

HOW TO DO THINGS WITH AN INFINITE REGRESS

Kevin T. Kelly
Department of Philosophy
Carnegie Mellon University

ABSTRACT

Scientific methods may be viewed as procedures for converging to the true answer to a given empirical question. Typically, such methods converge to the truth only if certain empirical presuppositions are satisfied, which raises the question whether the presuppositions are satisfied. Another scientific method can be applied to this empirical question, and so forth, occasioning an empirical regress. So there is an obvious question about the point of such a regress. This paper explains how to assess the methodological worth of a methodological regress by solving for the strongest sense of single-method performance that can be achieved given that such a regress exists. Several types of regresses are “collapsed” into corresponding concepts of single method performance in this sense. The idea bears on some other issues in the philosophy of science, including Popper’s falsificationism and its relationship to Duhem’s problem.

CONFIRMATION AND NATURALISM

Here is a familiar but unsatisfying approach to the philosophy of science. Science seeks to “justify” empirical hypotheses. Usually, evidence does not and never will entail them, so they must be “justified” in some weaker way. So there must be a relation of “partial support” or “*confirmation*” or “empirical rationality” falling short of full (deductive) support that justifies them. The principal task of the philosophy of science is to “explicate” the relation of empirical justification from practice and from historical examples. Any feature of scientific method or procedure that is not derived from this relation is extraneous to the philosophy of science *per se*, although it may be of tangential psychological, sociological, or purely computational interest. Thus, virtues such as confirmation, explanation, simplicity, and testing are normatively and philosophically relevant, but the logic of discovery (procedures for inventing new hypotheses) and procedural efficiency are beside the point (e.g., Laudan 1980).

The trouble with this approach is that explicating the justification relation (supposing it to be possible at all) does not begin to explain why justification *should* be as it is rather than some other way. Convincing, *a priori* answers are not forthcoming and attempts to provide them are no longer in fashion. One responds, instead, with the *naturalistic* view that if scientific standards are to be justified, that justification must *itself* be scientific (i.e., empirical). The next question is how scientific reasoning can justify itself. Circular justifications are more popular than infinite regresses of justification in the philosophical literature (somehow an infinite regress of justifications never “fires” or “gets started”), but it is hard to explain what the point of circles or regresses of justification could possibly be without first knowing what the point of justification, itself, is. And the

Formatted: Right: 0.25"

familiar, confirmation-theoretic philosophy of science under consideration provides no such explanation.

A PROCEDURAL PARADIGM

Consider an alternative paradigm for the philosophy of science, according to which scientific methods are *procedures* aimed at *converging to correct answers* rather than relations between hypotheses and finite bodies of evidence. Procedures are justified not by embodying some abstract relation of empirical justification between theory and evidence at every stage, but because they find correct answers both reliably and efficiently. Computational efficiency is relevant, since it brings one to the truth faster. The logic of discovery is also relevant, because the concept of convergence to a correct answer applies as much to methods producing hypotheses as to methods assessing given hypotheses. This is the perspective of *computational learning theory*, an approach to inductive inference proposed by Hilary Putnam (1963, 1965) and E. M. Gold (1965, 1967) and subsequently developed largely within computer science.¹

Here is a more precise formulation of the idea. Empirical methods are procedures or dispositions that receive successive inputs from nature and output successive guesses in response. Like computational procedures, inductive methods may be judged as solutions to problems. An *empirical problem* is not a particular situation but a *range of serious possibilities* over which the method is required to succeed. Success in a possibility means converging to a correct answer on the stream of inputs that would be received if that possibility were actual. Correctness may be truth or something weaker, such as empirical adequacy. It may also involve pragmatic components, such as being a potential answer to a given question; or, following Thomas Kuhn, it might be something like ongoing puzzle-resolving effectiveness in the unbounded future. The precise choice of the correctness relation is not the crucial issue. What does matter is that there is a potentially endless stream of potential inputs and that the standard notions of correctness transcend any finite amount of data. So achieving correctness reliably—in each of a broad range of cases—occasions the problem of induction: that no answer unverified by a finite sequence of inputs is guaranteed to be correct.

There are several different senses of convergent success, some of which are more stringent than others (cf. Kelly 1996). Let a hypothesis be given. It would be fine if a scientific procedure could eventually halt with acceptance or rejection just in case the hypothesis is respectively correct or incorrect. Call this notion of success *decision with certainty*. But some hypotheses are only *verifiable with certainty* (halt with acceptance if and only if the hypothesis is correct) or *refutable with certainty* (halt with rejection if and only if the hypothesis is false). Other hypotheses are only *decidable in the limit*, meaning that some method eventually stabilizes to acceptance if the hypothesis is correct and to rejection otherwise. There are also hypotheses for which it is only possible to stabilize to acceptance just in case the hypothesis is correct (*verification in the limit*) or to stabilize to rejection just in case the hypothesis is incorrect (*refutation in the limit*). Between

¹ For book-length reviews of the technical literature, cf. (Osherson et al. 1986), (Jain et al. 1999). For attempts to relate the ideas to the philosophy of science, cf. (Martin and Osherson 1998), (Kelly 1996), (Kelly 2000), and (Kelly and Glymour 2004).

decision in the limit and verification and refutation with certainty, one may refine the notion of success by asking how many *retractions* are necessary prior to convergence.² Kuhn and others have emphasized the tremendous social cost of retracting fundamental theories and, furthermore, the number of retractions required prior to convergence may be viewed as a notion of convergent success in its own right that bridges the concepts of certainty and limiting convergence with an infinite sequence of refined complexity concepts.

