemology of mathematics similarly requires a realist notion of truth, so that Kitcher's epistemology of mathematics is committed to (KC_2).

One might think that all this talk of realist conceptions of truth in connection with Kitcher's epistemology of mathematics is a non-starter, given the analogy Kitcher draws between mathematics and the ideal theory of gases [1984: esp. chap. 6].[7] The idea is to take statements of arithmetic

[7]Thanks to an anonymous referee for raising this issue.

(and of mathematics, generally) as true and grounded in certain sorts of actual operations performed by human agents, in the same way that the laws of ideal gas theory are true and grounded in facts about actual gases, and deny that the conception of truth at issue is realist, thereby freeing Kitcher of (KC_2). I postpone consideration of this issue until §III, when we will be in a better position to give it the attention it deserves.

II. Naturalistic Constructivism

Kitcher takes as his jumping-off point the recognition 'that mathematical knowledge is a historical product' [Kitcher 1988: 298]. On his view, each generation of mathematicians inherits its discipline from its predecessors, modifies that discipline with its own research, and passes the result to the succeeding generation. To make this process more precise, Kitcher introduces the notion of a *practice*. A mathematical practice P is a quintuple $\langle L^P, K^P, Q^P, A^P, V^P \rangle$, where the first component is the language used by practitioners of P; the second component is the set of statements already accepted by the practitioners of P; the third component is the set of unresolved questions considered worth investigating by the practitioners of P; the fourth component is the set of patterns of argument deployed by the practitioners of P to justify the members of K^P; and the fifth component is the set of views concerning methodological issues of mathematics (e.g., proper methods of proof and definition in mathematics, relative importance of sub-disciplines of mathematics, relationship(s) between mathematics and non-mathematical sciences, etc.) subscribed to by the practitioners of P [Kitcher 1984: 163–4].

Kitcher represents the development of mathematics by a sequence of mathematical practices, and he attempts to frame a naturalistic answer to (ME) in terms of this development. According to Kitcher, mathematical beliefs at a time t are justified just in case they are in K^{P_t} and there is a sequence of practices $P_{t_1}, P_{t_2}, \ldots, P_{t_n}$, such that $P_{t_n} = P_t, P_{t_1}$ is (somehow) empirically grounded, and transitions between practices are (on balance) rational [Kitcher 1988: 299].[8] To complete his account, Kitcher needs to (i) tell us how a mathematical practice can be empirically grounded and (ii) specify 'conditions under which transitions between practices preserve justification' [Kitcher 1988: 300]. Here I ignore (i) and focus on (ii).

Since preserving justification over interpractice transitions depends on the rationality of the transitions, we look to Kitcher's account of the rationality of interpractice transitions for an answer to (ii). Kitcher endorses a means–ends conception of rationality. He begins his account from the 'general thesis' that '[i]nterpractice transitions count as rational insofar as they maximize the chances of attaining the ends of inquiry' [Kitcher 1988: 304]. Kitcher recognizes two types of ends of inquiry: rational ends (which he calls *ends of rational inquiry*) and practical ends. The former consist in 'achieving

[8]Here P_t is mathematical practice at time t. The careful reader will note that I blur the distinction between beliefs and (accepted) statements with respect to the members of K. This is unproblematic for present purposes.

of truth and the attainment of understanding' [Kitcher 1988: 305], while the latter include 'the goals of providing for the welfare of present and future members of our species (and perhaps members of other species as well), of securing free and just social arrangements, and so forth' [ibid.].

Kitcher distinguishes internal and external interpractice transitions, and the distinction between rational and practical ends enables him to specify conditions under which each kind of transition is rational. Let $P = \langle L^P, K^P, Q^P, A^P, V^P \rangle$ be a mathematical practice, and suppose that some component or components of P are modified to yield a new practice P'. If P is modified in response to purely mathematical considerations, then the transition from P to P' is *internal*. If, on the other hand, P is modified in response to considerations which are in part non-mathematical—e.g., pressure from physics to provide a model for some newly observed phenomenon—then the transition from P to P' is *external*. The rationality of an internal interpractice transition consists in 'its advancement of the ends of rational inquiry *in mathematics*', while the rationality of an external interpractice transition consists in 'its advancement either of the ends of rational inquiry in some other branch of knowledge or of some practical ends' [Kitcher 1988: 306, original emphasis].

