Evidence of Evidence in Epistemic Logic

[draft of 24 August 2017; final version to appear in Mattias Skipper Rasmussen and Asbjørn Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford: Oxford University Press]

Timothy Williamson

The slogan 'Evidence of evidence is evidence' may sound plausible, but what it means is far from clear. It has often been applied to connect evidence in the current situation to evidence in another situation. The relevant link between situations may be diachronic (White 2006: 538): is present evidence of past or future evidence of something present evidence of that thing? Alternatively, the link may be interpersonal (Feldman 2007: 208): is evidence for me of evidence for you of something evidence for me of that thing? Such inter-perspectival links have been discussed because they can destabilize inter-perspectival disagreements. In their own right they have become the topic of a lively recent debate (Fitelson 2012, Feldman 2014, Roche 2014, Tal and Comesaña 2014).

This chapter concerns putative *intra*-perspectival evidential links. Roughly, is present evidence for me of present evidence for me of something present evidence for me of that thing? Unless such a connection holds between a perspective and itself, it is unlikely to hold generally between distinct perspectives. Formally, the single-perspective case is also much simpler to study. Moreover, it concerns issues about the relation between first-order and higher-order evidence, the topic of this volume.

The formulations in this chapter have not been tailored for optimal fit with previous discussions. Rather, they are selected because they make an appropriate starting-point, simple, significant, and not too *ad hoc*. In particular, I will not discuss existential generalizations to the effect that a given hypothesis has *some* evidential support, in the sense that some *part* of the evidence supports it. Such principles are usually too weak to be of interest, since from unequivocally negative evidence one can typically carve out a gerrymandered fragment that in isolation points the opposite way—for instance, by selecting a biased sample of data points. Instead, the focus will be on the *total* evidence.

For the sake of rigour and clarity, evidence will be understood in probabilistic terms. On some probabilistic readings of the slogan 'Evidence of evidence is evidence', it can be straightforwardly refuted by standard calculations of probabilities for playing cards, dice, and so on. For example, it might be interpreted as saying that if *p* is evidence for *q*, and *q* is evidence for *r*, then *p* is evidence for *r*. Counterexamples to such transitivity principles are easy to construct, whether '*p* is evidence for *q*' is understood as '*p* raises the probability of *q*' or as 'the conditional probability of *q* on *p*' is high (see also Fitelson 2012). This chapter is not concerned with principles refuted by standard first-order probabilistic calculations. The principles it discusses all involve readings of the phrase 'evidence of evidence' in terms of second-order evidence, evidence for propositions *about* evidence: more specifically, second-order probabilities, probabilities of propositions *about* probabilities.

Formal models will be used throughout, within the framework of epistemic logic, since it provides a natural way of integrating first-level epistemic conditions (such as evidence of a coming storm) and second-level epistemic conditions (such as evidence of evidence of a coming storm). An integrated framework is needed to give a fair chance to the idea that evidence of evidence is evidence. We will be asking questions like this: if the probability on the evidence that the probability on the evidence of a hypothesis H is at least 90% is itself at least 90%, under what conditions does it follow that the probability on the evidence of H is indeed at least 90%, or at least more than 0%?? Such principles may remind one of synchronic analogues of more familiar probabilistic reflection principles, and turn out to be sensitive to similar structural features of the underlying epistemic relations (compare Weisberg 2007 and Briggs 2009).

Bridge principles between first-level and higher-level epistemic conditions often turn out to imply versions of highly controversial principles in standard epistemic logic, most notably the contentious principle of so-called *positive introspection*, that if one knows something, one knows that one knows it, and the more obviously implausible principle of *negative introspection*, that if one doesn't know something, one knows that one doesn't know it (Williamson 2000: 228-237; 2014). To anticipate, various natural formalizations of the intra-perspectival principle that evidence of evidence is evidence also turn out to have such connections, although more complicated ones than usual. Since the overall argument of this chapter is against those principles, it is dialectically fair to use a formal framework that presents them with no unnecessary obstacles.

1. *The formal framework*

For clarity, the formal framework will first be explained, although some readers will already be familiar with it. The underlying models come from modal logic, as adapted to single-agent epistemic logic (Hintikka 1962), to which we can add the required probabilistic structure (Williamson 2000).

For present purposes, a *frame* is an ordered pair <W, R>, where W is a nonempty set and R is a dyadic relation over W, a set of ordered pairs of members of W. Informally, we conceive the members of W as *worlds*, or as relevantly but non-maximally specific states of affairs, mutually exclusive and jointly exhaustive. We model (coarse-grained) propositions as subsets of W; thus the subset relation corresponds to entailment, set-theoretic intersection to conjunction, union to disjunction, complementation in W to negation, and so on. If $w \in X$, the proposition X is true in the world *w*; otherwise, X is false in *w*. We use the relation R to model *evidence* for a given agent (expressed by 'one') at a given time (expressed by the present tense). More precisely, if *w* and *x* are worlds ($w, x \in W$) then R*wx* ($\langle w, x \rangle \in R$) if and only if it is consistent with one's total evidence in *w* that one is in *x*. We define the proposition R(*w*) as {*x*: R*wx*}, which is the strongest proposition to follow from one's evidence in *w*; in effect, R(*w*) is one's total evidence in *w*. Since our interest concern here is in what follows from one's evidence, which is automatically closed under multi-premise entailment, concerns about the logical omniscience imposed by such models are less pertinent than elsewhere in epistemic logic.

At the very least, the total evidence should be consistent, otherwise it both entails everything and excludes everything. In a probabilistic setting, we want to conditionalize on the evidence, which makes no obvious sense when the evidence is the empty set of worlds. Thus we require that $R(w) \neq \{\}$; in other words, the relation R is serial in the sense that each world has it to some world. The epistemic interpretation motivates the formal development but plays no further role in it. Formally, we will generalize over every (finite) nonempty set W of entities of any kind and every serial relation R over them.

We need to add probabilities to the frames. We assume W to be finite in order to avoid the complications inherent in infinite probability distributions. There is enough complexity and variety in finite probability distributions for most epistemological modelling purposes. A *probabilistic frame* is an ordered triple <W, R, Pr> where <W, R> is a frame, W is finite, R is serial, and Pr is a probability distribution over W. Thus Pr maps each subset of W to a real number between 0 and 1, where Pr(W) = 1 and $Pr(X \cup Y) = Pr(X) + Pr(Y)$ whenever X and Y are disjoint. We impose one further constraint on Pr: it is *regular*, in the sense that Pr(X) = 0 only if X = {} (the converse follows from the other axioms). The reason for that constraint will emerge shortly. Informally, we regard Pr as the prior probability distribution. For present purposes, it need not be *absolutely* prior; it may embody one's previously acquired background information.

Posterior probabilities in a world *w* are defined by conditioning Pr on one's total evidence in *w*: thus the posterior probability of X in *w*, the probability $Pr_w(X)$ of X on the evidence in *w*, is the prior conditional probability of X on R(*w*). These conditional probabilities are themselves defined in the usual way as ratios of unconditional probabilities, giving this equation:

EVPROB $\quad Pr_w(X) = Pr(X \mid R(w)) = Pr(X \cap R(w))/Pr(R(w))$

Of course, the ratio in EVPROB is well-defined only if Pr(R(*w*)) > 0. Since R is serial, R(*w*) is nonempty, but we still need regularity to conclude, from that, that Pr(R(*w*)) > 0. That is why the constraint was imposed. Informally, the probability of a proposition X on one's evidence in *w* is the weighted proportion of worlds consistent with one's evidence in *w* where X is true.

A stronger constraint on Pr than regularity is *uniformity*, which says that for any two worlds *w* and *x*, $Pr(\{w\}) = Pr(\{x\})$: all worlds have equal weight. In a finite probability space, uniformity entails regularity; uniformity equates Pr(X) with the unweighted proportion of members of X that are members of W. We do not impose uniformity, despite the simplicity it brings. For the members of W may represent less than maximally specific possibilities, which may themselves vary in their level of specificity: in that case, uniformity in the model would require non-uniformity at a level of resolution finer than that represented in the model. Permitting non-uniformity makes for more robust results.

Mathematically, this probabilistic framework is just like the one in Williamson 2000 and 2014. However, there R(*w*) was informally explained as what the agent *knows* in *w*, rather than as the evidence in *w*. This chapter does not impose the equation E=K of total evidence with total knowledge. The reason is not any loss of confidence in E=K, but simply a preference for addressing a wider range of views of evidence. Obviously, E=K remains compatible with the present framework. A formal correlate of this difference is that since knowledge entails truth, if R(w) is what the agent knows then $w \in R(w)$, so R is reflexive. All reflexive relations are serial but not *vice versa*. Even without E=K, the assumption that evidence consists of true propositions is attractive (Williamson 2000: 201-2). Nevertheless, for the sake of generality, we do not impose it. Thus the framework allows one's total evidence to be what one reasonably believes, rather than knows, so long as one's reasonable beliefs are jointly consistent.

The framework automatically includes within the models non-trivial propositions about probabilities on the evidence. For instance, for any proposition X and real number c, we may define $P_{\geq c}[X]$ as $\{w: Pr_w(X) \geq c\}$, the proposition that the probability on one's evidence of X is at least c, which may be true in some worlds and false in others. Thus $P_{\geq c}[X]$ itself receives a probability $Pr_w(P_{\geq c}[X])$ on one's evidence in a world w, so $P_{\geq c}[P_{\geq c}[X]]$ is in turn well-defined: it is the proposition that the probability on one's evidence that the probability on one's evidence of X is at least c is itself at least c. Other propositions about probabilities can be defined similarly; for instance, $P_{>c}[X]$ is $\{w: Pr_w(X) > c\}$, the proposition that the probability on one's evidence of X is greater than c.

When interpreting English renderings in which epistemic terms such as 'one's evidence' and 'the probability of X on one's evidence' occur within the scope of further epistemic vocabulary, it is crucial to remember that the embedded occurrences are here to be read *de dicto* rather than *de re*. Thus even if the probability of X on one's evidence *is* 70%, one cannot substitute '70%' for 'the probability of X on one's evidence' without loss in the scope of another probability operator, for doing so would in effect presuppose that it is certain on one's evidence that the probability of X on one's evidence is 70%. To read the terms for probability *de re* would preclude by *fiat* the non-rigid behaviour that represents uncertainty in an epistemic modal setting, and so prevent us from addressing the very epistemic issues we want to discuss. *De dicto* readings will therefore be understood throughout. They are unambiguously written into the formal notation, but one must bear that in mind when paraphrasing formulas into natural language.

Of course, the setup just described is not the only formal framework conceivable for theorizing about probabilities on evidence of probabilities on evidence. For instance, one could assign each world its own probability distribution, with no requirement that they all be derived by conditionalization on one's evidence in the world from a single prior probability distribution. However, that liberalization would have no chance of vindicating any interesting version of the idea that evidence of evidence is evidence. The present framework is simple, perspicuous, and tightly integrated. In particular, it avoids the *ad hoc* move of postulating a second probability distribution to handle second-order probabilities. In such respects it provides a comparatively hospitable environment for 'evidence of evidence' principles. If they do not thrive here, they are not robust. The framework is also mathematically tractable, permitting us to prove general results and through them to understand what the successes and failures of 'evidence of evidence' principles depend on.