A given, empirical problem may be solvable in one of the above senses but not in another. The best sense in which it is solvable may be said to be its *empirical complexity*. This is parallel to the theory of computability and computational complexity. In fact, the complexity classes so defined are already familiar objects in analysis and computability theory (Hinman 1978). In the philosophy of science, one speaks vaguely of *underdetermination* of theory by evidence. Elsewhere, I have proposed that degrees of underdetermination correspond to degrees of empirical complexity (Kelly 1996, 2000b). That yields a comprehensive framework for comparing and understanding different inference problems drawn from different contexts, as well as a unified perspective on formal and empirical inquiry (Kelly 1996, chapter 6, 7, 8, and 10, Kelly 2004), something that has bedeviled the confirmation-theoretic approach from the beginning.

Many methodological ideas familiar to philosophers of science emerge naturally from the procedural framework just described (cf. Kelly and Glymour 2004). One such idea is *Duhem's problem*, which turns on the observation that individual hypotheses in a scientific theory are refutable only in the context of other "auxiliary hypotheses." The problem is how to assign credit or blame to a hypothesis when falsifying instances may be due to a false auxiliary hypothesis with which the hypothesis has been forced to keep company. Here is how the problem looks from the perspective of learning theory. A hypothesis that is not refutable with certainty may be refutable with certainty *given* some other auxiliary hypotheses, which is the same as saying that the conjunction of the hypothesis with the other hypotheses is refutable with certainty. Indeed, there may be many potential sets of auxiliary hypotheses that make a given hypothesis refutable with certainty.

One can enumerate the possible systems of auxiliaries thought of so far and accept H as long as the first system of $H +$ auxiliaries consistent with receipt of the current inputs is not refuted. If the first such system is refuted, then H is rejected and one selects the first such system consistent with the data and with H . If new systems of auxiliaries are thought of, they can be added to the end of the queue of auxiliaries thought of already. This procedure verifies H in the limit so long as "creative intuition" eventually produces systems of auxiliaries covering all relevant possibilities admitted by H . So verifiability in the limit corresponds to the intuitive epistemic perplexity occasioned by Duhem's problem. That is important, because many issues in the philosophy of science (realism, conventionalism, observability, theory-ladenness, and paradigms) cluster around Duhem's problem.

² Cf. (Kelly and Glymour 2004) for an explanation of Ockham's razor in terms of retraction minimization.

One can (and, I suggest, should) think of the Kuhnian (1970) distinction between “normal” and “revolutionary” science along similar, procedural lines. A “paradigm” is a hypothesis that is not refutable in isolation but that becomes refutable when “articulated” with auxiliary hypotheses. Normal science involves the selection of auxiliary hypotheses compatible with the paradigm and with experience that make the paradigm refutable. Revolutionary science involves choice among paradigms. The crisp, stepwise solvability of “normal” problems reflects the constraints imposed by the presumed paradigm. Revolutionary science is far less crisp, since each paradigm can be articulated in an infinite variety of ways.³

The preceding points are illustrated naturally and concretely when the hypothesis in question concerns a *trend*. Questions about trends often generate controversy, whether in markets or in nature, because any finite set of evidence for the trend *might* be a local fluctuation around an unknown equilibrium. This sort of ambiguity permeated the debate between uniformitarian and progressionist geologists in the nineteenth century (cf. Ruse 1979). Progressionists⁴ held that geological history exhibits progress due to the cooling of the Earth from its primordial, hot state, whereas Lyell reinterpreted all apparent trends as local fluctuations on an immensely expanded time scale. If progressionism is articulated with a particular schedule of appearance for finitely many fossils, it can be decided with two retractions starting with “no”. Just say “no” until all the fossil types are seen to appear in the fossil record as early as anticipated (remember that it may take arbitrarily long to find a fossil that appears as early as anticipated); then say “yes” until some fossil type is observed to appear earlier than expected, after which say “no” again.⁵ In historical fact, Lyell claimed to have refuted progressionism when the Stonesfield mammals were found in Jurassic strata, prior to progressionist expectations, but it was open to the progressionists simply to “re-articulate” their paradigm with an accelerated schedule accommodating the new find, so progressionism is not really refutable with certainty. In fact, the progressionists were free to accelerate their schedule repeatedly, and no finite set of fossils could possibly refute all possible schedules, so without further reframing, the debate allows for a potentially endless give-and-take. To verify progressionism in the limit, do the following: enumerate the possible schedules of progress (assuming them to be presented as discretely presented rules). Apply the preceding two-retraction method to each schedule. If the first schedule for which the method says “yes” does not change when a new observation is made, say that

³ Much more can be said about this (Kelly 2000b). For example, one can also provide a naturalistic account of theory-laden data in learning theoretic terms.

⁴ The progressionists were called “catastrophists” because the early cooling of the Earth was supposedly accompanied by catastrophic changes unobserved nowadays.