Mathematical progress is measured in terms of 'movement' towards the epistemic ends of mathematics. One might reasonably think that chief among these ends is adding to our stock of mathematical knowledge (i.e., modifying a practice P to a practice P' in such a way that $K^P \subset K^{P'}$).[9] But clearly this can be done in response to purely mathematical considerations or in response to considerations that are in part non-mathematical. In other words, both internal and external transitions can add to our stock of mathematical knowledge. One who counts the growth of mathematical knowledge as mathematical progress *no matter whether that growth is produced by an internal or an external interpractice transition*, so long as the transition is rational, endorses what Kitcher calls a *liberal conception* of mathematical progress [Kitcher 1988: 311].

According to Kitcher, the liberal conception of mathematical progress is committed to there being two dimensions along which mathematical progress can be assessed, one epistemic and the other pragmatic.[10] Mathematical progress made as a result of an internal interpractice transition is *epistemically* praiseworthy, as it is a consequence of trying to meet a goal internal to mathematics. Mathematical progress made as a result of an external interpractice transition, on the other hand, is *pragmatically* praiseworthy, as it is a consequence of trying to meet a goal that is at least partially external to mathematics.

Kitcher suggests that the primary motivation for embracing the liberal conception of mathematical progress is a Platonist ontology of mathematics. For on the Platonist's view, '[t]he mathematician's task is to draw a map of Platonic heaven, and the acquisition of any "geographical" information constitutes progress' [Kitcher 1988: 311–12]. As Kitcher rejects

[9]Here '$\subset$' means proper subset.
[10]This terminology is Kitcher's.

Platonism about the ontology of mathematics, he also rejects the liberal conception of mathematical progress. He argues that in 'the context of a naturalistic approach to mathematical knowledge' [Kitcher 1988: 312] the Platonist conception of mathematics encounters problems beyond the well-known Benacerraf-style worries concerning our epistemic access to entities which exist outside the causal network. In particular:

> Like other theoretical realists, Platonists must explain why our ability to systematize a body of results provides a basis for belief in the existence of antecedently unrecognized entities.
>
> [Kitcher 1988: 312]

As long as mathematical ontology transcends mathematical practice in the way Platonism holds it does, no such explanation is likely to be forthcoming on Kitcher's view. But even if we jettison Platonism, the question does not become moot; it simply devolves onto the account of mathematical ontology substituted for Platonism in a slightly modified form. Instead of needing to explain why systematization justifies beliefs about previously unknown entities, we need to explain why systematization justifies new beliefs about previously acknowledged entities.

With this in mind, Kitcher opts for a quite different conception of mathematical ontology, one that closes the gap between mathematical ontology and mathematical practice. According to Kitcher,

> we should treat mathematics as an idealized science of human operations. The ultimate subject matter of mathematics is the way in which human beings structure the world, either through performing crude physical manipulations or through operations of thought. We idealize the science of human physical and mental operations by considering all the ways in which we could collect and order the constituents of our world if we were freed from various limitations of time, energy, and ability.
>
> [Kitcher 1988: 313]

So on Kitcher's view, the ontology of mathematics comprises (at least) operations we are actually able to perform on objects—'operations of collection, correlation, and so forth' [Kitcher 1984: 117]—as well as operations that we would be able to perform were we not limited in various ways. Performing and observing the performance of particular operations of the former sort yields empirically grounded, rudimentary 'proto-mathematical knowledge' which Kitcher takes to be 'epistemically unproblematic' [Kitcher 1984: 117]. Early mathematical practice or practices are filled out in an attempt to systematize this rudimentary knowledge, the point of which is to help us better understand the operations that gave rise to that knowledge. This yields a new practice, which is filled out in an attempt to systematize the knowledge codified in this practice. And so on. Thus, in addition to physical operations actually performed or performable on objects (or mental representations of them) or performable on such objects *modulo* certain physical limitations, we have in our mathematical ontology operations on those operations, operations on the operations on those operations, etc.