We can use the framework to formulate principles of forms such as this:

(i) $P_{\geq a}[P_{\geq b}[X]] \subseteq P_{\geq c}[X]$

This says that whenever the probability on the evidence that the probability on the evidence of X is at least *b* is itself at least *a*, the probability on the evidence of X is at least *c*. Such an inclusion is *valid* on a finite serial frame <W, R> if and only if it holds for every proposition X $\subseteq$W and every regular probability distribution Pr over W. Given real numbers *a*, *b*, *c* between 0 and 1, under what conditions on <W, R> is such a principle valid on that frame? Section 2 starts to explore the formal prospects for such principles, and to assess the epistemological issues that arise when they are interpreted in terms of evidence.

2. *Positive and negative introspection for evidence*

Here is a simple example of what can go wrong with evidence of evidence principles. Consider a frame <W, R> with just three worlds, where W = {0, 1, 2}, R(0) = {1}, R(1) = R(2) = {2} (so R is serial). Informally, if one is in world 1, it is certain on one's evidence that one is in world 1; if one is in world 1 or world 2, it is certain on one's evidence that one is in world 2. Let the proposition X be {2}, true in world 2 and false in the other two worlds. Then $P_{\geq 1}[X]$ is true in world 1, so $P_{\geq 1}[P_{\geq 1}[X]]$ is true in world 0. But in world 0, one is certain on one's evidence not to be in world 2, so $P_{>0}[X]$ is false. In this frame, even the weakest non-trivial evidence of evidence principles fail, because X is certain on one's evidence to be certain on one's evidence to be true, but also certain on one's evidence to be false. *A fortiori*, any principle of the form $P_{\geq a}[P_{\geq b}[X]] \subseteq P_{>c}[X]$ (for fixed *a*, *b*, *c* between 0 and 1) fails in this frame.

A relevant feature of that frame is that R is irreflexive; neither world 0 nor world 1 has R to itself. In both those worlds, one's total evidence is a falsehood. That is why probability 1 on one's evidence fails to entail truth. In a reflexive frame, if a proposition X is false at a world *w*, one's evidence R(*w*) contains *w*, so the probability of {*w*} conditional on R(*w*) is Pr({*w*})/Pr(R(*w*)), which is nonzero because Pr({*w*}) is nonzero, since Pr is regular, so the probability of X conditional on R(*w*) is less than 1 because X excludes {*w*}. Thus in a reflexive frame, both these principles hold:

(ii) $P_{\geq 1}[P_{\geq b}[X]] \subseteq P_{\geq b}[X]$

(iii) $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$

For $P_{\geq 1}[Y]$ entails Y for any proposition Y; what is certain is true. If a proposition is certain on one's evidence to be likely to some degree on one's evidence, then it *is* likely to that degree on one's evidence; similarly, if it is likely to some degree on one's evidence to be certain on one's evidence, then it *is* likely to at least that degree on one's evidence. By contrast, in any non-reflexive frame, $P_{\geq 1}$ is not a truth-entailing operator, for if a world *w* fails to have R to itself, then the proposition W−{*w*} (true everywhere except *w*) is true throughout R(*w*), and so certain on one's evidence in *w*, even though it is false in *w*. In effect, for regular probability distributions, $P_{\geq 1}$ is simply the familiar necessity operator $\Box$, and reflexivity corresponds to the T axiom $\Box p \supset p$.

The principles (ii) and (iii) demonstrate that one should resist any temptation to regard a requirement for evidence to be true as automatically 'externalist', and so in conflict with 'internalist' evidence of evidence principles. For the truth axiom entails the evidence of evidence principles (ii) and (iii).

Another relevant feature of the above frame is that R is non-transitive: world 0 has R to world 1, and world 1 has R to world 2, but world 0 does not have R to world 2. In modal logic, transitivity corresponds to the 4 axiom $\Box p \supset \Box\Box p$, which in epistemic logic is interpreted as the positive introspection principle that if one knows, one knows that one knows. In the present setting, the modal operator $\Box$ is interpreted as 'one's evidence entails that …', which is equivalent to 'it is certain on one's evidence that …' ($P_{\geq 1}$), because the regularity of the prior probability distribution guarantees that the conditional probability of a proposition X on one's evidence R(w) is 1 if and only if R(w) entails (is a subset of) X. Thus transitivity corresponds to the principle that if one's evidence entails a proposition, then one's evidence entails that one's evidence entails that proposition, or equivalently that if the proposition is certain on one's evidence, then it is certain on one's evidence to be certain on one's evidence. So stated, the principle is not of the form we have been considering, since the second-level condition is in the consequent rather than the antecedent. However, the validity of the 4 axiom is equivalent in any frame to the validity of its contraposed form $\Diamond\Diamond p \supset \Diamond p$, where the modal operator $\Diamond$ is interpreted in the present setting as 'it is consistent with one's evidence that …', which is equivalent to 'there is a nonzero probability on one's evidence that …' ($P_{>0}$), by regularity again. For any frame <W, R> and real number a between 0 and 1, the probability operators $P_{>a}$ and $P_{\geq 1-a}$ are dual to each other in a sense analogous to that in which $\Diamond$ and $\Box$ are: just as $\neg\Box\neg p$ is equivalent to $\Diamond p$ and $\neg\Diamond\neg p$ is equivalent to $\Box$X, so $W - P_{\geq 1-a}[W - X] = P_{>a}[X]$ and $W - P_{>a}[W - X] = P_{\geq 1-a}[X]$ (since $Pr_w(W - X) = 1 - Pr_w[X]$). In particular, for regular probability distributions, the operator $P_{>0}$ is simply the possibility operator $\Diamond$. Thus the transitivity of R is necessary and sufficient for the validity on the frame of this principle:

Positive introspection $\qquad$ $P_{>0}[P_{>0}[X]] \subseteq P_{>0}[X]$

If it is consistent with one's evidence that X is consistent with one's evidence, then X *is* consistent with one's evidence; nonzero evidence of nonzero evidence is nonzero evidence—though of course in a sense in which nonzero evidence for X is compatible with much stronger evidence against X.

Perhaps surprisingly, the validity of Positive Introspection on a frame <W, R> is equivalent to the validity of much weaker principles on that frame, of this form (for $0 < a < 1$):

Weak$_a$ Positive Introspection $\qquad$ $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[X]$

For example, Weak$_{99\%}$ Positive Introspection says that if it is at least 99% probable on one's evidence that a given proposition is at least 99% probable on one's evidence, then that proposition is more than 0% probable on one's evidence. Even that very mild-looking principle is valid only on transitive frames, and so requires the full power of Positive Introspection (proposition 1.0; all such references are to the appendix).

Weak$_a$ Positive Introspection is weaker than Positive Introspection in the sense that the former may hold on a given non-transitive frame <W, R> for every proposition $X \subseteq W$ while the latter does not for a *given* regular probability distribution over W. But for Weak$_a$ Positive Introspection to be *valid* on <W, R> is for it to hold for every proposition $X \subseteq W$ for *every* regular probability distribution over W. Generalizing over probability distributions erases the sensitivity to the quantitative threshold *a* in Weak$_a$ Positive Introspection because a sufficiently non-uniform distribution will exceed the threshold in any case of non-transitivity. We will encounter other examples below where the condition for a principle to be valid is strikingly insensitive to the specific numerical value of a probabilistic threshold, because one can concentrate almost all the weight of prior probability on a few worlds that behave 'badly' in the relevant way. Even requiring the prior probability distribution to be uniform would make less difference than one might expect, because one can often simulate the effect of a highly non-uniform distribution with a uniform distribution, by replacing individual worlds with clusters of worlds mutually indiscernible with respect to the accessibility relation R, where the comparative sizes of the clusters approximates the comparative non-uniform probabilities of the original worlds.

If one reads the positive introspection principle for evidence carelessly, it may sound more or less trivial. But it is not, for the truth of one's total evidence proposition X may not entail that X is at least part of one's total evidence. Given the equation E=K of one's total evidence with one's total evidence, positive introspection for evidence reduces to positive introspection for knowledge.[1] Elsewhere, I have argued against the latter principle (Williamson 2000: 114-130; 2014). Thus I am committed to rejecting positive introspection for evidence, though I will not rehearse the details of the argument here. Of course, positive introspection for knowledge still has defenders (Greco 2014; see also the exchange between Hawthorne and Magidor 2009; 2011 and Stalnaker 2009). Further reason to doubt positive introspection for evidence will emerge at the end of this section. However, the principle requires no extensive critique here, for salient evidence of evidence principles turn out to imply even more obviously problematic variants on *negative* introspection principles.

Negative introspection is equivalent to a principle formulated in the required way, with the second-level condition in the antecedent: if for all one knows one knows something, then one *does* know that thing. In the language of modal logic, that is the 5 axiom $\Diamond\Box p \supset \Box p$. It is valid on all and only frames where R is *euclidean*, in the sense that any points seen from the same point see each other (if R*xy* and R*xz* then R*yz*). Since our evidential interpretation of the frame makes $\Diamond$ equivalent to $P_{>0}$ and $\Box$ to $P_{\geq 1}$, negative introspection amounts to this principle:

Negative Introspection $\quad P_{>0}[P_{\geq 1}[X]] \subseteq P_{\geq 1}[X]$

If it is consistent with one's evidence that one's evidence entails X, then one's evidence *does* entail X.

Negative introspection for knowledge has been recognized by most philosophically sophisticated formal epistemologists since Hintikka (1962) as implausible. For example, consider a good case in which you know by sight that there is an apple on the table, and a corresponding bad case in which you appear to yourself to be in the good case, and still believe that there is an apple on the table, but your belief is false, because what looks like an apple is just a wax replica. By hypothesis, the bad case is indistinguishable from the inside from the good case. In the bad case, for all you know you are in the good case, so for all you know you know that there is an apple on the table; but you do not know that there is an apple on the table, for there is no apple on the table. Thus the bad case is a counterexample to negative introspection for knowledge. Given the limitations of our cognitive powers, the possibility of such mild sceptical scenarios follows almost inevitably from an anti-sceptical view of human knowledge. To idealize away such possibilities is to turn one's back on one of the main phenomena that epistemology is tasked with coming to understand.

Nevertheless, much mainstream epistemic logic outside philosophy, in computer science and theoretical economics, has continued to treat negative introspection for knowledge as axiomatic. This has done less harm than one might have expected, because the focus of such work has been on *multi-agent* epistemic logic, which is mainly concerned with what agents know about what other agents know, and iterations thereof, in particular with common knowledge. When modelling multi-agent epistemic phenomena, it is legitimate to idealize away single-agent epistemic complications, because they constitute noise with respect to the intended object of study. Similarly, in modelling dynamic epistemic phenomena, it is legitimate to idealize away synchronic epistemic complications, because they too constitute noise with respect to the intended object of study. But an idealization may be legitimate at one level of magnification and not at another: the astronomer can sometimes treat planets as point masses; the geologist cannot. Mainstream epistemology turns up the magnification on individual epistemic processes to a level at which negative introspection for knowledge is no longer a legitimate idealization.

What is the relation between negative introspection for knowledge and for evidence? Given E=K, the two principles stand or fall together. In the bad case, for all one knows one knows X, but one doesn't know X; similarly, it is consistent with one's evidence that one's evidence entails X, but one's evidence doesn't entail X. Even if one's evidence is restricted to the contents of one's non-inferential observational knowledge, mild sceptical scenarios can still arise for it, at least on the assumption that the contents of observation can concern one's physical environment. If all evidence must be true, then negative introspection requires all evidence to be in principle immune to sceptical scenarios, for example by concerning mental states that are essentially just as they appear to be to the agent. Such radical foundationalism depends on an antiquated view of the mind which there is no need to argue against here.