⁵ Attentive readers may have noticed that if the progressionists were to always posit exactly the currently observed earliest appearances for each fossil type then at most one retraction is required per re-articulation. That would be true if progressionism were flexible enough to articulate with arbitrary schedules. But the progressionist paradigm also included prior ideas about which fossil forms were more “advanced” than others, so if a “rudimentary” form were to appear earlier than expected, still more rudimentary forms would have to appear even earlier than that, and it is possible that no such examples had yet been found. In that case, progressionists would have to set a new schedule in which some rudimentary forms are expected earlier than the earliest known examples. And then the right rule is to say “no” until such examples are found, “yes” after they are found, and “no” after still earlier examples are found.

progressionism is true. Otherwise, say that progressionism is false. If progressionism is true, it is true in virtue of some schedule. Eventually, all the fossil types appear as soon as predicted and no fossil type ever appears earlier, so from that point onward the two-retraction method stabilizes to “yes” on the true schedule. For each schedule prior to the true one, the method eventually stabilizes to “no” (either because no fossil of a given type ever appeared early enough or because some fossil type is seen too early). So the enumeration method converges to “yes” when the first true schedule’s method has stabilized to “yes” and all prior schedules’ methods have stabilized to “no”. If progressionism is false, then every schedule’s method eventually converges to “no”, so the enumeration method outputs “no” infinitely often.

Uniformitarianism, on the other hand, is *refutable* in the limit: it looks bad as long as the two retraction method says “yes” for a fixed schedule and looks good when the current schedule is refuted.

Global warming (Berger and Labeyrie 1987) provides a more recent example of an awkward trend question. Is the current warming trend a chaotic spike no greater than historical spikes unaccompanied by corresponding carbon dioxide doses or is it larger than any historical spike unaccompanied by current carbon dioxide levels? Newly discovered high spikes in the glacial record make us doubt global warming and increasing temperatures higher than discovered spikes in the glacial record make us more confident that carbon dioxide levels are the culprit.

In a different domain, the cognitivist thesis that human cognition is computable is verifiable in the limit, for similar reasons. Each finite chunk of behavior is compatible with some computer program (cognitive theory), but each such theory is outrun eventually by uncomputable behavior. To verify the hypothesis of computability, enumerate all possible computable accounts and reject the computability hypothesis only when the first program compatible with human behavior is refuted. If behavior is computable, eventually the right program is first and the method converges to “yes”. Otherwise, each program is eventually refuted and the method says “no” infinitely often.

PROCEDURAL REGRESSES

The procedural outlook just described is subject to its own empirical regress problem. Many empirical problems are solvable, even in the limit, only if certain empirical presuppositions are satisfied. For example, knowing that a curve is polynomial allows one to infer its degree in the limit (from increasingly precise data), but if the question is expanded to cover infinite series, the answer “infinite degree” is only refutable in the limit. If empirical presuppositions are necessary for success, how can one determine whether they are satisfied? By invoking another method with its own empirical presupposition? And what about that one? So it seems that one is left with a regress of methods checking the presuppositions of methods checking the presuppositions of methods.... The point of a method guaranteed to converge to the truth is fairly clear. But what is the point of a regress of methods, each of which succeeds only under some material presupposition that might be false?

Formatted: Right: 0.25"

The basic idea developed in this article is a methodological *no free lunch principle*: the value of a regress can be no greater than the best single-method performance that could be achieved by looking at the outputs of the methods in the regress rather than at the data themselves. If the performance of the best such procedure is much worse than what could be achieved by looking at the data directly, one may justifiably say that the regress is methodologically *vicious*. If the best method that looks only at the outputs of the methods in the regress succeeds in the best feasible sense, then the regress is *vindicated*.

FINITE REGRESSES

Consider the empirical problem (H_0, K) of determining the truth of a given hypothesis H_0 over serious possibilities K . Fix a given sense of success (e.g., refutation with certainty, verifiability in the limit, etc.) Every method M_0 directed at assessing H_0 succeeds in the given sense over some set (possibly empty) of serious possibilities (input streams). The *empirical presupposition* H_1 of a given method M_0 for assessing H_0 is just the set of all serious possibilities (input streams) over which M_0 succeeds (i.e., H_1 is just the empirical proposition “ M_0 will succeed in the specified sense”). So let meta-method M_1 be charged with assessing the presupposition H_1 of method M_0 . Meta-method M_1 reads from the same input stream as M_0 , but instead of trying to determine the truth of the original hypothesis H_0 , M_1 tries to determine the truth of H_1 , the empirical presupposition of M_0 . With respect to the question H_1 , M_1 has its own empirical presupposition H_2 , of which M_2 determines the truth value of under empirical presupposition H_2 , and so forth.

