PREPRINT

Sic Transitivity: Reply to McGrew and McGrew

John Post and Derek Turner

The Journal of Philosophical Research, *25* (2000), 67-82

Lydia McGrew and Timothy McGrew's vigorous critique[1] is welcome for its effect of forcing greater clarity at a number of key points. We are grateful also for their recognition of how much is at stake – not just foundationalism, but any other "ism" which presupposes that all the relevant forms of inferential justification are transitive. Among the isms that presuppose transitivity – typically in their circularity arguments against opponents – are most varieties of relativism, skepticism, fideism, anti-realism, internalism, and more**:** a zoo. However, in the present paper we join McGrew and McGrew in concentrating on whether foundationalism can overcome the objection that because relevant forms of inferential justification are not transitive, the regress argument for foundationalism fails, together with the received foundationalist circularity arguments against many of foundationalism's opponents.

The regress argument is a deductive argument for, in the first place, the conclusion that at some point inferential justification or reason-giving must end, on pain of either vicious regress or vicious circularity. Indeed, many anti-foundationalists themselves use this much of the argument to reach the same conclusion. What distinguishes foundationalism is the further conclusion that the beliefs with which reason-giving ends can nonetheless be rationally justified, just non-inferentially. In any case, like all deductive arguments, the regress argument can be represented as a reduction to absurdity of conjoined assumptions**:**

(1) There are justified beliefs.
(2) Every justified belief is justified by inferring it from some justified belief or beliefs.
(3) No belief justifies itself.
(4) If a belief *x* justifies a belief *y*, and *y* justifies *z*, then *x* justifies *z*.
(5) There is no infinite sequence of beliefs each of which is justified by inferring it from its predecessor.

Each of these assumptions seems plausible in itself, yet jointly they entail a contradiction. In particular, (1)-(4) entail the contradictory of (5), namely

(6) There is an infinite sequence of beliefs each of which is justified by inferring it from its predecessor.

To escape this absurdity, foundationalists reject (2). Not every justified belief is justified by

[1] Mcgrew and McGrew (2000). All page references are to this paper unless otherwise noted.

inferring it from justified beliefs; some are justified non-inferentially, so that (1) is preserved, contrary to skeptics and relativists, who reject (1). What is crucial, however, is that the prior conclusion required by foundationalists (among others) – namely that reason-giving must end, contrary to (2) – follows only granted (4), which asserts the transitivity of inferential justification. Without the transitivity there is no absurdity.[2] The regress argument for foundationalism thus presupposes that all the relevant forms of inferential justification are transitive.

Note also that the regress argument contains a built-in circularity argument, as it should. For in order to exclude not only vicious regress, which is expressed by (6), but vicious circularity, which is unexpressed (except for the special case mentioned by (3)), there must be some argument to the effect that there cannot be "closed loops" of inferential justification, sequences of inferential justifications in which some proposition is in its own inferential ancestry. A special case is a sequence of inferential justifications in which Jn_1x (n_1 inferentially justifies x), Jn_2n_1, ... Jxn_n. In this sort of sequence, x is justified by something that is justified by ... something that is justified by x; hence x would be justified, since everything in its inferential ancestry (namely, every proposition in this sequence) would be inferentially justified, in line with (2). The built-in circularity argument against such closed loops is that since J is transitive by (4), a closed loop would entail that Jxx – that x inferentially justifies itself – which is intolerable and is duly excluded by (3).

That is, the circularity argument excludes closed loops of inferential justification by reducing them to the case in which a belief x is both the conclusion and a premise of a single, one-step justificatory argument (the case mentioned by (3)). And this case is a paradigm of vicious circularity, whereas something's occurring somewhere or other in its own inferential ancestry is not; a number of philosophers deny that the latter entails vicious circularity, whereas no one denies that a one-step justificatory argument whose conclusion is one of its premises is viciously circular (at least where the other premises do not themselves suffice to justify the conclusion). Indeed, the question of whether a justified belief may occur somewhere in its own inferential ancestry is one of the crucial issues between foundationalists and a number of their leading opponents (usually coherence theorists of some sort); foundationalists reject all such closed loops, a number of others emphatically do not. Hence foundationalists are not entitled to help themselves to the intuition that no justified belief may occur anywhere in its own inferential ancestry. The circularity argument built into the regress argument is meant to advance the debate beyond intuition fights, by reducing all cases in which a belief x occurs in its own ancestry to the case in which x is both a premise and the conclusion of a single, one-step justificatory argument – a paradigm case of vicious circularity about which everyone is likely to

[2]Cf. Post (1996), 243, Post (1980); and Black (1988). Post has realized only recently that Sanford (1975) was the first to note in print that regress arguments require the transitivity of the relation at issue, here the relation of inferential justification (see also Sanford (1984)). Post and Sanford have lately engaged in hilarious negotiations over proper punishment for Post, ranging from giving his next lecture in sack cloth and ashes, to walking to the next APA meeting on his knees, to ever more imaginative acts of contrition. Post's plea that the Sanford (1975) title – "Infinity and Vagueness" – was insufficiently informative has been of no avail.

agree (at least where the other premises do not themselves suffice to justify the conclusion). The trouble is that this built-in circularity argument presupposes that the inferential-justification relations involved in the ancestry are transitive, and there appear to be clear counter-examples to the would-be transitivity.