The position is more complex for views that allow false evidence. If my evidence may entail that there is an apple on the table, even though there is no apple on the table, then one's evidence in the bad case may include the proposition that there is an apple on the table after all. For instance, suppose that one's evidence is just what one rationally takes for granted, and that one can rationally take for granted something false. Then, in effect, negative introspection says that if it is consistent with what one rationally takes for granted that what one rationally takes for granted entails X, then what one rationally takes for granted *does* entail X. In contraposed form: if what one rationally takes for granted does not entail X, then what one rationally takes for granted entails that what one rationally takes for granted does not entail X. How to assess that principle may not be immediately obvious.

Although negative introspection does not imply that one's evidence is true, it does imply that *one's evidence entails that* one's evidence is true. More precisely, the consequence is that one's evidence entails that if one's evidence entails X, then X is true. Some such conditionals will be amongst one's evidence's false entailments, if it has any. Negative introspection implies the principle because the former corresponds to the euclidean frame condition that if Rxy and Rxz then Ryz; putting $z = y$ gives the condition that if Rxy then Ryy, which we may restate by calling R *quasi-reflexive*: every seen world sees itself. In modal logic, quasi-reflexivity corresponds to the quasi-truth principle $\Box(\Box p \supset p)$. It follows from negative introspection in any normal modal logic.[2,3] Whether it is rational to take for granted that something is rationally taken for granted only if it is true may again not be obvious. The quasi-reflexivity of the accessibility relation is also equivalent to the validity of the evidence of evidence principle (iii) above on the frame, for any given real number a strictly between 0 and 1 (proposition 3.0): if the probability on one's evidence that X is certain on one's evidence is at least a, then the probability of X on one's evidence is at least a. Again, the exact value of the numerical probability parameter a makes no difference, as long as it is not extremal.

We can consider negative introspection for evidence from a different angle. Imagine various pieces of evidence coming in from various sources. Normally, no one of these evidence propositions entails about itself that it *exhausts* one's total evidence. Suppose, for instance, that your evidence includes both the proposition FLASH that there is a flash and the proposition SQUEAK that there is a squeak. FLASH is simply neutral as to whether your evidence also includes the proposition SQUEAK that there is a squeak, and *vice versa*. Now suppose that no other evidence comes in. Then the conjunction FLASH & SQUEAK exhausts your evidence, but it does not itself entail that it exhausts your total evidence. For all the conjunction says, your evidence might also include the proposition HUM that there is a hum. Although some special evidence might somehow manage to entail of itself that it exhausts your evidence, there is no reason to expect evidence to do that in general. On almost any view, one's total evidence is usually much richer than FLASH & SQUEAK, but that does not mean that it entails its own totality. Nor is there any reason to postulate a meta-device, guaranteed to be in perfect working order, for surveying all one's evidence, including the evidence generated by the meta-device itself. Thus, for the sake of simplicity, we can work with the case where your total evidence is just FLASH & SQUEAK, since it is not structurally misleading. Although it is consistent with your total evidence that your total evidence entails (by including) HUM, your total evidence does *not* entail HUM, since FLASH & SQUEAK may be true while HUM is false. Hence the case is a counterexample to negative introspection for evidence. Moreover, as a template for counterexamples it works on a wide range of theories of evidence. It does not assume E=K; it does not even assume that all evidence is true. Thus, except under extreme idealizations, negative introspection for evidence is not a reasonable hypothesis.

That simple case also casts doubt on positive introspection for evidence. On the face of it, the flash and the squeak could occur without being part of your evidence, even implicitly. Thus FLASH & SQUEAK does not entail that your evidence entails FLASH & SQUEAK. So if FLASH & SQUEAK just is your evidence, your evidence may entail FLASH & SQUEAK even though your evidence does not entail that your evidence entails FLASH & SQUEAK. Someone might try to avoid that result by positing that one's evidence consists wholly of propositions that can only be true by being part of one's evidence, but that idea too threatens to degenerate into radically naive foundationalism.

So far, we have mainly assessed evidence of evidence principles that in effect transcribe principles from epistemic logic into the probabilistic idiom. Sections 3 and 4 discuss a wider range of evidence of evidence principles, and discusses whether they imply problematic forms of so-called introspection.

3. *Threshold Transfer*

The simplest interesting frames for single-agent epistemic logic are *partitional*: the accessibility relation R is an equivalence relation—it is reflexive, symmetric, and transitive—and so partitions W into mutually exclusive, jointly exhaustive subsets of the form R(*w*). Since any symmetric transitive relation is euclidean, and any reflexive euclidean relation is symmetric, R is an equivalence relation if and only if it is reflexive, transitive, and euclidean. Thus what the class of such frames validate about knowledge is that it entails truth and satisfies positive and negative introspection. We saw in section 2 how problematic positive and negative introspection are for both knowledge and evidence. However, simple cases make good starting-points, so we begin by reinterpreting partitional frames in terms of evidence, so that the evidence forms a partition.

Suppose that accessibility for evidence is an equivalence relation. Consequently, if one's evidence in a world *w* is consistent with one's being in a world *x*, so R*wx*, then one's evidence in *w* is the same as one's evidence in *x*, so any proposition has the same probability on the evidence in *w* as in *x*. This validates a strong principle about posterior probabilities:

Transfer$_a$ $\qquad\qquad$ $P_{>0}[P_{\geq a}[X]] \subseteq P_{\geq a}[X]$

In other words, if the probability on one's evidence that the probability on one's evidence of X is at least *a* is itself nonzero, then the probability on the evidence of X *is* at least *a*. Transfer$_a$ holds for every proposition X and real number *a* in any partitional frame. For if $P_{>0}[P_{\geq a}[X]]$ is true in *x*, then $P_{\geq a}[X]$ is true at some world *y* in R(*x*); if the frame is partitional, R(*y*) = R(*x*), so $Pr_y(X) = Pr_x(X)$, so $P_{\geq a}[X]$ is true in *x* too. In such frames, evidence of evidence is always *perfect* evidence of evidence, so no wonder evidence of evidence is evidence.

Many principles reminiscent of Transfer$_a$ hold *only* in partitional frames (for example, Williamson 2000: 311-15). However, what was needed to validate Transfer was only for R to be *quasi-partitional*, in the sense that whenever $y \in R(x)$, R(*x*) = R(*y*), in other words, if R*xy* then for all z, Rxz if and only if Ryz. As is easily seen, a relation is quasi-partitional if and only if it is both transitive and euclidean. Transfer$_a$ does not require reflexivity; it can hold even if some evidence is false. An example of a quasi-partitional but not partitional frame is one with just two worlds, 0 and 1, where R(0) = R(1) = {1}. It is not partitional because the world 0 is not in any set of the form R(w). The relation R is non-reflexive, because 0 does not have R to itself. R is also non-symmetric, because R01 but not R10. Nevertheless, Transfer$_a$ holds in such a frame. Conversely, for $0 < a < 1$, Transfer$_a$ is valid *only* on quasi-partitional frames (Appendix, proposition 2.0). Thus Transfer$_a$ is equivalent to quasi-partitionality. Again, the numerical value of the parameter a does not matter, as long as it is not extremal.

Since Transfer requires quasi-partitionality, it implies positive and negative introspection. Given their implausibility, we must seek weaker evidence of evidence principles, ones that do not make the outer probability operator redundant. Here is a natural candidate:

Threshold$_a$ Transfer $\quad P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq a}[X]$

In other words, whenever the probability on the evidence that the probability on the evidence of X is at least a is itself at least a, the probability on the evidence of X *is* at least a. Taking a as the threshold for something to be probable on the evidence (with $a > ½$), we can read Threshold$_a$ Transfer as saying that if it is probable on one's evidence that a hypothesis is probable on one's evidence, then that hypothesis *is* probable on one's evidence.

Transfer entails Threshold$_a$ Transfer for all a. For when $a = 0$, Threshold$_a$ Transfer is trivial, and when $a > 0$, $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[P_{\geq a}[X]]$. But not even Threshold$_a$ Transfer for all values of a together entails Transfer.

Here is an example of a frame on which Threshold$_a$ Transfer is valid while Transfer$_a$ is not. As before there are just two worlds, 0 and 1, where R01 but not R10, and R11, but this time R00 too. Thus R is reflexive as well as transitive, though not symmetric. R is also not euclidean (since R01 and R00 but not R10). Transfer$_a$ fails on this frame for any non-extremal value of a and any regular probability distribution. For $P_{\geq 1}[\{1\}]$ is true at 1 but false at 0, even though $P_{>0}[P_{\geq 1}[\{1\}]]$ is true at 0. Since R(0) = {0, 1} while R(1) = {1}, the total evidence propositions are not even mutually exclusive. Nevertheless, Threshold$_a$ Transfer is valid on this frame, whatever the probability distribution and the value of a. This follows from the general result that for all a between 0 and 1, Threshold$_a$ Transfer is valid on every frame that is *near-partitional* in this sense: whenever $y \in$ R(x), either R(y) = R(x) or R(y) = {y}. The new two-world frame is clearly near-partitional. The converse also holds for all a strictly between 0 and 1: every frame on which Threshold$_a$ Transfer is valid is near-partitional (proposition 4.0). This is another instance of the phenomenon already noted in relation to Weak$_a$ Positive Introspection: the validity on a frame of 'evidence of evidence' principles with a probabilistic threshold is often insensitive to the value of that parameter. In this case, whenever a and b are both strictly between 0 and 1, Threshold$_a$ Transfer and Threshold$_b$ Transfer are valid on exactly the same frames.

For R(y) to be {y} is for one's evidence in y to entail all and only truths in the world y: one's evidence tells the whole truth and nothing but the truth about y. That is a wildly idealized scenario. Thus, in practice, near-partitionality is very close to partitionality, and Threshold$_a$ Transfer very close to Transfer, too close for it to be a useful weakening.

Another way to assess the strength of Threshold$_a$ Transfer is by noting that the propositional modal logic of near-partitional (and serial) frames can be axiomatized by this set of axioms:

D $\quad \Box p \supset \Diamond p$

Q-T $\quad \Box(\Box p \supset p)$

4 $\quad \Box p \supset \Box\Box p$

Q-5 $\quad \Diamond p \supset \Box(\Diamond p \vee (q \supset \Box q))$

More precisely, the smallest normal modal logic with D, Q-T, 4, and Q-5 as theorems is sound and complete for the class of near-partitional serial frames. D corresponds to seriality, Q-T to quasi-reflexivity, 4 to transitivity, and Q-5 to a slight weakening of the euclidean property. The two disjuncts in the consequent of Q-5 correspond to the two disjuncts in the definition of near-partitionality (in the same order). On the relevant interpretation of the modal operators in terms of evidence, D says that one's evidence is compatible with what it entails, Q-T that one's evidence entails that one's evidence is true, and 4 that one's evidence obeys positive introspection. Q-5 weakens negative introspection (corresponding to the 5 principle $\Diamond p \supset \Box\Diamond p$, equivalent to the contraposed principle $\Diamond\Box p \supset \Box p$) by saying that if one's evidence is consistent with a proposition, then one's evidence entails that either one's evidence is consistent with the proposition or any given truth is entailed by one's evidence.