For example, let H_0 denote Lyell’s uniformitarian hypothesis, discussed earlier. After the Stonesfield find, Lyell declared victory for H_0 , which might be interpreted as *halting* definitively with “yes”.⁶ Lyell would also have had a reason to scoff at the progressionists if fossil types expected by a certain epoch (e.g., “missing links”) stubbornly refused to appear. On the other hand, perfect correspondence between the proposed schedule of progress and the actual fossil record would hardly be happy news for Lyell. So one might crudely reconstruct Lyell’s method M_0 as something like the following: until the currently fashionable schedule is *instantiated* (i.e., the earliest geological appearance of each fossil type matches the schedule exactly), scoff at the “anomalies” in the progressionist paradigm and say “yes” to uniformitarianism without halting.⁷ While the currently fashionable schedule is instantiated but not refuted, concede “no” without halting. Finally, when the schedule is refuted outright by a fossil that appears ahead of the current schedule (as the Stonefield find did), announce victory (i.e., halt with “yes”). So M_0 retracts at most twice, starting with “yes”, as the geological data pour in. But no possible strategy converges to the truth about uniformitarianism with just two retractions, since any number of successively accelerated schedules of first appearance for the various fossil types might appear perfectly instantiated for arbitrarily long periods of time before being shot down by a new find ahead of schedule. So M_0 finds the truth only under some empirical presupposition H_1 . Assuming that the serious

⁶ Of course, I oversimplify. He declared victory for a tangle of reasons that would defy any elegant logical representation, but the Stonesfield find seems to have been a significant rhetorical blow to progressionism (Ruse 1979).

⁷ More realistically, output “?” until the “anomalous” failure to find the missing fossils percolates into a “crisis”. That detail doesn’t really change anything in the following analysis.

possibilities K at the time were just those compatible with uniformitarianism and progressionism and that the earliest time of appearance in the fossil record is eventually observed for each fossil type, the presupposition H_1 of M_0 is that progressionism *implies* A_1 (where A_1 is the auxiliary hypothesis that fossil types will first appear according to the currently fashionable schedule) and, hence, that uniformitarianism is true if A_1 is not (i.e. $H_1 = H_0 \vee A_1$). For given H_1 , M_0 really does converge to the truth with just two retractions in the worst case and if H_1 is false, M_0 converges to the false conclusion that uniformitarianism is true.

In Kuhnian terms, the auxiliary hypothesis A_1 is an “articulation” of the progressionist paradigm and in the face of the uniformitarian competitor, the Stonesfield find constituted an anomaly for this particular articulation. As Kuhn takes pains to emphasize, the anomaly does not logically compel rejection of progressionism, since the schedule can be revised to accommodate the anomaly. But Lyell’s method M_0 halts with “yes” when A_1 is refuted, so M_0 fails to find the truth when progressionism is true in virtue of some revised schedule.

When the progressionists responded by revising their schedule to a new schedule A_2 , Lyell’s method was called rhetorically into question, for if the new schedule were true, his method would halt with “yes” even though progressionism is true. Since H_1 is an *empirical* hypothesis, Lyell might have responded to the challenge with a meta-method M_1 that checks the truth of H_1 . With A_2 on the table, it would be rhetorically pointless for Lyell to respond with a method that still presupposes H_1 ; he must at least entertain the revised schedule A_2 as a serious possibility, so that the presupposition of the meta-method M_1 is $H_2 = H_1 \vee A_2 = H_0 \vee A_1 \vee A_2$. Now over these extended possibilities, in which possibilities does M_0 succeed? Only in possibilities in which the revised schedule A_2 is false, since Lyell’s premature halting with “yes” upon the refutation of A_1 would be rescued by the falsity of A_2 . So M_1 should say “yes” until A_2 is instantiated, followed by “no” until A_2 is refuted, and should halt with “yes” as soon as A_2 is refuted. Notice that meta-method M_1 is pretty similar in spirit to Lyell’s original method since it still ungenerously entertains only finitely many possible schedules of progress. And like the original method, the meta-method converges to the truth with two retractions starting with “yes”, given that its empirical presupposition is true. The regress can be extended to any finite length, where meta-method M_i has presupposition $H_{i+1} = H_i \vee A_i$.

Say that a finite regress (M_0, \dots, M_n) *succeeds regressively* (relative to empirical problem (H_0, K)) in a given sense (e.g., verification with certainty) just in case there exist propositions H_1, \dots, H_n such that for each i no greater than n :

1. H_{i+1} is the presupposition of M_i with respect to H_i according to the specified sense of success and
2. K entails H_n .

So assuming that the relevant possibilities in the geological example are exhausted by H_n , the Lyellian regress may be said to succeed regressively concerning the uniformitarian

question over serious possibilities $K = H_n$ in the sense of convergence with two retractions starting with “yes”.

But sequential success is a far different matter from success with respect to the original question. How are the two related, if at all? The worry is that infinite regresses, like circles, do nothing at all but beg or forestall the original question under consideration. One way to answer this question is to construct a *regress collapsing function* $\Phi(a_1, \dots, a_n) = a$ that takes a sequence of n possible answers to a single answer. The collapsing rule Φ may be thought of as converting the regress (M_0, \dots, M_n) into a single method M^* such that for each finite input history e ,

$$M^*(e) = \Phi(M_0(e), \dots, M_n(e)).$$

Then if M^* succeeds in some ordinary, single-method sense in problem (H_0, K) , one can say that the regress is no worse in value than a single method that succeeds in that sense. In other words, the *methodological value* of a regress is the best single-method performance that could be recovered from the successive outputs of the constituent methods in the regress without looking at the inputs provided to these methods. The regress is *vindicated* if the best single method performance that can be achieved by collapsing it is also the best possible single method performance. Otherwise, the regress is *vicious* (a term often employed with no clear sense). Viciousness now comes in well-defined degrees, depending on how far short of optimal performance the best-performing regress collapse falls.