This proposal—that we begin with a practice that describes and systematizes our rudimentary shufflings of physical objects and that transitions between practices 'consist in introducing concepts, statements, problems, and reasonings' in order to 'help us understand the operations we are able to perform on our environments' and where this process sometimes 'generates new kinds of operations for later mathematics to consider' [Kitcher 1988: 314]—Kitcher refers to as *Naturalistic Constructivism*. Naturalistic Constructivism is a form of constructivism in that its attendant ontology is in some sense produced by our activity. It lays claim to being a form of naturalism in that it rejects a priori foundations for mathematics.

III. Correctness and Truth

We are supposed to prefer the ontology for mathematics proposed by Kitcher because it closes the gap between mathematical practice and mathematical ontology and so gives us some purchase on the problem of explaining how systematization justifies our mathematical beliefs. In order for Naturalistic Constructivism to honour (KC_1), the justification in question here needs to be truth-conducive. And to honour (KC_2), the notion of truth in play needs to be realist. So we would like to know whether or not systematization reliably produces mathematical beliefs, where the conception of truth involved in the reliability of systematization is realist.

Concerning mathematical truth, Kitcher writes:

> To say that a mathematical statement is true is to make a claim about the powers that are properly attributed to the ideal subject (or, more generally, to make a claim to the effect that the statement figures in a story that is properly told). What 'properly' means here is that, *in the limit of the development of rational mathematical inquiry, our mathematical practice contains that statement*. Truth is what rational inquiry will produce, in the long run.
>
> [Kitcher 1988: 314, emphasis added]

So on Kitcher's view, a mathematical statement σ is true just in case $\sigma \in K^{P_\infty}$,[11] where P_∞ is the practice at the limit of rational mathematical inquiry. Now Kitcher needs an argument to the effect that the regulative norms endorsed by Naturalistic Constructivism,[12] chiefly the direction to systematize, produce (mostly) true statements in the limit of mathematical inquiry.

Kitcher thinks he has such an argument. According to Kitcher, systematization yields 'a practice in which one has achieved a *unifying* perspective on the results, questions, or reasonings previously regarded as disparate. The new practice is justified in virtue of its providing this unified perspective' [Kitcher 1984: 217–18, emphasis added]. Thus Kitcher assimilates systematization to theoretical unification, which he judges is 'important to us because the unification of a field enhances our understanding of it'

[11]Recall that K^P is the set of statements accepted by the practitioners of P.
[12]I use 'regulative norms' as a synonym for 'belief-regulating norms'.

[Kitcher 1984: 218]. Enhanced understanding is important because (recall) the attainment of understanding is one of the ultimate ends of inquiry. But in the current context this is especially important. For on Kitcher's view, there is no concept of mathematical truth which is independent in the sense that it 'stands apart from the rational conduct of inquiry and our pursuit of nonmathematical ends, both epistemic and nonepistemic' [Kitcher 1988: 315], and '[s]ince there is no independent notion of mathematical truth, *the only epistemic end in the case of mathematics is the understanding of the mathematical results so far achieved*' [Kitcher 1988: 314–15, emphasis added].

For Kitcher, the rationality of an interpractice transition increases with the likelihood that the transition moves us closer to the ends of inquiry. Generally these ends are truth and understanding. But truth drops out as an end in the mathematical context. Or, as Kitcher puts it: 'Naturalistic constructivism *collapses the notions of justification and truth* in an interesting way' [Kitcher 1988: 314, emphasis added]. According to Naturalistic Constructivism, at the end of mathematical inquiry the class of justified mathematical statements (beliefs) is the same as the class of true mathematical statements (beliefs), viz., K^{P_∞}. This would be highly desirable if we had notions of truth and justification that were independent of one another. But what we have here appears to be a notion of truth that depends on a notion of justification: being true appears to *consist in* being justified in the limit of mathematical inquiry guided by the directive to systematize.

There is a question here about how we should understand the claim that mathematical truth is what we obtain in the limit of rational mathematical inquiry. This claim can be cast as:

(†) For every mathematical statement σ, σ is true just in case $\sigma \in K^{P_\infty}$.