Positive introspection for evidence was already discussed in section 2. As for axiom Q-5, it is no more plausible than negative introspection, even though logically it is slightly weaker. Under E=K, Q-5 implies claims like this about the bad case (at least if one knows E=K): if for all one knows there is no apple on the table, then one knows that either for all one knows there is no apple on the table or if there is life on other planets then one knows that there is life on other planets. But the antecedent is true in the bad case: since there is no apple on the table, for all one knows there is no apple on the table. Thus Q-5, read epistemically, generates the claim that, in the bad case, one knows that either for all one knows there is no apple on the table or if there is life on other planets then one knows that there is life on other planets. But since there is no useful connection between the disjuncts, one's only way of knowing the disjunction is by knowing one of the disjuncts (sometimes one knows a disjunction without knowing any disjunct, for instance when the disjunction is an instance of the law of excluded middle or the content of some disjunctive testimony, but in such cases there is a useful epistemic connection between the disjuncts). Hence either one knows that for all one knows there is no apple on the table, or one knows that if there is life on other planets then one knows that there is life on other planets. But one does not know that for all one knows there is no apple on the table, because for all one knows one knows that there is an apple on the table. One also does not know that if there is life on other planets then one knows that there is life on other planets, for one has no special access to whether there is life on other planets.

Even if we bracket E=K, and do not assume that all evidence is true, the case in section 2 where one's total evidence is just the conjunction FLASH & SQUEAK raises as severe a problem for Q-5, read in terms of evidence, as it does for negative introspection. Thus weakening negative introspection by the second disjunct in the consequent makes no significant difference to the plausibility of Q-5.

In brief, Threshold$_a$ Transfer principles weaken negative introspection too slightly to regain plausibility.

We can also consider principles intermediate between Threshold$_a$ Transfer and Weak$_a$ Positive Introspection, in other words, principles of the form $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq b}[X]$ for $a > b > 0$. For example, we might set $b = a^2$. Such principles are only valid on transitive frames, since they all entail Positive Introspection; the question is where the corresponding frame conditions come between transitivity and near-partitionality. In general, such intermediate principles require far more than transitivity. For instance, consider frames where $W = \{x\} \cup Y \cup Z$; $\{x\}$, Y, and Z are pairwise disjoint; $R(x) = W$, $R(y) = \{y\} \cup Z$ for $y \in Y$, and $R(z) = Z$ for $z \in Z$. Such frames are reflexive and transitive but neither symmetric nor euclidean. Let $|Y| = n^2$, $|Z| = n$, and Pr be the uniform probability distribution on W. Then $Pr_y(Z) = n/(n + 1)$ for $y \in Y$, and $Pr_z(Z) = 1$ for $z \in Z$. Thus $Y \cup Z \subseteq P_{\geq a}[Z]$ for sufficiently large n, in which case $Pr_x(P_{\geq a}[Z]) \geq (n^2 + n)/(n^2 + n + 1)$, so $x \in P_{\geq a}[P_{\geq a}[Z]]$ for sufficiently large n. But $Pr_x(Z) = n/(n^2 + n + 1)$, so $x \notin P_{\geq b}[Z]$ for sufficiently large n. Thus, for any given $a > b > 0$, for sufficiently large n the frame invalidates the principle $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq b}[X]$. For those who accept Positive Introspection but reject Negative Introspection for evidence, it is unclear what principled objection there might be to modelling evidence on such frames.[4]

An important frame condition for such principles is what we may call *quasi-nestedness*: <W, R> is quasi-nested if and only if whenever R*wx*, R*wy*, R*xz*, and R*yz*, then either R*xy* or R*yx*. Very roughly, if two points visible from a given point are invisible from each other, then their fields of vision are disjoint. Quasi-nestedness has a long history in epistemic logic, under varying terminology (Geanakoplos 1989, Dorst 2016). One can show that for all a, $b \in [0, 1]$, the natural-looking principle $Pr_{\geq a}[Pr_{\geq b}[X]] \subseteq Pr_{\geq ab}[X]]$ is valid on every finite serial transitive quasi-nested frame; conversely, any finite serial frame on which the principle is valid for all a, $b \in [0, 1]$ is transitive and quasi-nested (proposition 15.5). The difference in validity conditions between this principle and Threshold$_a$ Transfer is one way in which the numerical values of the thresholds do make a difference.

How plausible is quasi-nesting as an epistemic condition? It does not correspond to the validity of a formula in the language of propositional modal logic, for the modal system S4 is sound and complete both for the class of all reflexive transitive frames, many of which are not quasi-nested, and for the class of reflexive transitive tree frames, all of which are quasi-nested (see Blackburn, de Rijke, and Venema 2001: 353). Thus no formula of the language is valid in all and only quasi-nested frames. The distinctive consequences of quasi-nesting appear only in a more expressive language, such as one with probability operators. We shall see later in section 4 that many epistemically reasonable frames are *not* quasi-nested.

4. Comparative transfer

Probabilistic confirmation can be understood in two ways, absolute and comparative. In the absolute sense, evidence confirms a hypothesis if the probability of the hypothesis on the evidence reaches some fixed threshold (such as 80%). In the comparative sense, evidence confirms a hypothesis if the probability of the hypothesis on the evidence exceeds its prior probability. The absolute sense concerns *high* probability, the comparative sense high*er* probability, or probability raising. Various hybrids of the two standards are also conceivable, but it is better to start with the simple contrast. Clearly, evidence of evidence principles can also be read in either the threshold way or the comparative way. Again, various hybrid readings are also conceivable, and again it is better to start with the simpler ones. In the previous section, threshold readings of 'evidence of evidence' principles turned out to be unpromising. This section considers the comparative reading.

Probability-raising is easy to formalize in the present framework, since it provides both the prior probability distribution Pr and, for each world w, the posterior probability distribution Pr_w which conditionalizes Pr on one's (new) evidence in w, R(w). For present purposes, it does not matter whether Pr embodies earlier background evidence or not. Thus the proposition that the (new) evidence has raised the probability of a proposition X is simply $\{w: Pr_w(X) > Pr(X)\}$, which we notate as $P_>[X]$. It is just the set of worlds in which the posterior probability of X is greater than its prior probability. When the 'evidence of evidence' slogan is so understood in terms of probability-raising, it becomes this:

Comparative Transfer $\qquad P_>[P_>[X]] \subseteq P_>[X]$

In other words, if the evidence raises the probability that the evidence raises the probability of X, then the evidence *does* raise the probability of X. What are the prospects for Comparative Transfer? In which frames is it valid?

An initial observation is encouraging: Comparative Transfer is valid on some frames on which Negative and Positive Introspection both fail. In that way it is less demanding than the Threshold Transfer principles. In fact, the validity of the Threshold Transfer principles on a frame is neither necessary nor sufficient for the validity of Comparative Transfer (proposition 14.0). On closer inspection, the picture is more complicated and less rosy.

We can start with the case of Positive Introspection, or in modal terms the 4 axiom. Although its validity does not follow from that of Comparative Transfer, the validity of this weakening of it does (corollary 8.1):

(iv) $\qquad \Box\Box p \supset \Box\Box\Box p$

In other words, although the evidence may entail a proposition without entailing that it entails it, if the evidence entails that it entails a proposition, then it entails that it entails that it entails that proposition. In terms of the accessibility relation between worlds, the 4 axiom corresponds to transitivity, in other words, if you can get from *x* to *z* in two steps of accessibility, you can get there in one step. For comparison, (iv) corresponds to the feature that if you can get from *x* to *z* in *three* steps of accessibility, you can get there in two. It is hard to see what independent theoretical reason there might be for accepting the weaker principle that would not also be a reason for accepting the stronger. Relevant objections to Positive Introspection for knowledge generalize to objections to the weaker principle (Williamson 2000: 120-1).[5] If the proposition that one's evidence entails SQUEAK is added to one's evidence, it does not follow that the proposition that it has been added to one's evidence has itself been added to one's evidence.

As with previous 'evidence of evidence' principles, the problems for Comparative Transfer do not end with (slight weakenings of) the contested principle of Positive Introspection. They extend to (slight weakenings of) the far more generally rejected principle of Negative Introspection. For example, on any frame on which Comparative Transfer is valid, so is (v):

(v) $(\Box p \supset p) \vee (\Diamond q \supset \Box\Diamond q)$

Since *p* and *q* are independent variables, (v) says in effect that each world is either reflexive or euclidean: if the former, all instances of the T axiom hold at it, if the latter, all instances of the 5 axiom do. For those (like the author) who hold that all evidence is true, (v) is unproblematic, because it is true thanks to the first disjunct. But for those who hold that evidence may be false, (v) has the effect that Negative Introspection must hold of all propositions whatever whenever one has some false evidence. Epistemologically, it is quite unclear why the T and 5 axioms should play such complementary roles. To turn the screw, (v) requires that if one has false evidence about one topic, then one's evidence about some completely different topic conforms perfectly to Negative Introspection. But if the T axiom for evidence has false instances in some situations, and Negative Introspection has false instances, it is almost inevitable that in some combined situations both principles will have false instances, perhaps about unrelated topics, which is enough to violate (v). For example, if one's evidence is consistent with the proposition that one's evidence entails HUM, even though one's evidence does not in fact entail HUM, why should that prevent one from getting false evidence about something else, if false evidence is in general an option?

However, even for those who accept the T axiom, Comparative Transfer has problematic consequences. They include weakenings of the B axiom $p \rightarrow \Box\Diamond p$, which corresponds to the symmetry of the accessibility relation. B is implausible in epistemic logic, for reasons closely connected to the implausibility of Negative Introspection. Read in terms of knowledge, B says that if something obtains, one knows that for all one knows it obtains. But consider any pair of a good case, in which all goes well, things are as they seem and one has plenty of knowledge, and a bad case, a sceptical scenario in which one seems to oneself to be in the good case but things are not as they seem, much of what one seems to oneself to know is false, and one knows very little. In the good case, one knows things incompatible with being in the bad case. In the bad case, for all one knows one is in the good case, so for all one knows one knows things incompatible with being in the bad case. Thus, if one is in the bad case, one does *not* know that for all one knows one is in the bad case. Hence the B axiom fails for knowledge. Consequently, on the equation E=K, the B axiom also fails for evidence.

Even if evidence is not equated with knowledge, the B axiom faces similar counterexamples to those for knowledge. There are two cases to consider: either false evidence is disallowed (R must be reflexive), or false evidence is allowed (R can be non-reflexive).

First, suppose that false evidence is disallowed. Start with a good case where E is true and part of one's evidence. Typically, there will be a bad case where one's evidence is consistent with being in the good case, but E is false and so not part of one's evidence. The bad case may be one which seems from the inside just as the good case does, or it may simply be that in both cases one's capacity to survey one's evidence as a whole is imperfect, and in the bad case one cannot ascertain that one's evidence lacks E and add that fact to one's evidence. Either way, in the bad case, ¬E is true, but one's evidence does not entail that one's evidence is consistent with ¬E, for one's evidence in the bad case is consistent with being in the good case, where one's evidence is inconsistent with ¬E. Thus, on its evidential reading, the B axiom fails in the bad case.

Now suppose instead that false evidence is allowed. Typically, the motivation for allowing it is to let one's evidence in both good and bad cases be the content of appearances common to those cases. Thus one's evidence is the same in the two cases, but true in the good case and false in the bad case. Since one's evidence is false in the bad case, one's evidence in both cases entails that one is not in the bad case. Therefore, since one's evidence is true in the good case, and entails that one is not in the bad case, one's evidence in both cases does not entail that one's evidence is consistent with one's being in the bad case. Thus, on its evidential reading, the B axiom fails again in the bad case.