For illustrations of vindication and viciousness, turn once again to the Lyellian regress (M_0, M_1) of length 2. Assuming that $K = H_2$, this regress succeeds *regressively* with two retractions starting with “yes”, since that is what each method achieves given its presupposition. Now consider the best sort of method one could build from this regress without peeking at the inputs. Let e be an arbitrary, finite data history compatible with K . When $M_1(e)$ converges to “yes”, whatever $M_0(e)$ converges to is true and when $M_1(e)$ converges to “no”, whatever $M_0(e)$ converges to is false. So it is sensible to define the collapsing rule so that the first answer a_1 is repeated or reversed, depending on whether M_1 answers “no” or “yes”:

$$\begin{aligned} \Phi(a_1, \text{“yes”}) &= a_1; \\ \Phi(a_1, \text{“no”}) &= \text{reverse}(a_1). \end{aligned}$$

Over possibilities in $K = H_2$, the following histories may occur:

1. neither schedule is instantiated;
2. schedule A_1 is instantiated but not schedule A_2 ;
3. schedule A_1 is refuted but schedule A_2 is not instantiated;
4. schedule A_1 is refuted and schedule A_2 is instantiated;
5. both schedules are refuted.

In such an input stream, method $M^*(e) = \Phi(M_0(e), M_1(e))$ retracts at most four times, starting with “yes”. Notice that this is the *sum* of the worst-case retractions of the constituent methods in the regress (suggesting a general pattern) and that the initial answer is the same as that of the constituent methods.

Does M^* achieve the best possible single-method performance in the problem under consideration? Let M be an arbitrary method that converges to the truth about uniformitarianism over serious possibilities in $K = H_2$. Nature can withhold instantiation of both schedules until M is forced (on pain of failing to converge to the truth) to say “yes”, since otherwise both schedules are false, so according to H_2 , uniformitarianism is true. At that point, Nature is free to continue to present data instantiating but never refuting schedule A_1 until M concludes “no”, since otherwise A_1 is true and implies that uniformitarianism is false. Nature is free to continue to present data refuting schedule A_1 without instantiating schedule A_2 (since A_2 is faster than A_1) until M concludes “yes”. Nature is now free to continue to present data refuting schedule A_2 and instantiating schedule A_2 without refuting it until M concludes “no”. Finally, Nature is free to continue to present data refuting A_2 , forcing M to conclude “yes”. So an arbitrary method that converges to the true answer in the problem also requires at least four retractions starting with “yes”. Starting with “no” would require yet another retraction and starting with “?” would still require four (even not counting the change from “?” to “yes” (by arguments similar to the one just given). So the best possible single method performance in this problem is four retractions starting with “yes”. Hence, Lyell’s regress is *vindicated*, since it can be collapsed into a single method with the best possible performance.

Vindication is not trivial. For example, Lyell might have been a lunatic who reversed his answer every day, whereas his meta-method (physician) might have been perfectly rational and said “no” *a priori* concerning insane Lyell’s success. This regress succeeds regressively in the strongest possible sense (each method succeeds over its respective presupposition with no retractions) but it is *entirely vicious* because both methods ignore the data entirely, precluding all attempts to collapse the regress into a method that even converges to the truth in the limit.

The preceding example illustrates that even extremely strong regressive success does not suffice for vindication. That is because the crazy method fails in an unnatural way. Real science loves to “frame” messy questions to appear crisper than they really are by specifying evidential triggers for when to reject a hypothesis or paradigm that is not really refuted (as in the case of Lyell’s identification of progressionism with a particular schedule of progress). This opens the door to some failures to converge to the truth, since the trigger for halting may be premature. But in spite of this obvious risk of failure, reliance on determinate empirical “triggers” for rejection or acceptance has a silver lining: it ensures that failure of the method occurs in an *orderly way* that has implications concerning the truth of the original hypothesis. In the jargon coined above, reliance on evidential triggers links mere *regressive success* to genuine *methodological value*, as in the Lyell example. In that example, the trigger for dumping progressionism is refutation of schedule A_1 , which fails only when progressionism is true (over serious possibilities K

Formatted: Right: 0.25"

$= H_0 \vee A_1 \vee A_2$). Hence, the meta-method's determination that the method fails has a bearing on the original question and that information is exploited by the collapsing function. This is a new explanation of why induction *should* proceed by means of crisp triggers or defaults, for otherwise empirical regresses would be methodologically worthless, as in the case of the insane regress.

More generally, say that method M *converges* with at most n retractions starting with "yes" just in case M never starts with any other answer and never retracts more than n times in *any possible input stream* in K . Due to its reliance on concrete, empirical "triggers", Lyell's method converges with at most two retractions starting with "yes" over all possibilities in K , even though it does not converge to the truth in all of these possibilities.

For concreteness, the discussion so far has focused on a particular example, but the conclusions drawn are far more general, depending only on the logical relationships between the various success criteria. To lift the discussion to this more general, methodological level, let R be a relation between regresses and problems (e.g., convergence and regressive success in a given sense) and let Q be a relation between single methods and problems (e.g., success in some other sense).