McGrew and McGrew's foundationalist response to all of this is basically threefold. (i) They define a relation E of positive evidence, which they argue is necessary for inferential justification regardless of the mode of inference involved, "maps directly over the definition of 'evidence for' most widely in use in confirmation theory" (MS p. 3), and is transitive and irreflexive. From all of this it is to follow, by way of a kind of circularity argument, that a belief may never occur anywhere in its own inferential ancestry, so that there can be no closed loops whatever. (ii) In light of this no-closed-loops conclusion, they construct a revised regress argument for foundationalism in order to "express more explicitly the foundationalist challenge" (MS p. 22). (iii) They argue that certain counter-examples Post has raised against the transitivity of various forms of inferential justification do not work against the transitivity of their relation E (we grant that they do not work against E), and indeed do not work in the first place against the transitivity of at least some of the forms inferential justification targeted by Post.

Were McGrew and McGrew's revised regress argument to succeed, together with the circularity argument it contains, the door would be open again to the whole unlovely zoo of philosophies, often relativist, according to which there must always come a point where there no further argumentative recourse is possible. Fortunately, as we will argue, one can show that their relation E is not transitive after all, by means of a counter-example constructed in their own terms, and that this non-transitivity undermines their circularity argument against all closed loops, insofar as their argument requires E to be transitive. The same counter-example also undermines their revised regress argument for foundationalism, which expressly requires E to be transitive; and a related example shows how a proposition can easily occur in its own inferential ancestry without the sort of vicious circularity they argue must follow. Further, if their relation E is necessary for any mode whatever of inferential justification, as they claim (but we doubt), then our argument that E is non-transitive amounts to a further argument that inferential justification is not in general transitive. Finally, their reasons for rejecting certain counter-examples to the transitivity of some forms of inferential justification are less than persuasive.

McGrew and McGrew's argument (or one of their arguments) against all closed loops of inference parallels the circularity argument built into the regress argument. If a proposition *p* were in its own inferential ancestry, they argue, then *p* would bear the positive-evidence relation E to something that bears E to ... something that bears E to *p*; since E is transitive, it would follow that *p* bears E to itself, which would amount to "a paradigm case of vicious circularity" (MS pp. 1, 21).

But is their relation E of positive evidence transitive? According to their definition of E, if E*xy* (*x* bears E to *y*), then

(b) given the other propositions and their credibilities in *y*'s inferential ancestry, *y* would

> have a lesser level of credibility than it currently has if *x* were not present in *y*'s inferential ancestry or were present with a lesser degree of credibility (MS p. 2).

By the same token, if E*xy*, then given the propositions and their credibilities in *y*'s inferential ancestry other than *x*, *y*'s level of credibility is increased by *x*'s presence and degree of credibility in *y*'s ancestry.

As McGrew and McGrew say, "It is condition (b) ... which carries the epistemic 'clout,'" when it comes to arguing for E's transitivity (MS p. 2, note 6). Further, according to their account, the degree of credibility provided by *x* to *y* is related to conditional probability in such a way that if *x* raises *y*'s degree of credibility, on a given background *b*, then P(*y*/*x*) > P(*y*) on *b*. Indeed, they go so far as to say that "a contingent proposition *B* is evidence for a contingent proposition *A* iff, on a given background ... P(*A*/*B*) > P(*A*)" (MS p. 3), thus embracing Carnap's positive relevance criterion. The background, according to condition (b), includes all the propositions in *y*'s inferential ancestry, other than *x*, together with their credibilities. Hence if E*xy*, then P(*y*/*x*) > P(*y*) on the relevant background *b*, namely on all the propositions in *y*'s inferential ancestry, other than *x*, together with their credibilities.