It is natural to ask what makes (†) true (if indeed it is). Specific to the case at hand, we might answer:

(A_1) (†) is true because what makes a mathematical statement σ true is its being a member of K^{P_∞}, i.e., its being such that we eventually come to have a settled belief that σ by following the directive to systematize.[13]

Or we might answer:

(A_2) (†) is true because the directive to systematize is such that, for every true mathematical σ, following this norm will eventually result in a settled belief that σ.

Is either of these answers open to Kitcher in the light of his commitments to (KC_1) and (KC_2)?

Consider (A_1). According to this answer, mathematical statements have their truth values in virtue of our rationally believing them (or not) in the

[13]I treat 'acceptance' and 'belief' as synonyms.

long run, rational belief being a matter of believing in accordance with the directive to systematize. Our rationally believing a mathematical statement σ (in the long run) *makes* σ true. This is clearly an epistemic conception of truth of the sort Goldman explicitly rejects as inadequate for purposes of a reliabilist conception of justification [Goldman 1986b: chap. 7, esp. §7.2]. Since honouring (KC_1) involves explaining the reliability of cognitive processes licensed by belief-regulating norms (whether or not Kitcher is a full-blown reliabilist), and Kitcher endorses Goldman's conception of reliability, Kitcher cannot accept (A_1) and maintain his commitment to (KC_1). In addition, as the conception of truth deployed in (A_1) is not realist, Kitcher also cannot accept (A_1) and maintain his commitment to (KC_2). Thus, so long as Kitcher is committed to (KC_1) or (KC_2) he cannot accept (A_1). What about (A_2)?

According to (A_2), the regulative norms of mathematical inquiry are such that, in the long run, every truth of mathematics will be believed by mathematicians following those norms. We can grant that the notion of truth at issue here is adequate to (KC_2), so whatever makes mathematical propositions true is independent of the minds and epistemic activities of mathematicians. But why think that the regulative norms of mathematical inquiry will ultimately lead us to truth in this sense? One needs an argument here, an argument for the correctness of the regulative norms of mathematics. This, however, is just the sort of argument we thought we were getting, viz., an argument for the correctness of the regulative norms guiding mathematical practice according to Naturalistic Constructivism. We are in danger of begging the question. To honour (KC_1), Naturalistic Constructivism needs to explain the correctness of the regulative norms of mathematics, which involves explaining the reliability of the cognitive processes licensed by those norms. The naturalistic constructivist's explanation of this reliability relies on (†), which in turn relies on (A_2) as a result of Kitcher's commitment to a realist conception of truth. But the explanatory challenge posed in (KC_1) reasserts itself in (A_2): (A_2) does not explain the correctness of the belief-regulating norms operative in mathematics; it merely asserts it. Naturalistic Constructivism honours (KC_2) only at the expense of honouring (KC_1).

Kitcher might respond by arguing that his conception of mathematical truth is a hybrid of the epistemic conception rejected above and the realist conception that seems to push him into begging the question, arguing further that being neither purely epistemic nor so robustly realist allows his conception of truth to thread a course between (A_1) and (A_2). In a certain sense, we can understand the conception of truth Kitcher deploys as realist: the operations that figure in the mathematical facts according to Naturalistic Constructivism are real, in the sense that whatever ontological status they have doesn't depend on us or our cognitive activity,[14] and Kitcher might say that mathematical statements have the truth values they do in virtue of facts about these operations and the facts about these

[14]For the sake of argument, I grant that no problems arise for the ontological status of these operations as the process of taking operations on operations on operations, etc. is iterated.

operations are sufficiently objective, i.e., sufficiently independent of the minds and epistemic activities of mathematicians, to underwrite a notion of truth adequate to an account of the reliability of the relevant cognitive processes. I have two worries about this response.[15]