Relation R is *methodologically collapsible* to relation Q if and only if for each problem p and for each regress $(M_0(e), \dots, M_n(e))$ satisfying R with respect to p , there exists a collapsing function Φ such that the single method $M^*(e) = \Phi(M_0(e), \dots, M_n(e))$ satisfies Q with respect to p .

It is also interesting to turn tables and investigate whether single-method success can be *stretched* into some notion of regressive success and convergence. A *stretching* function is a mapping $\Psi(a) = (a_1, \dots, a_n)$ from answers to sequences of answers. The *stretching* of method M by Ψ is the regress defined by:

$$(M_0(e), \dots, M_n(e)) = (\Psi(M(e))_0, \dots, (\Psi(M(e))_n).$$

Then

relation Q is *methodologically stretchable* to relation R if and only if for each problem p and for each method M satisfying Q with respect to p , there exists a methodological stretching function Ψ such that the regress $(M_0(e), \dots, M_n(e)) = (\Psi(M(e))_0, \dots, (\Psi(M(e))_n)$ satisfies R with respect to p .

Now one can define methodological equivalence between regressive and single-method performance as follows:

R for regresses is *methodologically equivalent* to Q for single methods if and only if R is collapsible to Q and Q is stretchable to R .

For example, a regress of two methods that converge and succeed sequentially with one retraction starting with "yes" (i.e., a regress of two refuters) is methodologically

Formatted: Right: 0.25"

equivalent to one method that succeeds with two retractions starting with “yes”.⁸ More generally, the pattern hinted at earlier amounts to this:

Proposition 1: The following are methodologically equivalent.⁹

- regressive success and convergence under a finite retraction bound n_i for each constituent method M_i .
- Single method success under the sum of the bounds n_i starting with “no” if an even number of the M_i start with “no” and starting with “yes” otherwise.

That settles the matter for finite regresses of methods with bounded retractions. Moving on to weaker senses of convergence, it is easy to see that any finite regress of methods that succeed regressively and converge in the limit is equivalent to a single method that decides in the limit: just accept if an even number of the methods in the sequence reject and reject otherwise. Regresses of methods that verify in the limit or refute in the limit are not reducible to any of our notions of success and may be thought of as a natural way to build methodological success criteria applicable to more complex hypotheses. The situation simplifies when all of the presuppositions of methods in the regress are entailed by H_0 or by its complement. Then one may speak of an H_0 -entailed or co - H_0 -entailed regress, respectively.

Proposition 2 The following are methodologically equivalent:

- an H_0 -entailed regress (M_0, M_1) such that M_1 refutes [verifies] in the limit the presupposition H_1 of M_0 as a limiting refuter [verifier].
- a single method M that refutes [verifies] H_0 in the limit.

Proposition 3 the following are methodologically equivalent:

- a co - H_0 -entailed regress (M_0, M_1) such that M_1 verifies [refutes] in the limit the presupposition H_1 of M_0 as a limiting refuter [verifier],
- a single method M that refutes [verifies] H_0 in the limit.

⁸ Here’s the trick. Both methods start out with “yes”. Let the constructed method M start with “yes” because M_1 will succeed and M_1 currently says that M_0 will succeed and M_0 now says “yes”. If M_1 ever says “no”, then let M reverse whatever M_0 says because M_1 is right in saying that M_0 is wrong (since M_1 has already used its one retraction and has therefore converged to the truth). At worst, both retract and M retracts once each time. So M retracts at most twice, starting with “yes”. Methodological equivalence requires that one can also produce a regress of two refuters M_0, M_1 , from an arbitrary method M that succeeds with two retractions starting with acceptance. Here’s how to do it. Let M_0 say “yes” until M retracts once and say “no” thereafter. Let M_1 say “yes” until M retracts twice and say “no” thereafter. Let H_1 be the proposition that M_0 successfully refutes H_0 with certainty. That is true just in case M retracts at most once. Thus, M_1 really succeeds in refuting H_1 with certainty over all possibilities M succeeds over, as required.

⁹ The proofs of all the propositions may be found in (Kelly 2000a).

INFINITE REGRESSES

Suppose it is required that *every* challenge to an empirical presupposition be checked empirically, so that there is a potentially *infinite* regress of methods testing the assumptions of methods.... The point of such practice is far less obvious than that of finite regresses, since finite regresses are “anchored” or “founded” by genuine success of the terminal meta-method. Infinite regresses have no final “court of appeals” in this sense to anchor them. Are they, therefore, necessarily vicious? This is no longer a matter of mere philosophical opinion. It is a logically precise question about methodological equivalence that will now be explored.

Recall that in the Lyellian regress, each method covers more possibilities than its predecessor, for a method that did not cover more possibilities would hardly be an effective rhetorical response to skeptical objections. Say that such a regress is *nested*, since the presuppositions of the successive meta-methods get ever weaker.¹⁰ This does not imply that H_0 entails H_1 , for H_0 is not a presupposition. Then

Proposition 4 The following are methodologically equivalent:

- an infinite, nested regress (M_0, \dots, M_n, \dots) of sequential refuters.
- a single method M that decides H_0 with at most two retractions, starting with acceptance, over the disjunction ($H_1 \vee \dots \vee H_n \vee \dots$) of all the presuppositions of the methods in the regress.