In brief, whether false evidence is allowed or not, the B axiom fails on reasonable views of evidence.

As already noted, the validity of Comparative Transfer does not require the validity of the B axiom for evidence. Comparative Transfer is valid on some reflexive but non-symmetric frames, such as the two-world frame <W, R> where W = {good, bad}, R(good) = {good} and R(bad) = {good, bad} (by proposition 9.0). Since that is exactly the problematic structure under discussion, Comparative Transfer may look to be out of the wood. But it is not. For although its validity on a frame does not require the accessibility relation to be symmetric *everywhere*, it does limit how widespread the failures can be. More specifically, if Comparative Transfer is valid on a finite serial frame <W, R>, then for all worlds w, x, y, z in W, if Rwx, Rxy, and Ryz, then either Rxw, or Ryx, or Rzy. In other words, if there is a chain of three successive links of accessibility, at least one of those links is bidirectional, an instance of symmetry. In the language of modal logic, that frame condition corresponds to this axiom (proposition 7.9):

(vi) $$p \supset \Box(q \supset (\Diamond p \vee \Box(r \supset (\Diamond q \vee \Box\Diamond r)))$$

The three variables p, q, and r are mutually independent.

To see why (vi) is problematic, note that the contrast between good and bad cases is not all-or-nothing. A case may be good in one respect, bad in another. For example, one may suffer a minor illusion about the size, shape, or distance of a particular building, or the direction from which a particular sound is coming, while continuing to gain large amounts of perceptual evidence about one's environment in many other respects. Indeed, such mixed cases may be usual in everyday life. For simplicity, consider worlds which differ from each other only in goodness or badness in three independent minor respects, 1, 2, and 3. Let each subset of {1, 2, 3} label the world (or scenario) which is good in the respects it contains and bad in the other respects. Thus world {} is bad in all three respects, while world {1, 2, 3} is good in all three respects. In line with the previous discussion, we naturally identify accessibility with the subset relation, for $S \subseteq S^*$ ($\subseteq W$) just in case there is no respect in which S is good and S* is not (which is what would block accessibility). This accessibility relation is reflexive and transitive, but not symmetric.[6] For instance, {1} has R to {1, 2} but not *vice versa*, because {1} stands to {1, 2} as the bad case to the good case in respect 2, while they do not differ in respects 1 and 3. Thus {} has R to {1}, {1} has R to {1, 2}, and {1, 2} has R to {1, 2, 3}, and none of these R links is reversible: {1} lacks R to {}, {1, 2} lacks R to {1}, and {1, 2, 3} lacks R to {1, 2}. Thus the frame condition corresponding to (vii) is violated, so Comparative Transfer is violated. In terms of (vi) itself, there is a counterexample where *p* is the proposition that one is in a bad case in respect 1, *q* is the proposition that one is in the bad case in respect 2, and *r* is the proposition that one is in the bad case in respect 3.

For reflexive frames, where evidence is always true, the counterexamples can be slightly simplified, for we only require a chain of two steps of R for one of the steps to be bidirectional. More specifically, if Comparative Transfer is valid on a finite reflexive frame <W, R>, then for all worlds *w*, *x*, *y* in W, if R*wx* and R*xy*, then either R*xw* or R*yx* (proposition 7.10) In the language of modal logic, that frame condition corresponds to this axiom (proposition 7.11):

(vii) $p \supset \Box(q \supset (\Diamond p \vee \Box\Diamond q))$

In this case, counterexamples can make do with two respects rather than three. Since the frame just described is already reflexive, there is no need to elaborate.

In brief, although Comparative Transfer is on balance weaker than the Threshold$_a$ Transfer principles, its epistemological consequences are still implausibly strong.

We briefly return to the family of multiplication principles $Pr_{\geq a}[Pr_{\geq b}[X]] \subseteq Pr_{\geq ab}[X]]$ considered in section 3. A finite serial frame validates all those principles if and only if it is transitive and quasi-nested. But no frame of the kind just considered with sceptical scenarios in more than one respect is quasi-nested, for {} has the accessibility relation R to both {1} and {2}, and they both have R to {1, 2}, but neither of {1} and {2} has R to the other. Although quasi-nesting holds in the two-world non-symmetric frame corresponding to a sceptical scenario in only one respect (since that frame is connected), it is hard to think of any principled reason for permitting the one-respect set-up but rejecting all the multiple-respects set-ups. Thus a principled defence of quasi-nesting will be forced to reject the natural modelling of sceptical scenarios. On its epistemic reading, the multiplication principle $Pr_{\geq a}[Pr_{\geq b}[X]] \subseteq Pr_{\geq ab}[X]]$ is correspondingly unreasonable, for at least some values of a and b (it is trivial when $a = 0$ or $b = 0$).

Back to Comparative Transfer. From a technical point of view, it behaves rather differently from the other principles considered in this chapter. Unlike them, it essentially involves the prior probability distribution Pr, in ways that cannot be reduced to the posterior distributions Pr_w. Whether it holds at a given world depends on global features of the frame, and can depend on what happens at worlds to which the starting world does not even bear the ancestral of the accessibility relation, for those worlds contribute to the priors with which the posteriors are being compared. One manifestation of this global aspect is that in some cases Comparative Validity is valid on each of two mutually disjoint frames but invalid on their union, even when the two original frames are isomorphic to each other (proposition 13.0). That cannot happen in the standard model theory of ordinary modal logic or the other principles considered in this paper, because at every stage the generalizations over worlds are restricted by the accessibility relation, and so never cross the boundary between one of the original frames and the other. As a corollary, by contrast with the Threshold$_a$ Transfer principles, there is no set of formulas in the standard language of propositional modal logic whose validity on a frame is equivalent to the validity of Comparative Transfer, for the class of frames on which such formulas are valid is closed under disjoint unions.

In part for the reason just discussed, the model theory is more intricate for Comparative Transfer than for the principles discussed earlier. The partial results in the appendix give some indication of its complexity. They do not include a necessary and sufficient condition, in non-probabilistic terms, for Comparative Transfer to be valid on a frame <W, R>. That is left as an open problem for the interested reader. The epistemological implausibility of Comparative Transfer makes it non-urgent for present purposes.

5. *Conclusion*

This paper has considered by no means all imaginable renderings of the slogan 'evidence of evidence is evidence' in the framework of single-agent synchronic epistemic logic with probabilities. However, it has assessed the most natural candidates, with bleak results: they all have epistemologically implausible consequences, which fail even in the mildest sceptical scenarios.

To overturn this provisional negative verdict, it will not be enough to formulate another candidate rendering that has not been considered here. Such moves come far too cheap. Rather, what would be needed are serious results to the effect that the candidate principle is valid in a good variety of frames invalidating the implausible introspection principles.

If that challenge is unmet, we may tentatively conclude that, in any reasonable sense, evidence of evidence is not always evidence—not always in the single-agent synchronic case, and *a fortiori* not always in the multi-agent and diachronic cases either.

Of course, none of this means that evidence of evidence is not *typically* evidence. It would also be useful to have positive results giving less demanding sufficient conditions in the present framework for *most* evidence of evidence to be evidence, in senses close to those in this paper.[7]

Appendix

For convenience, definitions in the main text are repeated here as they arise.

Weak$_a$ Positive Introspection $\qquad P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[X]$

Proposition 1.0: For $a \in (0, 1)$, Weak$_a$ Positive Introspection is valid on a finite serial frame iff it is transitive.

Proof: Suppose that $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[X]$ is valid on a finite serial frame <W, R>. Suppose further that Rxy, Ryz, but not Rxz. Thus $y \neq z$. If W = {y, z}, define a probability distribution Pr over W thus:

$\Pr(\{y\}) = 1 - a$

$\Pr(\{z\}) = a$

If W ≠ {y, z}, so $|W| = n > 2$, instead define Pr thus:

$\Pr(\{y\}) = a(1 - a)$

$\Pr(\{z\}) = a$

$\Pr(\{u\}) = (1 - a)^2/(n - 2) \qquad$ for $u \notin \{y, z\}$

The following argument works on both definitions.

Since $z \in R(y)$, $\Pr_y(\{z\}) = \Pr(\{z\} \cap R(y))/\Pr(R(y)) = \Pr(\{z\})/\Pr(R(y)) \geq \Pr(\{z\}) = a$, so $y \in P_{\geq a}[\{z\}]$.

Moreover, since $y \in R(x)$ and $z \notin R(x)$:

$\Pr_x(\{y\}) = \Pr(\{y\})/\Pr(R(x)) \geq \Pr(\{y\})/\Pr(W - \{z\}) = \Pr(\{y\})/(1 - a) \geq a(1 - a)/(1 - a) = a$

Hence $x \in P_{\geq a}[\{y\}]$, so $\Pr_x(\{y\}) \geq a$. But $\{y\} \subseteq P_{\geq a}[\{z\}]$ so $\Pr_x(P_{\geq a}[\{z\}] \geq a$, so $x \in P_{\geq a}[P_{\geq a}[\{z\}]]$. Thus, by hypothesis, $x \in P_{>0}[\{z\}]$, so $P_x(\{z\}) > 0$; but R is regular, so $z \in R(x)$, contrary to hypothesis. Thus R is transitive after all.

Conversely, suppose that R is transitive and Pr is a regular probability distribution over W. Suppose that $x \in P_{\geq a}[P_{\geq a}[X]]$. Thus $\Pr_x(P_{\geq a}[X]) \geq a > 0$, so there is a $y \in P_{\geq a}[X] \cap R(x)$. Thus $\Pr_y(X) \geq a > 0$, so there is a $z \in X \cap R(y)$. But then $z \in R(x)$ because $y \in R(x)$, $z \in R(y)$, and R is transitive. Hence $z \in X \cap R(x)$, so $\Pr_x(X) > 0$ because Pr is regular, so $x \in P_{>0}[X]$. This shows that $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[X]$, as required.

Corollary 1.1: For any a, $b \in (0, 1)$ and finite serial frame <W, R>, Weak$_a$ Positive Introspection is valid on <W, R> iff Weak$_b$ Positive Introspection is.

Transfer$_a$ $\qquad P_{>0}[P_{\geq a}[X]] \subseteq P_{\geq a}[X] \qquad$ (for $a \in [0, 1]$)

A frame <W, R> is *quasi-partitional*: for all $x \in W$, if $y \in R(x)$ then $R(x) = R(y)$ (equivalently, R is transitive and euclidean).

Proposition 2.0: For $a \in (0, 1)$, Transfer$_a$ is valid on a finite serial frame iff it is quasi-partitional.