First, if facts concerning the relevant operations are objective enough to ground a notion of truth that is sufficiently realist, in the sense that it's nothing about what mathematicians think or do that makes mathematical propositions true or false, we get essentially the same problem with accounting for the correctness of regulative norms that came up in the discussion of (A_2) above. This turns on the ontological status of the operations Kitcher takes to be the subject matter of mathematics. If those operations are 'out there' among the possibilia, existing independently of us and our mathematical activities, then facts about these operations can likely ground an appropriately realist conception of truth. But then why think that our mathematical activities get us on to those operation-facts? That is, why think that the finitely many operations we ever carry out on the finitely many objects in our environment are a good guide to what is the case with the proper-class many possible operations that, according to Kitcher, constitute the subject matter of mathematics in all its richness? We have an underdetermination problem. If, on the other hand, the relevant operations are not 'out there' but are instead somehow generated by our mathematical activity, then it's hard to see how facts about them can ground an appropriately realist conception of truth, since, in this case, mathematical propositions are true in virtue of operation-facts which themselves depend on us and our mathematical activity. Kitcher seems to suggest the first of these options in some places [1984: 120–2] and the second in others [1988: 313–14]. However, neither option is satisfactory given the commitments of his naturalism. This brings me to my second worry.

One sometimes has the impression that Kitcher intends an intermediate position according to which our mathematical activity in some sense gives rise to the relevant operations, yet facts about those operations do not depend on our activity. He invites us to 'conceive of mathematics as a collection of stories about the performances of an ideal subject to whom we attribute powers in the hope of illuminating the abilities we have to structure our environment' [Kitcher 1988: 313]. On this view operation-facts are codified in the stories that constitute mathematics and, though we have considerable leeway in constructing those stories, there are limits on what counts as an acceptable story. Those limits induce limits on operation-facts, hence facts about the relevant operations gain some independence from our mathematical activity (i.e., our story construction).

It is natural to ask after the source and status of the factors constraining story construction. One such factor is that stories must be compatible with what we might call *ground-level* stories, i.e., those stories that codify operations we are actually able to perform on our environment.[16] This somewhat restricts the types of stories that are acceptable. For instance, no

[15]Cf. the discussion of truth and the ideal gas theory at the end of this section.

[16]The compatibility at issue here requires at least logical or conceptual consistency. I bracket the question of what it may additionally require.

story according to which the operation O' is a successor operation of O and operations O and O' are matchable[17] will be acceptable. But at the same time, this compatibility condition does relatively little to constrain the development of set theory. To push past $V_{\omega+\omega}$—where set surrogates for the applied parts of mathematics arguably live [Feferman 1993]—some other limiting factor on the acceptability of stories is required. On this, Kitcher does not say much, and what he does say is of dubious worth to him in the present context.

In characterizing mathematical truth, Kitcher indicates that one way to understand a claim that some mathematical statement σ is true is as a claim that σ 'figures in a story that is properly told' [Kitcher 1988: 314]. Presumably, a properly told story is an acceptable story, and vice versa. If so, then whatever factors determine whether or not a story is properly told will be limiting factors on the acceptability of stories. But as I noted at the top of this section where I set out Kitcher's view of mathematical truth, σ figures in a properly-told story just in case $\sigma \in K^{P_\infty}$. Indeed, appearing in K^{P_∞} is what it *means* for σ to figure in a properly-told story.[18] Thus, it appears that story construction is regulated by the same norms as mathematical inquiry. A properly-told story is one that codifies the operation-facts (statements of which are) found in K^{P_∞}. This, of course, does not resolve anything; it just brings us back to the beginning.

The arguments of this section are variations on a theme. Honouring (KC_1) requires honouring (KC_2): explaining the correctness of belief-regulating norms involves accounting for the reliability of certain cognitive processes, and a notion of truth adequate to such an account will be realist. Attempts to give Kitcher a plausible story for Naturalistic Constructivism's honouring (KC_1) come up short in one of two ways. Accounts that honour (KC_2) rely on already having an account of the correctness of the relevant norms, and so beg the question. Non-question-begging ways of honouring (KC_1) avail themselves of notions of truth which are insufficiently realist. Consequently, non-question-begging ways of honouring (KC_1) fail to honour (KC_2). Hence, the claim that Naturalistic Constructivism provides a naturalistic answer to (ME), on Kitcher's understanding of naturalism, is false.[19]

One might now raise the issue left unaddressed at the end of §I, viz., that despite my arguments in that section to the contrary Kitcher does not endorse a realist conception of truth for mathematics. Recall that this objection rests on an analogy between mathematics and the ideal gas theory. According to Kitcher:

> Arithmetic owes its truth not to the actual operations of actual human agents, but to the ideal operations performed by ideal agents. In other words, I construe arithmetic as an *idealizing theory*: the relation between arithmetic and

[17]Intuitively, this says that adding one object to a finite collection C yields a collection equinumerous with C. The relevant definitions can be found in Kitcher [1984: chap. 6, §III].