In the special case in which the infinite refuting regress is **H_0 -entailed**, it is equivalent to a single method that really refutes H_0 over the disjunction of presuppositions of constituent methods in the regress. More generally, if M_0 succeeds with n retractions, the refuting regress is equivalent to a single method that succeeds with one more retraction, starting with the same initial conjecture as M_0 .

Formatted: Highlight

Several points should be emphasized. First, the collapsing construction used to prove the preceding results is not a single, infinitary collapsing function $\Phi(M_0(e), \dots, M_n(e), \dots)$ that looks at the outputs of all the methods in the regress at once. It is, rather, a sequence of finitary collapsing functions of increasing arity that are invoked at successive stages of inquiry

$$\begin{aligned} &\Phi_0(M_0()), \\ &\Phi_1(M_0(e_0), M_1((e_0))), \\ &\Phi_2(M_0((e_0, e_1)), M_1((e_0, e_1))), \\ &\cdot \\ &\cdot \\ &\cdot \end{aligned}$$

so the collapsed output at a given stage of inquiry is constructed out of only finitely many of the outputs of the methods in the regress. Hence, the equivalences hold even if the

¹⁰ This does not imply that H_0 entails H_1 , since H_0 is not a presupposition.

Formatted: Right: 0.25"

infinite regress is built up through time in response to specific, skeptical challenges, instead of being given all at once. Second, no method in the regress has a presupposition as weak as the presupposition of the regress itself, so appealing to a regress is a way to weaken presuppositions of inquiry overall *after* a method with given presuppositions has been chosen. Third, although such regresses yield greater reliability, they are feasible only for hypotheses that are decidable with two retractions, which falls far short of dealing with Duhem's problem, which gives rise to problems that are only verifiable in the limit.

Formatted: Highlight

The last point is ironic for Popper's (1968) "falsificationist" philosophy of science. Popper started with the common insight that universal laws are refutable but not verifiable. But his "falsificationist" philosophy was not that naïve. He was aware of Duhem's problem of blame-assignment and of the fact that an isolated hypothesis can be sustained come-what-may by twiddling other auxiliary hypotheses. He held that this "conventionalist stratagem" of preserving a pet hypothesis at the expense of changes elsewhere is a bad idea because it ensures convergence to the wrong answer if the hypothesis is false. Better, he thought, to *stipulate* crisp conditions under which the (non-refuted) hypothesis should be rejected in advance. Of course, the stipulation involves another hypothesis: that the rejection is not in error. But one can also set up falsification conditions for that hypothesis, etc. Carried to its logical conclusion, this recommendation amounts to an infinite refutation regress. I am unaware whether Popper somewhere addressed the question of nesting, but it would be quite natural for someone vaguely concerned with truth-finding to add this requirement. Now the whole point of Popper's philosophy was to find the truth in the face of Duhem's problem. But the preceding result shows that Popper falls far short. Questions involving even concrete auxiliary hypotheses like uniformitarianism's schedules of progress are not even decidable in the limit, but an infinite Popperian regress of nested refuters exists only when the question is decidable with just two retractions.

The irony is worse than that. For Popper, the falsificationist, *could* have addressed Duhem's problem had he been a regressive *verificationist* rather than a regressive falsificationist.

Proposition 5 The following are methodologically equivalent:

- An infinite, directed regress (M_0, \dots, M_n, \dots) of methods that converge and succeed regressively in any of the following senses: verification with certainty, n retractions starting with "yes", decision in the limit, or refutation in the limit.
- A regress (M_0, M) such that M refutes the presupposition H_1 of M_0 in the limit over the disjunction $(H_2 \vee \dots \vee H_n \vee \dots)$ of all the other presuppositions in the regress.

Recall that regresses of limiting methods are irreducible to simpler success criteria. If the regress is H_0 -entailed or co- H_0 -entailed, however, then one obtains the following, cleaner results.

Formatted: Right: 0.25"

Proposition 6 The following are methodologically equivalent:

- An infinite, H_0 -entailed, directed regress (M_0, \dots, M_n, \dots) of methods that converge and succeed regressively in any of the following senses: verification with certainty, n retraction starting with “yes”, decision in the limit, or refutation in the limit.
- A single method M such that M refutes H_0 in the limit over the disjunction ($H_2 \vee \dots \vee H_n \vee \dots$) of all the other presuppositions in the regress.

Proposition 7 The following are methodologically equivalent:

- An infinite, co- H_0 -entailed, directed regress (M_0, \dots, M_n, \dots) of methods that converge and succeed regressively in any of the following senses: verification with certainty, n retraction starting with “yes”, decision in the limit, or refutation in the limit.
- A single method M such that M verifies H_0 in the limit over the disjunction ($H_2 \vee \dots \vee H_n \vee \dots$) of all the other presuppositions in the regress.