Proof: Suppose that Transfer_a is valid on a finite serial frame <W, R>. If $|W| = 1$ the result is trivial, so we may assume that $|W| = n \geq 2$. Let $x, z \in W$, and $y \in R(x)$. Choose b so that $\max\{a, 1 - a\} < b < 1$. Define a probability distribution Pr over W by setting:

$\Pr(\{z\}) = b$

$\Pr(\{u\}) = (1 - b)/(n - 1)$ for $u \neq z$

If $z \in R(y)$, $\Pr_y(\{z\}) = \Pr(\{z\})/\Pr(R(y) \geq \Pr(\{z\}) = b > a$, so $y \in P_{\geq a}[\{z\}]$. Since $y \in R(x)$: $x \in P_{>0}[P_{\geq a}[\{z\}]]$, so by Transfer_a, $y \in P_{\geq a}[\{z\}]$, so $z \in R(x)$. On the other hand, if $z \notin R(y)$, then $\Pr_y(W-\{z\}) = 1$, so $y \in P_{\geq a}[W-\{z\}]$. Since $y \in R(x)$, $x \in P_{>0}[P_{\geq a}[W-\{z\}]]$, so $x \in P_{\geq a}[W-\{z\}]$ by Transfer_a, so $1 - \Pr_x(\{z\}) = \Pr_x(W-\{z\}) \geq a$. Hence $\Pr_x(\{z\}) \leq 1 - a$. But, as before, if $z \in R(x)$ then $\Pr_x(\{z\}) \geq b > 1 - a$. Thus $z \notin R(x)$. So for all $z \in W$, $z \in R(y)$ iff $z \in R(x)$. Thus $R(y) = R(x)$. Hence R is quasi-partitional. The converse is routine.

Corollary 2.1: For any $a, b \in (0, 1)$, Transfer_a is valid on a finite serial frame iff Transfer_b is.

Proposition 3.0: For $a \in (0, 1)$, the principle $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$ is valid on a finite serial frame iff it is quasi-reflexive.

Proof: Suppose that $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$ is valid on a finite serial frame <W, R>. If $|W| = 1$ the result is trivial, so we may assume that $|W| = n \geq 2$. Let $x, z \in W$. Define a probability distribution Pr over W just as in the proof of 2.0. Suppose that $z \notin R(z)$. Thus: $R(z) \subseteq W-\{z\}$, so $\{z\} \subseteq P_{\geq 1}[W-\{z\}]$, so $\Pr_x(P_{\geq 1}[W-\{z\}]) \geq \Pr_x(\{z\}) = \Pr(\{z\})/\Pr(R(x)) \geq \Pr(\{z\}) = b > a$. Hence $x \in P_{\geq a}[P_{\geq 1}[W-\{z\}]]$. By hypothesis, $P_{\geq a}[P_{\geq 1}[W-\{z\}]] \subseteq P_{\geq a}[W-\{z\}]$. Thus $x \in P_{\geq a}[W-\{z\}]$. By an argument as in the proof of 2.0, we can show that $z \notin R(x)$. By contraposition, if $z \in R(x)$ then so $z \in R(z)$. Thus R is quasi-reflexive.

Conversely, suppose that R is quasi-reflexive and Pr is a regular probability distribution over W. Suppose that $z \in P_{\geq a}[P_{\geq 1}[X]]$. So $\Pr_z(P_{\geq 1}[X]) \geq a$. Let $u \in P_{\geq 1}[X] \cap R(z)$. Since $u \in R(z)$, by quasi-reflexivity $u \in R(u)$. Hence if $u \notin X$, $\Pr_u(X) < 1$ because Pr is regular. So, since $u \in P_{\geq 1}[X]$, $u \in X$. Thus $P_{\geq 1}[X]] \cap R(z) \subseteq X$, so $a \leq \Pr_z(P_{\geq 1}[X]) \leq \Pr_z(X)$, so $z \in P_{\geq a}[X]$. Thus $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$, as required.

Corollary 3.1: For any $a, b \in (0, 1)$, one of these principles is valid on a finite serial frame iff the other is: $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$ and $P_{\geq b}[P_{\geq 1}[X]] \subseteq P_{\geq b}[X]$.

Threshold_a Transfer $\quad P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq a}[X]$ $\quad$ (for $a \in [0, 1]$)

A frame <W, R> is *near-partitional*: for all $x \in W$, if $y \in R(x)$ then either $R(y) = R(x)$ or $R(y) = \{y\}$.

Proposition 4.0: For $a \in (0, 1)$, Threshold_a Transfer is valid on a finite serial frame iff it is near-partitional.

Proof: Suppose that Threshold_a Transfer is valid on a finite serial frame <W, R>, but <W, R> is not near-partitional. Thus for some w and $x \in R(w)$, $R(x) \neq R(w)$ and $R(x) \neq \{x\}$. Since Threshold_a Transfer is valid on <W, R>, so is the principle $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$, for:

$P_{\geq 1}[X]] \subseteq P_{\geq a}[X]$, so $P_{\geq a}[P_{\geq 1}[X]] \subseteq P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq a}[X]$.

Thus, by 3.0, R is quasi-reflexive, so $\{x\} \subseteq R(x)$, for $x \in R(w)$. Since $R(x) \neq \{x\}$, there is a $y \in R(x)$ with $y \neq x$. Furthermore, by Threshold$_a$ Transfer, the principle $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{>0}[X]$ is also valid on <W, R>, because $P_{\geq a}[X] \subseteq P_{>0}[X]$. Thus, by 1.0, R is transitive, so $R(x) \subseteq R(w)$; since $R(x) \neq R(w)$, there is a $z \in R(w)-R(x)$, so $z \neq y$. Moreover, since $x \in R(x)$, $z \neq x$. Thus x, y, and z are mutually distinct. There are two cases:

Case (i): $W = \{x, y, z\}$. Define a probability distribution Pr over W thus:

$$\Pr(\{x\}) = a(1-a)$$
$$\Pr(\{y\}) = a^2$$
$$\Pr(\{z\}) = 1-a$$

Since $R(x) \subseteq W-\{z\}$, $\Pr(R(x)) \leq \Pr(W-\{z\}) = 1 - \Pr(\{z\}) = a$. Thus, since $y \in R(x)$:

$$\Pr_x(\{y\}) = \Pr(\{y\})/\Pr(R(x)) \geq a^2/a = a$$

Since $y \in R(x)$, $R(y) \subseteq R(x)$ because R is transitive, so $\Pr(R(y)) \leq \Pr(R(x)) \leq a$. Also $y \in R(y)$, because R is quasi-reflexive. Thus:

$$\Pr_y(\{y\}) = \Pr(\{y\})/\Pr(R(y)) \geq \Pr(\{y\})/\Pr(R(x)) \geq a$$

Thus $\{x, y\} \subseteq P_{\geq a}[\{y\}]$. Moreover, $\{x, y\} \subseteq R(w)$ because R is transitive. Since $z \in R(w)$, $R(w) = W$. Hence $\Pr_w(P_{\geq a}[\{y\}]) \geq \Pr_w(\{x, y\}) = \Pr(\{x, y\})/\Pr(R(w)) = \Pr(\{x, y\}) = a(1-a) + a^2 = a$.

So $w \in P_{\geq a}[P_{\geq a}[\{y\}]]$. Hence $w \in P_{\geq a}[\{y\}]$ because Threshold$_a$ Transfer is valid on <W, R> by hypothesis. But $R(w) = W$, so $\Pr_w(\{y\}) = \Pr(\{y\}) = a^2 < a$ because $a < 1$. Hence $w \notin P_{\geq a}[\{y\}]$, which is a contradiction. Thus <W, R> is near-partitional after all.

Case (ii): $|W| = n > 3$. Define a probability distribution Pr over W thus:

$$\Pr(\{x\}) = a(1-a)/(1+a)$$
$$\Pr(\{y\}) = 2a^2/(1+a)$$
$$\Pr(\{z\}) = (1-a)/(1+a)$$
$$\Pr(\{u\}) = a(1-a)/(n-3)(1+a) \qquad \text{for } u \in W-\{x, y, z\}$$

The argument resembles that of case (i); some overlapping parts are omitted.

$\Pr(R(x)) \leq 1 - \Pr(\{z\}) = 2a/(1+a)$. Thus:

$\Pr_x(\{y\}) = \Pr(\{y\})/\Pr(R(x)) \geq (2a^2/(1+a))/(2a/(1+a)) = a$.

Similarly, $\Pr_y(\{y\}) \geq \Pr(\{y\})/\Pr(R(x)) \geq a$. Thus $\{x, y\} \subseteq P_{\geq a}[\{y\}]$. But:

$\Pr_w(P_{\geq a}[\{y\}]) \geq \Pr_w(\{x, y\}) = \Pr(\{x, y\})/\Pr(R(w)) \geq \Pr(\{x, y\}) = (a(1-a) + 2a^2)/(1+a) = a$

So $w \in P_{\geq a}[P_{\geq a}[\{y\}]]$. Hence $w \in P_{\geq a}[\{y\}]$ by Threshold$_a$ Transfer. But $\{x, y, z\} \subseteq R(w)$. Thus: $\Pr_w\{x, y, z\} \leq \Pr_w(R(w))$, so:

$\Pr_w(\{y\}) = \Pr(\{y\})/\Pr(R(w)) \leq \Pr((\{y\})/\Pr(\{x, y, z\}) = 2a^2/(a(1-a) + 2a^2 + 1 - a) = 2a^2/(1+a^2)$.

But $(1-a)^2 > 0$, so $1 > 2a/(1+a^2)$, so $a > 2a^2/(1+a^2)$, so $\Pr_w(\{y\}) < a$. Hence $w \notin P_{\geq a}[\{y\}]$, again a contradiction. Thus <W, R> is near-partitional after all.

Conversely, suppose that <W, R> is semi-partitional. Let $w \in W$ and $X \subseteq W$. Suppose that $\Pr_w(X) < a$. We first note (1):

(1) $\{x: \Pr_x(X) \geq a\} \cap R(w) \subseteq X \cap R(w)$

For if $\Pr_x(X) \geq a > \Pr_w(X)$ then $R(x) \neq R(w)$, so by near-partitionality $R(x) = \{x\}$; since $\Pr_x(X) > 0$, some member of $R(x)$ belongs to X, so $x \in X$. But (1) entails (2):

(2) $\Pr_w(\{x: \Pr_x(X) \geq a\}) \leq \Pr_w(X) < a$

By contraposition, if $\Pr_w(\{x: \Pr_x(X) \geq a\}) \geq a$ then $\Pr_w(X) \geq a$. In other words, $P_{\geq a}[P_{\geq a}[X]] \subseteq P_{\geq a}[X]$, so Threshold$_a$ Transfer at a is valid on <W, R>, as required.

Corollary 4.1: For $a, b \in (0, 1)$, Threshold$_a$ Transfer is valid on a finite serial frame iff Threshold$_b$ Transfer is.

Lemma 5.0: Let <W, R> be a finite serial frame <W, R>, $w \in W$, and $X \subseteq W$. Then $\Pr(X) < \Pr_w(X)$ for some regular probability distribution Pr on <W, R> iff $R(w) \cap X \neq \{\}$ and $R(w) \cup X \neq W$.
Proof: Suppose that $R(w) \cap X \neq \{\}$ and $R(w) \cup X \neq W$. Let $x \in R(w) \cap X$ and $y \in W-(R(w) \cup X))$, so $x \neq y$. Thus $|W| = n \geq 2$. Define a probability distribution Pr over W thus:

If $n = 2$ $\Pr(\{x\}) = ½$ and $\Pr(\{y\}) = ½$.

If $n > 2$, $\Pr(\{x\}) = 1/3$, $\Pr(\{y\}) = ½$, and $PR(\{z\}) = 1/6(n-2)$ for $z \in W-\{x, y\}$.

Either way, since $X \subseteq W-\{y\}$, $\Pr(X) \leq \Pr(W-\{y\}) = ½$. Since $R(w) \subseteq W-\{y\}$, $\Pr(R(w)) \leq \Pr(W-\{y\}) = ½$. But $\Pr_w(X) = \Pr(X \cap R(w))/\Pr(R(w)) \geq \Pr(\{x\})/\Pr(R(w)) \geq (1/3)/(1/2) = 2/3 > ½ \geq \Pr(X)$.