Now assume that E*xy* and E*yz*. If E were transitive, it would follow that E*xz*. But from E*xz* it would follow, by their account, that P(*z*/*x*) > P(*z*) on the relevant background *b*, namely on all the propositions in *z*'s inferential ancestry, other than *x*, together with their credibilities. Accordingly, if there are cases in which E*xy* and E*yz* but *not* P(*z*/*x*) > P(*z*), on the relevant background, then there are cases in which E*xy* and E*yz* but not E*xz*, from which it follows that E is not transitive. There are such cases. Consider

g. Gcd
f. Fcd
h. (*x*)(*y*)(F*xy* 6 G*xy*)
k. (*y*)(Fa*y* 6 Ga*y*)[3]

Suppose that for *S* at *t*, Egh (g is evidence for h) on background f (consistent with McGrew and McGrew's endorsement of confirmation theory), together with whatever is in the conjunction *b* of the beliefs in the inferential ancestry of g and f (possibly none, since g and f might be basic for *S* at *t*, consistent with their foundationalism). Thus the beliefs in the inferential ancestry of h for *S* at *t* are g, f and *b*. Also, (i) since by hypothesis Egh on f and *b*, and (ii) since Egh implies P(h/g) > P(h) according to their account, it follows that (iii) P(h/g) > P(h) on f and *b*, or equivalently that P(h/g&f&*b*) > P(h/f&*b*), where f and *b* are all the propositions other than g in h's inferential ancestry, in compliance with their condition (b). Note that g&f&*b* is consistent (as will be required below) because by hypothesis, Egh on f and *b*, which implies that

---

[3]The example is from Merrill (1979), 110-111, who uses it to argue against Hempel's consequence condition ("If an observation report confirms every one of a class K of sentences, then it also confirms any sentence which is a logical consequence of K").

P(h/g&f&*b*) > P(h/f&*b*), which would be undefined if g&f&*b* were inconsistent.

Suppose further that Ehk for *S* at *t* – that *S* at *t* both recognizes that h entails k and believes k on the basis of this entailment. The beliefs in the inferential ancestry of k for *S* at *t* are g, f, *b* and h, so that h is evidence for k on g, f and *b*. Also, P(k/h) > P(k) on g, f and *b*, or P(k/h&g&f&*b*) > P(k/g&f&*b*), where g, f and *b* are all the propositions other than h in k's inferential ancestry, again in compliance with condition (b). And h&g&f&*b* is consistent, because by hypothesis, Ehk on g, f and *b*, which implies that P(k/h&g&f&*b*) > P(k/g&f&*b*), which would be undefined if g&f&*b* were inconsistent.

So we have both Egh and Ehk, for *S* at *t*. If E were transitive, it would follow that Egk. And from Egk it would follow that P(k/g) > P(k) on the relevant background, according to McGrew and McGrew's condition (b) and their view that if E*xy*, then P(*y*/*x*) > P(*y*) on the relevant background. The relevant background here consists of all the propositions other than g in k's inferential ancestry for *S* at *t*, namely h, f and *b*; as they put it, "one proposition is only evidence for another in any given case when background knowledge within the tree is taken into account" (MS p. 5). Thus if E were transitive, it would follow from Egh and Ehk that P(k/g&h&f&*b*) > P(k/h&f&*b*). By the same token, if Egh and Ehk but *not* P(k/g&h&f&*b*) > P(k/h&f&*b*), then E is not transitive.

This is indeed what happens. For h entails k, from which it follows that h&f&*b* entails k. By a law of the probability calculus, if *x* is consistent and entails y, then P(*y*/*x*)=1. By substitution of 'k' for '*y*' and 'h&f&*b*' for '*x*' in this law, P(k/h&f&*b*)=1. Since 1 is the maximum probability a proposition can have, P(k/g&h&f&*b*) cannot be greater than P(k/h&f&*b*), contrary to the inequality 'P(k/g&h&f&*b*) > P(k/h&f&*b*)'. Hence Egh and Ehk but not P(k/g&h&f&*b*) > P(k/h&f&*b*), from which it follows that Egh and Ehk but not Egk. The relation E of positive evidence, as defined by McGrew and McGrew, is not transitive.

Note that *b*, f, g, h and k are all in the inferential ancestry of k for *S* at the same time *t*. This means that they are all in the same "synchronic evidence tree," as McGrew and McGrew call it. Note also that the tree involves no loops of mutual support; their claim (on MS p. 19) that cases of mutual support should not be represented by a synchronic evidence tree has no bite here. This spells trouble for the crucial premise (5) of their revised regress argument for foundationalism**:**

> If a belief *x* is positive evidence for a belief *y*, and, in the same synchronic evidence tree, *y* is positive evidence for *z*, then *x* is positive evidence for *z* (MS p. 23).

Our counter-example to the transitivity of E undermines this crucial premise without resorting to contested examples of mutual support.