[18]For 'means', see the block quotation at the beginning of this section.

[19]One might be tempted to help Naturalistic Constructivism simultaneously honour both (KC_1) and (KC_2) by modifying it to accommodate the claim that mathematics is a priori. However, such a move would conflict with Kitcher's express rejection of mathematical apriorism [1984].

the actual operations of human agents parallels that between the laws of ideal gases and the actual gases that exist in our world.

[1984: 109, original emphasis]

What goes for arithmetic goes for mathematics generally here. So on this view, mathematical truth is of the same sort as truth for the laws of the ideal gas theory. On the latter, Kitcher gives us the following [1984: 116–17].

Consider the Boyle–Charles Law, $PV=RT$.[20] A formalized precise statement of this law is

$$\text{(BC)}\ \forall x[G(x) \rightarrow (P(x) \times V(x) = R \times T(x))]$$

As Kitcher notes, if the extension of 'G' includes actual gases, then (BC) is false: actual gases don't obey (BC). However, if we restrict the extension of 'G' to ideal gases, then (BC) is true, not because there are ideal gases and they obey (BC), but because there are no ideal gases. In other words, (BC) construed as about only ideal gases is vacuously true. Kitcher takes this second reading as the correct one, for ideal gas theory and for mathematics [1984: 117, n. 18].

So the account of mathematical truth on offer has it that mathematical claims are vacuously true in virtue of being about operations performed by ideal agents, of which there are none. Call this the *ideal theory of truth* (or just the *ideal theory*). On analogy with (BC), considered as a statement of the truth conditions of the Boyle–Charles Law, we can use the axiomatization of arithmetic given by Kitcher [1984: 113–14, (1)–(15)], so-called *Mill Arithmetic*, to formalize the truth conditions of statements of arithmetic. For simplicity, we consider the truth conditions of the axioms themselves. The details of the axiomatization need not detain us. Let L_K be the language of Kitcher's axiomatization plus the unary predicate 'I'. Intuitively, '$I(x)$' holds of all and only ideal operations (i.e., operations performed by ideal agents). For any statement σ of Mill Arithmetic, σ^I is the relativization of σ to 'I', i.e., the result of relativizing the quantifiers of σ to 'I' in the standard way [Chang and Keisler 1973: 242]. All but two of the axioms of Mill Arithmetic are universally quantified statements in prenex normal form (i.e., all quantifiers at the front). The truth conditions for those axioms according to the ideal theory are straightforwardly given by their respective relativizations to 'I'. For example, axiom (2) says that $\forall x \forall y[M(x,y) \rightarrow M(y,x)]$.[21] The relativization of this to 'I' is:

$$(2^I) \quad \forall x[I(x) \rightarrow (\forall y[I(y) \rightarrow (M(x,y) \rightarrow M(y,x))])].$$

Since there are no ideal agents, there are no ideal operations. So (2^I), and hence (2), is vacuously true, just as (BC) is. Does the ideal theory get Kitcher off the hook with respect to (KC2)? I don't think so.

[20]This says that pressure times volume is equal to the so-called 'ideal gas constant' times temperature.
[21]For present purposes it doesn't matter what 'M' is.

First, consider one of the axioms that is not a prenexed universal quantification, in particular the one which says that $\exists x U(x)$.[22] This is (13) in Kitcher's axiomatization, and its relativization to '*I*' is:

$$(13^I) \quad \exists x[I(x) \wedge U(x)].$$

Since there are no ideal operations, (13^I) is false. So whether or not the conception of truth in play here is realist, it's inadequate to Kitcher's needs. I think Kitcher may have a reply to this problem in that *U*-type operations only involve segregating a single object, something that actual human agents can clearly accomplish. This means that Kitcher doesn't need to rely on ideal agents for the truth of (13). Notice, though, that such a move would yield a bifurcated conception of truth for mathematics, some statements falling under one conception of truth and others falling under a different conception of truth. Such a situation is prima facie unsatisfying, but I won't pursue the point here.