For the converse, let Pr be any regular probability distribution on <W, R>.
First, suppose that $R(w) \cap X = \{\}$. Then $\Pr_w(X) = 0$ so $\Pr(X) \geq \Pr_w(X)$.
Second, suppose that $R(w) \cup X = W$. Then $W-R(w) \subseteq X$ so $\Pr(W-X \mid W-R(w)) = 0$. Hence, by total probability:
$\Pr(W-X) = \Pr(W-X \mid R(w))\Pr(R(w)) + \Pr(W-X \mid W-R(w))\Pr(W-R(w)) = \Pr(W-X \mid R(w))\Pr(R(w)) \leq \Pr(W-X \mid R(w)) = \Pr_w(W-X)$ so $\Pr(X) = 1 - \Pr(W-X) \geq 1 - \Pr_w(W-X) = \Pr_w(X)$, as required.

Notation: $R^{-1}(x) = \{w: Rwx\}$ and $N_R = \{w: R(w) \neq W\}$.

Comparative Transfer: $P_>[P_>[X]] \subseteq P_>[X]$

Proposition 6.0: Let <W, R> be a finite serial frame <W, R> and $x \in W$. Then:
$P_>[P_>[\{x\}]] \subseteq P_>[\{x\}]$ for every regular probability distribution Pr on <W, R> iff:
for all $w, x \in W$: if $R(w) \cap (R^{-1}(x) \cap N_R) \neq \{\}$ and $R(w) \cup (R^{-1}(x) \cap N_R) \neq W$ then $x \in R(w)$.
Proof: First note that for any regular probability distribution Pr and $x \in W$:

$P_>[\{x\}] = R^{-1}(x) \cap N_R$

For if $x \in R(w)$ and $R(w) \neq W$ then $w \in P_>[\{x\}]$ because $\Pr(R(w)) < 1$ so $\Pr(\{x\}) < \Pr(\{x\})/\Pr(R(w)) = \Pr_w(\{x\})$. If $R(w) = W$ then $w \notin P_>[\{x\}]$ because $\Pr(\{x\}) = \Pr_w(\{x\})$. If $x \notin R(w)$ then $w \notin P_>[\{x\}]$ because $\Pr_w(\{x\}) = 0$ and $\Pr(\{x\}) > 0$. Thus it suffices to prove:

$P_>[P_>[\{x\}]] \subseteq P_>[\{x\}]$ for every regular probability distribution Pr on <W, R> iff:

for all $w \in W$ if $R(w) \cap P_>[\{x\}] \neq \{\}$ and $R(w) \cup P_>[\{x\}] \neq W$ then $x \in R(w)$.

Suppose that $P_>[P_>[\{x\}]] \subseteq P_>[\{x\}]$ for every regular probability distribution Pr on <W, R>, and that $R(w) \cap P_>[\{x\}] \neq \{\}$ and $R(w) \cup P_>[\{x\}] \neq W$. By lemma 5, $\Pr(P_>[\{x\}]) < \Pr_w(P_>[\{x\}])$ for some regular probability distribution Pr on <W, R>, so $w \in P_>[P_>[\{x\}]]$, so $w \in P_>[\{x\}]$ by hypothesis, so $w \in R^{-1}(x) \cap N_R$, so $x \in R(w)$.

Conversely, suppose: whenever $R(w) \cap P_>[\{x\}] \neq \{\}$ and $R(w) \cup P_>[\{x\}] \neq W$, $x \in R(w)$. Let Pr be a regular probability distribution on <W, R>. Suppose that $w \in P_>[P_>[\{x\}]]$. Thus $\Pr(P_>[\{x\}]) < \Pr_w(P_>[\{x\}])$. Then, by 5.0, $R(w) \cap P_>[\{x\}] \neq \{\}$ and $R(w) \cup P_>[\{x\}] \neq W$, so by hypothesis $x \in R(w)$. Therefore $w \in R^{-1}(x)$, and $R(w) \neq W$, so $w \in N_R$, so $w \in P_>[\{x\}]$. Thus $P_>[P_>[\{x\}]] \subseteq P_>[\{x\}]$, as required.

Corollary 6.1: If Comparative Transfer is valid on a finite serial frame <W, R>, then for all *v*, *w*, *x*, *y*, *z*, if R*wy*, R*yx*, not R*yv*, not R*wz*, and either not R*zx* or for all *u* R*zu*, then R*wx*.

Corollary 6.2: If Comparative Transfer is valid on a finite serial frame <W, R>, then for all *x*, *y*, z, *v*, if neither R*xx* nor R*yv*, then if R*xy* and R*yz* then R*xz*.
Proof: Substitute *x* for *w* and *z,* and *z* for *x* in 6.1.

Proposition 7.0: Let <W, R> be a finite serial frame <W, R> and $x \in W$. Then:
$P_>[P_>[W-\{x\}]] \subseteq P_>[W-\{x\}]$ for every regular probability distribution Pr on <W, R> iff:
for all $w, x \in W$, if $x \in R(w)$ then either $R(w) \subseteq R^{-1}(x)$ or $R^{-1}(x) \subseteq R(w)$.
Proof: Note that $P_>[W-\{x\}] = W-R^{-1}(x)$. For if $w \in R^{-1}(x)$ and $R(w) \neq W$ then $w \notin P_>[W-\{x\}]$ because $Pr(\{x\}) < Pr_w(\{x\})$ (as in 6.0) so $Pr(W-\{x\}) > Pr_w(W-\{x\})$. If $R(w) = W$ then, again as in 6.0, $w \notin P_>[\{x\}]$. If $w \notin R^{-1}(x)$ then $w \in P_>[W-\{x\}]$ because $Pr_w(W-\{x\}) = 1$ and $Pr(W-\{x\}) < 1$.
Thus it suffices to prove:

> $P_>[P_>[W-\{x\}]] \subseteq P_>[W-\{x\}]$ for every regular probability distribution Pr on <W, R> iff:
> for all $w \in W$ if $R(w) \cap P_>[\{x\}] \neq \{\}$ and $R(w) \cup P_>[\{x\}] \neq W$ then $w \in W-R^{-1}(x)$.

The rest of the proof is similar to that of 6.0

Corollary 7.1: If Comparative Transfer is valid on a finite serial frame <W, R>, then for all *w*, *x*, if R*wx* then either for all *z* if R*wz* then R*zx* or for all *z* if R*zx* then R*wz*.

Corollary 7.2: If Comparative Transfer is valid on a finite serial frame <W, R>, then whenever R*wx*, either R*ww* or R*xx*.
Proof: By 7.1, substituting *x* for *z* in the first disjunct and *w* for *z* in the second.

Corollary 7.3: On any finite serial frame on which Comparative Transfer is valid, so is $(\Box p \supset p) \vee \Box(\Box q \supset q)$.
Proof: Suppose that $\Box p \supset p$ is false at $w \in W$ in a model based on such a frame <W, R>. Thus not R*ww*. Hence, by 7.2, R*xx* for all *x* such that R*wx*, so $\Box(\Box q \supset q)$ is true at *w* in the model.

Corollary 7.4: On any finite serial frame <W, R> on which Comparative Transfer is valid, for all *w*, *x*, if R*wx* then there is a *u* such that R*wu* and R*ux*
Proof: From 7.2

Corollary 7.5: On any finite serial frame on which Comparative Transfer is valid, so is $\Box\Box p \supset \Box p$.
Proof: From 7.4.

Corollary 7.6: On any finite serial frame <W, R> on which Comparative Transfer is valid, for all *w*, *x*, *y*, if R*wx* and R*wy* but not R*ww* then R*xy*.
Proof: By 7.1, since R*wy*, either for all *z* if R*wz* then R*zy* or for all *z* if R*zy* then R*wz*. The first disjunct implies that R*xy* because R*wx*. The second disjunct implies that R*ww* because R*wy*.

<u>Corollary 7.7</u>: On any finite serial frame on which Comparative Transfer is valid, so is $(\Box p \supset p) \vee (\Diamond\Box q \supset \Box q)$.
<u>Proof</u>: Suppose that $\Box p \supset p$ is false at some *w* in a model based on such a frame <W, R>. Then not R*ww*. Suppose that $\Diamond\Box q$ is true at *w* in the model, so $\Box q$ is true at some *x* such that R*wx*. Then whenever R*wy*, by 7.6 R*xy*, so *q* is true at *y*. Thus $\Box q$ is true at *w*.

<u>Corollary 7.8</u>: On any finite serial frame <W, R> on which Comparative Transfer is valid, for all *w*, *x*, *y*, *z*, if R*wx*, R*xy*, and R*yz*, then either R*xw*, or R*yx*, or R*zy*.
<u>Proof</u>: Suppose that R*wx*, R*xy*, and R*yz*. By 7.2, since R*xy*, either R*xx* or R*yy*. Suppose that R*xx*. So by 7.1 either for all *u* if R*xu* then R*ux* or for all *u* if R*ux* then R*xu*. Hence either if R*xy* then R*yx* or if R*wx* then R*xw*; thus either R*xw* or R*yx*. Similarly, if R*yy* then either R*yx* or R*zy*. Hence either R*xw*, or R*yx*, or R*zy*.

<u>Corollary 7.9</u>: On any finite serial frame on which Comparative Transfer is valid, so is $p \supset \Box(q \supset (\Diamond p \vee \Box(r \supset (\Diamond q \vee \Box\Diamond r))))$.
<u>Proof</u>: Suppose that the formula is false at some *w* in some model based on the frame. Then *p* is true and $\Box(q \supset (\Diamond p \vee \Box(r \supset (\Diamond q \vee \Box\Diamond r)))$ at *w*. Hence $q \supset (\Diamond p \vee \Box(r \supset (\Diamond q \vee \Box\Diamond r))$ is false at some *x* such that R*wx*. Hence *q* is true and both $\Diamond p$ and $\Box(r \supset (\Diamond q \vee \Box\Diamond r))$ false at *x*. The former implies that not R*xw*. The latter implies that $r \rightarrow (\Diamond q \vee \Box\Diamond r)$ is false at some *y* such that R*xy*, so *r* is true and both $\Diamond q$ and $\Box\Diamond r$ are false at *y*. The former implies that not R*yx*. The latter implies that $\Diamond r$ is false at some *z* such that R*yz*. Hence not R*zy*. That is impossible by 7.8.

<u>Corollary 7.10</u>: On any finite reflexive frame <W, R> on which Comparative Transfer is valid, for all *w*, *x*, *y*, if R*wx* and R*xy*, then either R*xw* or R*yx*.
<u>Proof</u>: A simplification of the proof of 7.8.

<u>Corollary 7.11</u>: On any finite reflexive frame on which Comparative Transfer is valid, so is $p \supset \Box(q \supset (\Diamond p \vee \Box\Diamond q))$.
<u>Proof</u>: A simplification of the proof of 7.9.

<u>Corollary 7.12</u>: Suppose that Comparative Transfer is valid on a finite serial frame <W, R>. Then if R*wx* but not R*ww*, either R(*x*) = W or R(*x*) = R(*w*).
<u>Proof</u>: From 6.2 and 7.6.