If our argument so far is sound, there must be some flaw in McGrew and McGrew's argument for the transitivity of E. Without attempting a complete diagnosis, we suggest there is a non-sequitur in their move (on MS p. 4) from**:**

> if E*xy*, [then] a lowering of the credibility of *x* on the background would result in a lowering of the credibility of *y* on the background. And ... if E*yz*, and the credibility of y is lowered, [then] the credibility of z must be lowered

to**:**

> So, if the credibility of *x* is lowered and E*xy* and E*yz*, the credibility of *z* must be lowered.

The latter statement is incomplete, since it abbreviates a statement that mentions the background**:**

> So, if the credibility of *x* is lowered and E*xy* and E*yz*, the credibility of *z* must be lowered *on the background*.

What background? According to condition (b), it is the background on which E*xz*, namely all the propositions other than *x* in the ancestry of *z*, together with their credibilities. Only if one stops to consider *just what this background might include in particular cases* is one likely to think of situations in which part of the background properly between *z* and *x* entails *z*. In such situations the probability of *z* on the background is 1, so that the credibility of *z* on the background is maximal by their account and thus could not be increased by adding *x* to this background, or by increasing *x*'s credibility, whence it follows that *z* would not have a lesser level of credibility than it currently has if *x* were not present in *z*'s inferential ancestry or were present with a lesser degree of credibility on this background, contrary to condition (b). In such situations, it is not the case that E*xz* even if E*xy* and E*yz*.

One might hope to escape the foregoing argument against transitivity of E by decoupling degree of credibility from the inequality 'P(*y*/*x*) > P(*y*)'. But this escape is not available to McGrew and McGrew, who are heavily committed to Carnap's positive relevance criterion that contingent *x* is evidence for contingent *y* on *b* only if P(*y*/*x*) > P(*y*) on *b*. Indeed, their argument that nothing can bear E to itself depends on some such asymmetric measure of degree of credibility (MS p. 21), and their use of the probability calculus commits them to *y*'s credibility maxing out when consistent *x* entails y.

By contrast, Wesley Salmon is able to argue, consistent with Kolmogorov's theorem on total probability, that we do "have a transitivity relation" in case P(*y*/*x*) is high and *y* entails *z* (because in that case it follows that P(*z*/*x*) is also high).[4] Salmon is able to argue this because, unlike McGrew and McGrew, he is talking about a relation of inductive support free of the requirement that *x* is evidence for *z* only if P(*z*/*x*) > P(*z*) on the background they count as relevant; he requires only that P(*z*/*x*) $ P(*y*/*z*) when *y* entails *z*. Since there are other cases involving this relation of inductive support in which we do not have transitivity, as he argues, he concludes that "The probability calculus shows unequivocally that the relation of inductive

[4] Salmon (1965), 168.

support must not be construed as a transitive relation." [5]

Suppose, however, that McGrew and McGrew tried the following tactic. They might try defining a relation E2 of positive evidence decoupled from the inequality 'P(*y*/*x*) > P(*y*)', which relation nonetheless would do the work that their relation E is supposed to do in precluding closed inferential loops, this time by reducing them to the case in which something bears E2 to itself (which they would argue is intolerable). The trouble is that any such tactic is vulnerable to the following general difficulty, which the foregoing argument against the transitivity of E exploits.

Let us say that a relation *R* implies a relation *Q* if and only if, for any *x* and y, if *Rxy* then *Qxy* . Then it can be proved that

> TH. If *R* implies *Q*, and there are *x, y* and *z* such that *Rxy* and *Ryz* but not *Qxz,* then *R* is not transitive.[6]

We relied on an instance of TH when we inferred from 'For any *x* and *y*, if E*xy*, then P(*y*/*x*) > P(*y*),' and 'Egh & Ehk but not P(k/g) > P(k)' on the background, to E's not being transitive. What is general about TH is that it entails that given *any* would-be epistemic relation *R* – whether of evidence, justification, support, dependence, confirmation, basing, you name it – *R* is transitive only if for *none* of *R*'s implications *Q* is it the case that there are *x, y, z* such that *Rxy* and *Ryz* but not *Qxz*. And it is by no means easy to show than none of *R*'s implications suffers this fate, as the present troubles with relation E illustrate.

Furthermore, suppose someone does manage to show, for some especially conspicuous relation *Q*1 implied by *R*, that there are no *x, y, z* such that *Rxy* and *Ryz* but not *Q*1*xz*, thus successfully defending *R*'s transitivity against attack from this particular direction. Even so, there may well be some other, unforeseen relation *Q*2 implied by *R* such that *Rxy* and *Ryz* but not *Q*2*xz*, so that *R* is not transitive after all. Thus even if McGrew and McGrew's relation E could somehow be defended by decoupling credibility from the inequality 'P(*y*/*x*) > P(*y*)', E could still fail to be transitive because it implies some other relation *Q*2 such that *Rxy* and *Ryz* but not *Q*2*xz* for some *x, y, z*, We suspect that this is indeed the case, but need not pursue the matter here.