To illustrate proposition 6, recall Lyell’s uniformitarian hypothesis. Extending the finite regress discussed earlier without end, one obtains an infinite regress in which M_n says “no” for H_n if the $n + 1$ th schedule of progress is instantiated but non-refuted, and says “yes” otherwise, thereby presupposing $H_n = H_0 \vee A_1 \vee \dots \vee A_{n+1}$. This is, evidently, an H_0 -entailed nested regress of methods that converge and succeed regressively with two retractions starting with “yes” and, hence (by proposition 6), is equivalent to a single limiting refutation procedure M for H_0 that succeeds over the disjunction of the presuppositions, which is in turn equivalent to the disjunction of the two competing paradigms (i.e., uniformitarianism \vee progressionism). Here is how to construct M in this particular case. Method M maintains a queue of the methods added to the regress so far. Each time a new method is added to the regress, it gets added to the end of the queue (the regress is only “potentially” infinite). If the method at the head of the queue says “yes”, it is placed at the end of the queue (ahead of any new methods added at that stage). Each time the method at the head of the queue is shuffled to the back, M says “yes”. Otherwise, M says “no”. Suppose that some proposition H_i is true. Let H_n be the first such. Suppose that $n > 0$, so that H_0 is false. Then H_{n-1} is false. Since H_n is true, M_n converges correctly to “no”. If $k < n$, then H_{k-1} is false and H_k is false, so M_k converges incorrectly to “yes”. If $k > n$, then H_{k-1} is true and H_k is true (by nesting), so M_k converges correctly to “yes”. Hence, M_n is the unique method in the sequence that converges to “no”. So eventually M_n comes to the head of the queue after it has converged to “no” and M converges to “no” at that stage, as required, since H_0 is false. Now suppose that $n = 0$, so that H_0 is true. Then all of the presuppositions are also true, by nesting, and the hypothesis H_0 is true, so all of the methods converge correctly to “yes”. Hence, M says “yes” infinitely often. So M refutes uniformitarianism in the limit, in accordance with proposition 6. Proposition 7 is illustrated by progressionism in the same example, if one exchanges “yes” with “no”.

Formatted ... [1]

Formatted: Right: 0.25"

Observe how the collapsing construction in this example unwinds the rhetorical game of sequentially responding to challenges with methods that entertain more possibilities into a single, ongoing process of inquiry that finds the truth over all the possibilities covered by the constituent methods in the regress. This, I propose, is a new and interesting model of how rhetorical and reliabilist conceptions of science can be reconciled and systematically compared.

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

CONCLUSION

Scientific method may be conceived as a justifying argument or as a procedure aimed at finding a correct answer. Both conceptions raise a natural question about the propriety of infinite empirical regresses, whether of “evidential justification” or of methods checking methods checking methods. Since it is hard to say what evidential justification is *for*, it is hard to bring the notion of infinite regresses of justification under firm theoretical control. The procedural concept of methodological equivalence, on the other hand, allows one to “solve” for the best single-method performance that a given kind of regress is equivalent to. Some motivated conditions on regresses result in nontrivial regresses that achieve sufficient power to address Duhem’s problem.

ACKNOWLEDGEMENT

The author is indebted to C. Glymour for detailed comments on the penultimate draft of this paper.

Bibliography

Berger, W. and L. Labeyrie (1987) *Proceedings of the NATO Advanced Research*

Workshop on Abrupt Climatic Change: Evidence and Implications, Dordrecht: Reidel.

Gold, E. M. (1965) “Limiting Recursion”, *Journal of Symbolic Logic* 30: 27-48.

Gold, E. M. (1967) “Language Identification in the Limit”, *Information and Control* 10: 447-474.

Hinman, P. (1978) *Recursion Theoretic Hierarchies*, New York: Springer.

Jain, S., D. Osherson, J. Royer, and A. Sharma (1999) *Systems that Learn*, second edition, Cambridge: MIT Press.

Kelly, K. (1996) *The Logic of Reliable Inquiry*, New York: Oxford University Press.

Kelly, K. (2000a) “Naturalism Logicized”, in After Popper, Kuhn, and Feyerabend, R. Nola and R. Sankey, eds., Dordrecht: Kluwer pp. 177-210.

Formatted: Right: 0.25"

- Kelly, K. (2000b) "The Logic of Success", *British Journal for the Philosophy of Science*, forthcoming in the millennial issue.
- Kelly, K. (2004) "Uncomputability: The Problem of Induction Internalized," *Theoretical Computer Science* 317: 227-249.
- Kelly, K. and C. Glymour (2004) "Why Probability Does Not Capture the Logic of Scientific Justification", C. Hitchcock, ed., *Contemporary Debates in the Philosophy of Science*, Oxford: Blackwell pp. 94-114.
- Kuhn, T. (1970) *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.
- Laudan, L. (1980) "Why Abandon the Logic of Discovery?", in *Scientific Discovery, Logic, and Rationality*, T. Nickles, ed., Dordrecht: D. Reidel.
- Martin, E. and D. Osherson (1998) *Elements of Scientific Inquiry*, Cambridge: MIT Press.
- Osherson, D., M. Stob, and S. Weinstein (1986) *Systems that Learn*, Cambridge: MIT Press
- Popper, K. (1968), *The Logic of Scientific Discovery*, New York: Harper.
- Putnam, H.(1963) "Degree of Confirmation and Inductive Logic", in *The Philosophy of Rudolph Carnap*, ed. A. Schilpp. LaSalle: Open Court.
- Putnam, H. (1965) "Trial and Error Predicates and a Solution to a Problem of Mostowski", *Journal of Symbolic Logic* 30: 49-57.

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight

Page 14: [1] Formatted tempadmin 6/14/2005 7:03:00 PM
Highlight