Second, notice that the ideal theory of truth isn't obviously non-realist, at least not in the operative sense. Recall from §I that we are concerned with verification-transcendent truth: statements have the truth values they do independently of us and our cognitive activities. According to the ideal theory of truth, mathematical statements have the truth values they do in virtue of facts concerning ideal operations, in particular the fact that there are none. But this fact doesn't depend on us or our activities. Hence, the ideal theory appears to be verification-transcendent, i.e., realist.

One might think otherwise, given that Kitcher talks about the powers of ideal agents, and hence ideal operations, in some sense being stipulated by us via our axiomatizations [1984: chap. 6]. The idea is that the axioms of Mill Arithmetic implicitly define the powers of ideal agents as group axioms implicitly define the class of groups. The relevant stipulations are not entirely conventional; they must be 'appropriately grounded' [Kitcher 1984: 116–17] in experience. This allows Kitcher to avoid mathematical knowledge being a priori [1984: chap. 4], but it is also supposed to introduce an element of objectivity into arithmetic truth. In particular, the notion of truth in play is supposed to deploy reference and satisfaction in the standard (Tarskian) way [Kitcher 1984: 140–1, response to Objection 3]. Thus it appears that the conception of truth at issue is indeed realist. That the non-existence of ideal operations is a matter of fact, independent of us and our cognitive activities, only strengthens this appearance.

Lastly, the ideal theory of truth makes mathematics inconsistent. As Charles Chihara notes [1990: 228–30], not only will universally quantified statements we want to come out true do so on the ideal theory, so will universally quantified statements we don't want to come out true. For example, both of the following statements come out (vacuously) true on the ideal theory.

[22] I ignore the other axiom that is not a prenexed universal quantification (axiom (10)), which is effectively the induction schema.

(P3E) Every prime less than 3 is even

(P3NE) Every prime less than 3 is not even

But (P3E) and (P3NE) plus

(2UP3) 2 is the unique prime less than 3

entail

(‡) 2 is even and 2 is not even.

Since (P3E) and (P3NE) are both true according to the ideal theory, and (2UP3) had better be if the ideal theory stands a chance of being adequate,[23] (‡) is true according to the ideal theory. But then (‡) is in K^P for our current mathematical practice P, and K^P is inconsistent.[24] So not only is it dubious that the ideal theory of truth frees Kitcher from (KC_2), Kitcher's stated project of accounting for current mathematical knowledge [1984: 3] is incompatible with the ideal theory. The ideal theory of truth holds no help for Kitcher's epistemology of mathematics.[25]

IV. Concluding Remarks

The difficulty I have isolated for Naturalistic Constructivism, viz., the tension between honouring both a commitment to explaining the correctness of the regulative-norms of mathematics and a commitment to a realist conception of truth, can be put into a broader context. A similar tension is found in the general project of naturalizing mathematics. In particular, it's unlikely that any epistemology of mathematics can satisfy naturalistic desiderata concerning theory revision and ontology while simultaneously accommodating any theory of mathematical truth, realist or not. I won't defend these claims here. The relevant arguments can be found in my [manuscript b]. My intention is rather merely to note that the problems with Kitcher's specific attempt to naturalize mathematics are pieces of a larger picture, a picture which makes for unpleasant viewing for mathematical naturalists.[26]

Louisiana State University

Received: May 2006
Revised: April 2007

[23]Note that since (2UP3) is part of currently accepted arithmetic practice, Kitcher will not object to this.

[24]This problem arises from Kitcher's use of the material conditional. So one might try to help Kitcher by instead using a subjunctive conditional, though Kitcher himself self-consciously rejects this [Chihara 1990: 231–2, esp. n. 12]. The challenge for one who advocates this move is to provide an account of truth conditions for subjunctive conditionals that are neither realist nor violate Kitcher's empiricist scruples (e.g., by countenancing possible worlds).