In any event, assume for the sake of argument that McGrew and McGrew's relation E of positive evidence, as they define it, is "a necessary condition for inferential justification, regardless of the particular inference form used.... [N]o proposition can be inferentially justified if no proposition bears this relation to it" (MS pp. 1-2). This entails that any inferential relation *J* implies the relation E, in the sense defined above**:** for any *x* and y, *Jxy* 6 E*xy*. It follows by TH that *J* is not transitive if there are cases in which *Jxy* and *Jyz* but not E*xz*.

---

[5]Salmon (1965), 167.

[6]Outline proof**:** Assume that (*x*)(*y*)(R*xy* 6 *Qxy*) and *R*ab & *R*bc but not *Q*ac. It follows by universal instantiation and tautological inference that *R*ab & *R*bc but not *R*ac, which entails that *R* is not transitive.

There are such cases. For example, suppose that g above, on background f combined with other background *b*, is (just) enough to justify h for *S* at *t*, so that *J*gh on f and *b*. Suppose further that h, by virtue of entailing k, justifies k for *S* at *t* on g, f and *b*, so that *J*hk on g, f and *b* for *S* at *t*. Then we have *J*gh on f and *b*, and *J*hk on g, f and *b*, for *S* at *t*. But once again it is not the case that Egk on the relevant background, namely on the propositions in k's ancestry other than g, which are f, *b* and h. For h entails k, so that h&f&*b* does too and P(k/h&f&*b*)=1, from which it follows that the inequality implied by 'Egk' according to McGrew and McGrew's account – namely 'P(k/g&h&f&*b*) > P(k/h&f&*b*)' – is not satisfied. Consequently, if their relation E of positive evidence is a necessary condition for inferential justification regardless of the particular inference form used, then inferential justification is not in general transitive.

The foregoing argument also shows that their relation "sufficient to justify_1" (defined on MS p. 7) is not transitive either, as can easily be checked. Thus even when one takes into account the potential ambiguity in the term 'given', about which they suggest Post is confused, nonetheless "sufficient to justify_1" – if it implies E – is no more transitive than "sufficient to justify_2." In addition, neither of these two relations they define of inferential justification is the one Post is talking about in the example they target, so that much of their discussion here is wide of the mark.

So far we have focused on E's non-transitivity and its consequences for McGrew and McGrew's view. A major consequence is that their circularity argument fails against a belief's occurring in its own inferential ancestry, insofar as the argument requires E to be transitive. But they advance a further argument, which may not require E to be transitive. The argument is that if *P* were present in its own inferential ancestry,

> then it would be epistemically legitimate for P to be treated as "positive evidence for" itself – i.e., for a reasoner to believe P as the apex of a synchronic tree with a greater degree of credence because of the presence of P as evidence farther down in its own ancestry. Yet obviously, it is *not* epistemically legitimate for a reasoner to raise the credence he gives to P on the grounds that P. If granting P a greater degree of credence than one would otherwise give to it on the basis of P is not vicious, then nothing is (MS p. 21).

The crux of the argument appears to be that if *P* were in its own ancestry, then *P* would automatically raise its own credibility, which would be a viciously circular state of affairs. We agree of course that there are kinds of inferential ancestries in which *P*'s being in its own ancestry would result in *P*'s raising its own credibility (whether or not this would in every case be a viciously circular state of affairs). But there is at least one important kind of case in which *P*'s being in its own ancestry does not result in *P*'s raising its own credibility, on McGrew and McGrew's understanding of the relations among credibility, probability, and their evidence relation E. Suppose that

A. *x* entails *y* on background *b* (that is, *x* conjoined with *b*, or elements of *b*, entails *y*).
B. there is substantial evidence e_x for *x* that is independent of *y*.

C. there is substantial evidence e_y for *y* that is independent of *x*.
D. *y* confirms *x* on *b* and e_y.

For example, consider

f. Fcd
h. (*x*)(*y*)(F*xy* 6 G*xy*)
g. Gcd

In line with B and C, assume that there is both substantial evidence e_h for h independent of g, and substantial evidence e_g for g independent of h. Because h entails g on background f and evidence e_h, in line with A, h is evidence for g on f and e_h. Further, g confirms h on f and e_g, in line with D. Thus g is in its own inferential ancestry. Of course McGrew and McGrew would reject both the particular example and the general pattern defined by A-D, on the ground that just by being in its own ancestry, g would automatically raise its own credibility. But....

Does g raise its own credibility? No, not on McGrew and McGrew's understanding of the relations among credibility, probability, and their evidence relation E. For h entails g on f and e_h, so that $P(g/h\&f\&e_h)=1$. Hence the credibility of g given h maxes out on the background, because g is entailed by h on the background. This means that nothing in g's ancestry beyond h, including g itself, can increase g's credibility. Consequently, even though g occurs in its own ancestry, g occurs in its own ancestry beyond the step at which h is evidence for g by virtue of entailing g on the background; since g's credibility maxes out at this step, neither g nor anything else beyond this step can increase g's credibility. Hence it does not follow from a proposition's being in its own ancestry that it would raise its own credibility, on their understanding of the relations among credibility, probability, and their evidence relation E. This may follow for some kinds of ancestry, depending on what forms of inference are involved in the ancestry and in what order, but hardly all.

The general pattern defined by A-D occurs frequently in the sciences and elsewhere. Among its many instances is the case of mutual support between Newton's theory and the discovery of Neptune. McGrew and McGrew object to Post's use of this case to counter-example the transitivity of a certain kind of inferential justification. As Post sees it, Newton's theory T_N, on background *b* that includes various auxiliaries and substantial evidence e_T for T_N, entails

ON. Planet *X* (later named Neptune) appears at orbital position *p*N at clock time *t*N,

where e_T is independent of ON; hence A and B are satisfied. Further, there is substantial observational evidence e_{ON} for ON independent of T_N, and ON confirms T_N on *b* and e_{ON}, in line with C and D. Post claims that T_N inferentially justifies ON on *b*, that ON is part of what inferentially justifies T_N, and that ON does not inferentially justify itself even in part, contrary to what transitivity in this sort of case would require.

McGrew and McGrew reply by rejecting the idea that in this sort of case (or in any other instance of A-D),

> The correct representation of mutual support involves giving a proposition [like ON] evidential weight in its own [synchronic] ancestry (MS p. 20).

They argue at length that foundationalists can adequately depict reasoning from mutual support by representing it not as synchronic, as does the pattern A-D, but as a diachronic process of thinking. Unfortunately, this would show at most that the foundationalist is not committed to the assumption McGrew and McGrew reject, not that it must be rejected. Nonetheless they assert that "as we have pointed out, if 'inferential ancestry' refers to synchronic ancestry, [the assumption] is false" (MS p. 20). What they seem to have in mind is their earlier argument against any proposition's ever appearing in its own evidential ancestry. This argument, as we have noted, is a circularity argument which requires their positive evidence relation E to be transitive. Since E is not transitive, the argument fails. There is a later argument, also noted, which may not require E to be transitive. Rather, the argument is that if a proposition *P* were in its own ancestry, it would automatically raise its own credibility. This too we have shown to fail – on their understanding of the relations among credibility, probability, and E – in any case in which *P*'s credibility maxes out on the relevant background by virtue of being entailed by it.

Nor is it clear in the first place that foundationalists should always go diachronic, as McGrew and McGrew evidently require, when depicting reasoning that involves mutual support. Many foundationalists do not, often in order to accommodate the objective concept of evidence they find in science and elsewhere. According to Laura Snyder's characterization of the concept,

> whether e is evidence for h does not depend upon anyone's beliefs or knowledge about e, h, or anything else. Hence if some e is evidence for h, it is so regardless of what any person knows or believes.[7]

As Snyder remarks, "Any plausible theory of evidence (that is, one which captures how the concept is used in science) must account for this impersonal use of evidence in science." Carnap and Hempel, who are foundationalists, are among those who advance theories of evidence that conform to this general condition. Yet the condition is incompatible with the diachronic, historical approach McGrew and McGrew would have us apply to any case involving mutual support. The reason it is incompatible is that "on any objective view, the *time* at which e is known relative to the invention of h cannot be relevant to whether e is evidence for h." [8] Hence adopting the diachronic approach in every case would exclude the objective concept of evidence at work in science and elsewhere. This seems an exorbitant price to pay, if the motive is to preserve transitivity for all forms of positive evidential support. The point is not that there are no contexts of inferential reasoning involving mutual support to which a synchronic approach

---

[7] Snyder (1998), 469.

[8] Snyder (1998), 470.

does not apply. Rather, sweeping rejection of a synchronic approach regardless of the context requires far more argument than is on offer.