[25]Notice that the problems raised for the ideal theory of truth here don't obviously derail Kitcher's epistemology of science more generally. They do present a challenge to his account of knowledge where ideal theories are concerned, but I'm aware of no reason to think that the bulk of theoretical knowledge on Kitcher's view concerns ideal theories.

[26]Thanks to Richard Boyd, Jon Cogburn, Eric Gilbertson, Harold Hodes, Brendan Jackson, Richard A. Shore, Zoltán Szabó, Jessica Wilson, and anonymous referees for this journal for valuable discussion and comments.

References

Chang, Chen Chung and H. Jerome Keisler 1973. *Model Theory*, Amsterdam: North-Holland.
Chihara, Charles 1990. *Constructibility and Mathematical Existence*, Oxford: Clarendon Press.
Feferman, Solomon 1993. Why a Little Bit Goes a Long Way: Logical Foundations of Scientifically Applicable Mathematics, *Proceedings of the Philosophy of Science Association,* 2. 1993: 442–55.
Field, Hartry 1989 (1988). Realism, Mathematics, and Modality, in *Realism, Mathematics, and Modality*, Oxford: Basil Blackwell: 227–81.
Goldman, Alvin 1976. Discrimination and Perceptual Knowledge, *Journal of Philosophy* 73: 771–91.
Goldman, Alvin 1986a (1979). What is Justified Belief?, in *Empirical Knowledge*, ed. Paul K. Moser, Savage: Rowman and Littlefield Publishers: 171–92.
Goldman, Alvin 1986b. *Epistemology and Cognition*, Cambridge: MA: Harvard University Press.
Goldman, Alvin 1992. Epistemic Folkways and Scientific Epistemology, in *Liaisons: Philosophy Meets the Cognitive Sciences*, Cambridge: MA: MIT Press: 155–75.
Kitcher, Philip 1976. Explanation, Conjunction, and Unification, *Journal of Philosophy* 73: 207–12.
Kitcher, Philip 1981. Explanatory Unification, *Philosophy of Science* 48: 507–31.
Kitcher, Philip 1984. *The Nature of Mathematical Knowledge*, New York: Oxford University Press.
Kitcher, Philip 1985. Two Approaches to Explanation, *Journal of Philosophy* 82: 632–9.
Kitcher, Philip 1988. Mathematical Naturalism, in *Essays on the History and Philosophy of Modern Mathematics, Minnesota Studies in the Philosophy of Science*, ed. W. Aspray and P. Kitcher, 11: 293–328.
Kitcher, Philip 1989. Explanatory Unification and the Causal Structure of the World, in *Studies in Explanation, Minnesota Studies in the Philosophy of Science*, ed. P. Kitcher and W. Salmon, 13: 410–505.
Kitcher, Philip 1992. The Naturalist's Return, *Philosophical Review* 101: 53–114.
Kitcher, Philip 1993. *The Advancement of Science*, New York: Oxford University Press.
Kitcher, Philip 1995. Author's Response, *Philosophy and Phenomenological Research* 55: 653–73.
Kitcher, Philip 2000. A Priori Knowledge Revisited, in *New Essays on the A Priori*, ed. P. Boghossian and C. Peacocke, Oxford: Oxford University Press: 65–91.
Kitcher, Philip 2002. On the Explanatory Role of Correspondence Truth, *Philosophy and Phenomenological Research* 64: 346–64.
Kornblith, Hilary 1980. Beyond Foundationalism and the Coherence Theory, *Journal of Philosophy* 72: 597–612.
Maddy, Penelope 1990. *Realism in Mathematics*, Oxford: Clarendon Press.
Plantinga, Alvin 1993a. *Warrant and Proper Function*, New York: Oxford University Press.
Plantinga, Alvin 1993b. *Warrant: The Current Debate*, New York: Oxford University Press.
Roland, Jeffrey W. manuscript a. How To Be a (Thoroughgoing) Naturalist.
Roland, Jeffrey W. manuscript b. On the Coherence of Naturalizing Mathematics.