Finally, we turn to the matter of whether inference to the best explanation is transitive. Post readily concedes that the argument in Post (1996) against the would-be transitivity of inference to the best explanation needed work, though not along the lines urged by McGrew and McGrew. The main argument there was essentially that inference to the best explanation is a form of inferential justification *J* such that '*Jxy*' implies '*y* is the best explanation of *x*,' and that since the relation "best explanation of" is not transitive, neither is the relation *J* that implies it. This was a mistake, as Andrew Cling subsequently pointed out in correspondence; there are counter-examples to 'if relation *Q* is not transitive, neither is any relation *R* that implies *Q*.' It was in response to Cling that Post derived TH above**:** if *R* implies *Q*, and there are *x*, *y* and *z* such that *Rxy* and *Ryz* but not *Qxz*, then *R* is not transitive. For instance, if there are *x*, *y* and *z* such that *x* inferentially justifies *y* by virtue of *y*'s being the best explanation of *x*, and *y* inferentially justifies *z* by virtue of *z*'s being the best explanation of *y*, but *z* is not the best explanation of *x*, then inference to the best explanation is not transitive. The argument in Post (1996) then goes through when revised in light of TH.

To see why, consider the following more explicit version of the argument. The point of inference to *the* best explanation is not that we are to infer only *one* explanation of some phenomena, but that we are to infer only one from among the *competing* explanations. Typically the same affair can be explained in several different but compatible ways. When we notice that the water we put on the stove is still not boiling, we may infer that we forgot to turn on the burner, or we may infer that the mean kinetic energy of the water molecules is too low. The two explanations are compatible; neither can be said, in the relevant sense, to be a better explanation than the other of why the water is not boiling. By contrast, an explanation in terms of the burner's not working does compete with the explanation that we forgot to turn it on. Of course there is no way to tell which is the better explanation given only that the water is not boiling. So we enlarge our data set – we check to see whether the burner is turned on – in order to narrow the field of plausible explanations and decide whether to undertake repairs or lament our absent-mindedness. In general, "Given our data and our background beliefs, we infer what would, if true, provide the best of the competing explanations we can generate of those data (so long as the best is good enough for us to make any inference at all)." [9]

Now suppose that *x* inferentially justifies *y* by virtue of *y*'s being the best explanation of *x*, a better explanatory story than any competitor. Let C_X be the set of the competitors we can generate. C_X might contain the explanation that we forgot to turn the burner on, together with any explanation we can generate that competes with it, including the burner's not working. Or, if we are interested in a different type of explanation altogether – say at a different "level of abstraction," in McGrew and McGrew's phrase – then C_X might instead contain the explanation that the mean kinetic energy of the molecules is too low, together with any of its competitors we

[9] Lipton (1991), 58.

can generate (compare the two levels of abstraction in the balloon case McGrew and McGrew cite from Salmon). In any event, the relation at issue is a specific form of inferential justification such that '*x* inferentially justifies *y*' implies that *y* is the best explanation of *x* relative to some already specified set C_X of competitors we can generate. In general, talk of the best explanation of *x*, hence of inference to it, makes sense, and is unambiguous, *only relative to some already specified set* C_X.

Next, suppose that *y* in turn inferentially justifies *z* in the same sense**:** *y* inferentially justifies *z* by virtue of *z*'s being the best explanation of *y*, where the competing explanations of *y* we can generate are in C_Y. Now if this specific relation of inferential justification were transitive, it would follow that *x* inferentially justifies *z* by virtue of *z*'s bing the best explanation of *x*, hence that *z* is the best explanation of *x*. But *z* is not the best explanation of *x*. For by hypothesis it is *y* that is the best explanation of *x*, which is to say, as seen, that we have already specified some set C_X of the competing explanations we can generate of *x*, and that *y* is better than any other explanation in C_X. If *z* were to be the best explanation of *x* in the same sense, *z* would have to be in C_X and be better than any other explanation in C_X. Since by hypothesis it is *y* that is better than any other explanation in C_X, *z* cannot be the best explanation of *x*. Transitivity fails for inference to the best explanation.

True, if we were given only that (i) *y* is the best explanation of *x* relative to C_X (say relative to competitors at the mechanical level of a pressure gradient in the air molecules, in the balloon example), and (ii) *z* is the best explanation of *y* relative to C_Y, then *z* would not have to be among the explanations in C_X, hence would not have to compete with *y* as an explanation of *x* (instead, *z* could be an explanation at the higher level of Einstein's equivalence principle, in the balloon example). Thus it might look as though *z* could be the best explanation of *x* after all, consistent with transitivity, which is what McGrew and McGrew argue.

However, we must be careful not equivocate on 'the best explanation of *x*', which is fatally easy to do here. In *precisely what sense* is *z* to be the best explanation of *x*? That is, *precisely which set* is the one relative to which *z* is to be the best explanation of *x*? If it is C_X – the one we started with by hypothesis when we assumed that *y* is the best explanation of *x* relative to C_X – then *z* is in C_X along with *y*, if *z* is to be the best explanation of *x*, so that *z* does have to compete with *y* as an explanation of *x*, hence cannot be the best explanation of *x* because by hypothesis *y* is.

On the other hand, suppose that the set relative to which *z* is to be the best explanation of *x* is not C_X but some other set C_X*. Then we would be using the expression 'the best explanation of *x*' in two different senses, hence equivocally. If such equivocation were allowed, one could reject any counterexample whatever to any transitivity claim for any relation, which is absurd. For instance, consider the relation "*x* is the best professor *y* ever had," which is non-transitive; it does not follow from '*x* is the best prof *y* ever had, and *y* is the best prof *z* ever had' that *x* is the best prof *z* ever had. Suppose nonetheless that someone insists that "best prof of" is transitive. The natural way to argue against this alleged transitivity is by presenting a counterexample of the form

Smith is the best prof Jones ever had, Jones is the best prof Brown ever had, but Smith is *not* the best prof Brown ever had.

Now what if this sort of counterexample is rejected on the ground that while Smith is the best prof Jones ever had in the classroom, and Jones is the best prof Brown ever had in the classroom, still Smith is the best prof Brown ever had in a correspondence course? This is the sort of equivocation involved when McGrew and McGrew try to deflect Post's counterexample to the relation "the best explanation of" (MS p. 17). We conclude that on pain of such equivocation, the set relative to which *z* would have to be the best explanation of *x* is C_X , not some other set C_X*.

McGrew and McGrew also object to Post's treatment of a relation of inference to the best explanation which makes use of a quite specific notion of best explanation**:** *x* inferentially justifies *y* by virtue of *y*'s being the best interlevel explanation of *x* which satisfies (among other things) the conditions (1)-(7) McGrew and McGrew quote from Post (2000). In this particular form of inferential justification, '*x* inferentially justifies *y*' implies, by (1)-(7), that the interlevel explanation is given by an interlevel theory that connects the theories *T1* and *T2* at the levels at which *x* and *y* occur, respectively, where *T1* and *T2* are heuristically, confirmationally and explanatorily dependent on each other to a certain threshold degree *N* . Then it is easy to show that there are cases in which *x* inferentially justifies *y* and *y* inferentially justifies *z*, but there is no leap-frog interlevel theory connecting the two theories at the levels at which *x* and *z* occur, where these two theories are heuristically, confirmationally and explanatorily dependent on each other to degree *N* .[10] It follows by TH that this particular form of inferential justification, which is a quite specific kind of inference to the best explanation, is not transitive.

True, it might be "quite possible to define a far more minimal sense of 'explains' than that characterized [with the help of] (1)-(7), which sense would be at least *prima facie* sufficient to sustain an epistemic relation like E" (MS p. 12). But aside from the fact that E proves non-transitive after all, as we have shown, the question of whether there is some such minimal sense of 'explains' would be beside the point, which is that a certain quite specific but important form of inference to the best explanation, like a number of other specific forms of inferential justification, is not transitive. Sic transitivity.

References

Black, O. (1988). "Infinite Regresses of Justification," *International Philosophical Quarterly, 28*, 421-437.

Hempel, Carl G. (1966). *Philosophy of Natural Science* (Englewood Cliffs: Prentice-Hall).

Lipton, Peter (1991). *Inference to the Best Explanation* (London: Routledge).

[10]Post (2000), Section 3.

McGrew, Lydia, and McGrew, Timothy (2000). "Foundationalism, Transitivity and Confirmation," *The Journal of Philosophical Research*, *26*.

Merrill, G. H. (1979). "Confirmation and Prediction," *Philosophy of Science, 46*, 98-117.

Post, John F. (1980). "Infinite Regresses of Justification and of Explanation," *Philosophical Studies, 38*, 31-52.

Post, John F. (1996). Chapter on "Epistemology," *Discourse on Method and Meditations in First Philosophy: René Descartes*, ed. David Weissman (New Haven: Yale University Press, 1996), 236-271.

Post, John F. (2000). "Is Supervenience Asymmetric?" for L. C. Pereira and M. Wrigley, eds., *Festshcrift in Honor of Oswaldo Chateaubriand* (Manuscrito). This is the paper McGrew and McGrew cite as Post (forthcoming).

Salmon, Wesley C. (1965). "Consistency, Transitivity, and Inductive Support," *Ratio*, *7*, 164-169.

Sanford, David H. (1975). "Infinity and Vagueness," *Philosophical Review*, *84*, 520-535.

Sanford, David H. (1984). "Infinite Regress Arguments," in J. H. Fetzer, ed., *Principles of Philosophical Reasoning* (Totowa, NJ: Rowman & Allenheld), 94-117.

Snyder, Laura J. (1998). "Is Evidence Historical?" in M. Curd and J. A. Cover, eds., *Philosophy of Science: The Central Issues* (New York: W. W. Norton), 460-480.