<u>Corollary 7.13</u>: If a serial frame <W, R> with $|W| \leq 3$ satisfies the necessary conditions in 6.0 and 7.0 for Comparative Transfer to be valid, then Comparative Transfer is valid on <W, R>.
<u>Proof</u>: Suppose that $X \subseteq W$. If $|X| = 1$ then by 6.0 $P_>[P_>[X]] \subseteq P_>[X]$; if $|X| = 2$ then by 7.0 $P_>[P_>[X]] \subseteq P_>[X]$. If $|X| = 0$ or $|X| = 3$ then $P_>[P_>[X]] = \{\}$ so the case is trivial.

<u>Proposition 8.0</u>: If Comparative Transfer is valid on a finite serial frame <W, R>, then for all *w*, *x*, *y*, *z*, if R*wx*, R*xy*, and R*yz*, then for some *u*, R*wu* and R*uz*.
<u>Proof</u>: It suffices to derive a contradiction from the hypothesis that R*wx*, R*xy*, R*yz*, and there is no *u* such that R*wu* and R*uz*. First note that there are no shortcuts in the chain from *w* to

z, in the sense that not Rwy, not Rxz, and not Rwz. The first two are easy because we can put $u = y$ and $u = x$ respectively; the third follows from 7.4. But by an instance of 6.1:

if Rwx, Rxy, not Rxz, not Rwz and not Rzy then Rwy.

Thus, since Rwx, Rxy, not Rwz, and not Rwy, it follows that Rzy. By another instance of 6.1:

for all v, if Rxy, Ryz, not Ryv, not Rxw, and not Rwz, then Rxz.

Thus, since Rxy, Ryz, not Rwz, and not Rxz, this follows:

(*) either Rxw or Ryv for all v.

But by an instance of 7.1:

if Rxy then either for all u if Rxu then Ruy or for all u if Ruy then Rxu.

Thus, since Rxy, either if Rxw then Rwy or if Rzy then Rxz. But on the initial hypothesis we already have that Rzy, not Rxz, and not Rwy. Thus not Rxw. Hence, by (*), Ryv for all v. But by yet another instance of 7.1:

if Rwx, Rxy, not Rxz, not Rwy, and for all v Ryv, then Rwy.

This is a final contradiction, since on the initial hypothesis we already have that Rwx, Rxy, not Rxz, not Rwy, and for all v Ryv.

<u>Corollary 8.1</u>: If Comparative Transfer is valid on a finite serial frame, then so too is $\Box\Box p \equiv \Box\Box\Box p$.

<u>Proof</u>: From 7.5 and 8.0 by standard model theory of modal logic.

<u>Proposition 9.0</u>: If $\langle W, R\rangle$ is a finite serial frame such that whenever $w \in W$, $x \in R(w)$, either $R(w) = W$ or $R(x) = W$ or $R(x) = R(w)$, then Comparative Transfer is valid on $\langle W, R\rangle$.

<u>Proof</u>: If $R(w) = W$ then for all $Y \subseteq W$ $Pr_w(Y) = Pr(Y)$ so $w \notin P_>[Y]$.

Suppose that $w \notin P_>[X]$ and $w \in P_>[P_>[X]]$. Thus $R(w) \neq W$ (let $Y = P_>[X]$).

Suppose that for all $x \in R(w)$ either $R(x) = W$ or $R(x) = R(w)$.

If $R(x) = W$, then $x \notin P_>[X]$. If $R(x) = R(w)$ then $Pr_x(X) = Pr_w(X)$ so $x \notin P_>[X]$ because $w \notin P_>[X]$.

Thus for all $x \in R(w)$, $x \notin P_>[X]$. Hence $Pr_w(P_>[X]) = 0$, so $w \notin P_>[P_>[X]]$, contrary to hypothesis.

<u>Corollary 9.1</u>: On some finite serial frames, Comparative Transfer is valid but none of $\Box p \supset \Box\Box p$, $\Diamond p \supset \Box\Diamond p$, $p \supset \Box\Diamond p$, and $\Box p \supset p$ is valid.

<u>Proof</u>: Consider the three-world frame where $W = \{a, b, c\}$, $R(a) = \{b\}$, $R(b) = W$, and $R(c) = \{c\}$. By 9.0, Comparative Transfer is valid on $\langle W, R\rangle$. R is non-transitive, because Rab and Rbc but not Rac, so $\Box p \supset \Box\Box p$ is not valid. R is non-euclidean because Rba and Rbc but not Rac, so $\Diamond p \supset \Box\Diamond p$ is not valid. R is non-symmetric because Rbc but not Rcb, so $p \supset \Box\Diamond p$ is not valid. R is non-reflexive because not Raa, so $\Box p \supset p$ is not valid.

<u>Proposition 10.0</u>: Let $\langle W, R\rangle$ be a finite serial frame such that for some $x, z \in W$, $x \neq z$, $x \in R(x)$ and for all $y \in W$, if $y \neq x$ then $R(y) = \{z\}$. Then Comparative Transfer is valid on $\langle W, R\rangle$.

<u>Proof</u>: Let $\langle W, R\rangle$, x, z be as specified. Suppose that $y \notin P_>[X]$ for some $y \in W$ and $X \subseteq W$. If $y \neq x$ then $R(y) = \{z\} = R(z)$, so $y \notin P_>[P_>[X]]$ (as in the proof of 4.0). Thus it suffices to show that if $x \notin P_>[X]$ then $x \notin P_>[P_>[X]]$. There are two cases.

Case (i): $z \in X$. We may assume that $X \neq W$, otherwise $P_>[X] = \{\}$ so $P_>[P_>[X]] = \{\}$ and we are done. Hence $Pr(X) < 1$. Now if $y \neq x$, $R(y) = \{z\} \subseteq X$ so $Pr_y\{X\} = 1$, so $y \in P_>[X]$. Since $x \in R(x)$, $W - R(x) \subseteq P_>[X]$, so $Pr(P_>[X] \mid W - R(x)) = 1$. Hence, by the principle of total probability:

$$\begin{aligned}\Pr(P_>[X]) &= \Pr(P_>[X] \mid R(x))\Pr(R(x)) + \Pr(P_>[X] \mid W{-}R(x))\Pr(W{-}R(x))\\ &= \Pr_x(P_>[X])\Pr(R(x)) + \Pr(W{-}R(x))\\ &\geq \Pr_x(P_>[X])\Pr(R(x)) + \Pr_x(P_>[X])\Pr(W{-}R(x)) = \Pr_x(P_>[X])\end{aligned}$$

Hence $x \notin P_>[P_>[X]]$, as required.
Case (ii) $z \notin X$. Hence if $y \neq x$, $\Pr_y[X] = 0$, so $y \notin P_>[X]$. Thus if $x \notin P_>[X]$, then $\Pr_x(P_>[X]) = 0$, hence $x \notin P_>[P_>[X]]$ and we are done.

<u>Proposition 11.0</u>: Let <W, R> be a finite serial frame, with some $z \in W$ such that for all $x \in W$, $R(x) \subseteq \{x, z\}$. Then Comparative Transfer is valid on <W, R>.
<u>Proof</u>: Let <W, R>, z be as specified, and $x \notin P_>[X]$. It suffices to show that $x \notin P_>[P_>[X]]$. There are two cases.
Case (i): $z \in X$. As in the proof of 10.0, we may assume that $X \neq W$, so $\Pr(X) < 1$. Observe that $X \subseteq P_>[X]$, for if $y \in X$ then $R(y) \subseteq \{y, z\} \subseteq X$, so $\Pr_y(X) = 1 > P_>[X]$, so $y \in P_>[X]$. Thus:

$$\Pr(X) \leq \Pr(P_>[X])$$

Since $x \notin P_>[X]$ and $z \in X$, $P_>[X] \cap R(x) \subseteq X \cap R(x)$, so:

$$\Pr_x(P_>[X]) \leq \Pr_x(X)$$

Since $x \notin P_>[X]$:

$$\Pr_x(X) \leq \Pr(X)$$

Thus $\Pr_x(P_>[X]) \leq \Pr_x(X) \leq \Pr(X) \leq \Pr(P_>[X])$, so $x \notin P_>[P_>[X]]$, as required.
Case (ii): $z \notin X$. Since $R(z) = \{z\}$, $\Pr_z[X] = 0$, so $z \notin P_>[X]$. Hence if if $x \notin P_>[X]$ then $P_>[X] \cap R(x) = \{\}$, so $\Pr_x(P_>[X]) = 0$, so $x \notin P_>[P_>[X]]$, and we are done.

<u>Proposition 12.0</u>: The condition on R for Comparative Transfer to be valid on a finite serial frame <W, R> is not expressible in first-order logic without identity.
<u>Proof</u>: Let $W = \{0, 1, 2\}$, $R(0) = \{0, 1\}$, $R(1) = \{1\}$, $R(2) = \{2\}$, $W^* = \{0, 1, 2, 3\}$, f be the function from W^* onto W such that $f0 = 0$, $f1 = 1$, $f2 = f3 = 2$, and for all j, k in W^*, R^*jk iff $Rfjfk$. Thus 2 and 3 are indiscernible in terms of R^*. By standard non-modal model theory, the models <W, R> and <W*, R*> satisfy the same formulas in the language of first-order logic with just one predicate letter and without identity, which is binary. Since for all j in W, $R(j) \subseteq \{j, 2\}$, by 11.0 Comparative Validity is valid on <W, R>. But Comparative Transfer is not valid on <W*, R*>. For let Pr be the uniform probability distribution over W^*, and $X = \{1, 3\}$, so $\Pr(X) = ½$. Now $\Pr_0(X) = \Pr_2(X) = \Pr_3(X) = ½$, $\Pr_1(X) = 1$, so $\Pr_>[X] = \{1\}$, so $\Pr(\Pr_>[X]) = ¼$. Consequently, $\Pr_0(\Pr_>[X]) = ½$, so $0 \in \Pr_>[\Pr_>[X]]$, but $0 \notin \Pr_>[X]$. Since <W*, R*> differs from <W, R> with respect to the validity of Comparative Transfer but not with respect to which formulas of first-order logic without identity are satisfied, the former cannot be expressed in terms of the latter.

<u>Corollary 12.1</u>: The necessary conditions in 6.0 and 7.0 for Comparative Transfer to be valid on a finite serial frame are not jointly sufficient.
<u>Proof</u>: By 6.1 and 7.1, those necessary conditions are expressible in first-order logic without identity, so by 12.0 they are not jointly sufficient.

<u>Proposition 13.0</u>: There is a finite serial frames on which Comparative Transfer is valid such that Comparative Transfer is not valid on the union of two disjoint copies of that frame.

Proof: Let W = {0, 1}, R(0) = {0, 1}, R(1) = {1} W* = {2, 3}, R*(2) = {2, 3}, R*(3) = {3}; thus <W, R> and <W*, R*> are isomorphic. By 9.0, Comparative Transfer is valid on <W, R> and <W*, R*>. But it is not valid on <W $\cup$ W*, R $\cup$ R*>. For let Pr be the uniform probability distribution over W and X = {1, 2}, so Pr(X) = ½. Thus Pr_0(X) = Pr_2(X) = ½, Pr_1(X) = 1, and Pr_3(X) = 0, so $Pr_>$[X] = {1}, so Pr($Pr_>$[X]) = ¼. Thus Pr_0($Pr_>$[X]) = ½, so $0 \in Pr_>[Pr_>[X]]$, but $0 \notin Pr_>[X]$.

Proposition 14.0: For $a \in (0, 1)$, the validity of Threshold$_a$ Transfer and the validity of Comparative Transfer are independent conditions on a frame.
Proof: The union frame in the proof of 13.0, on which Comparative Transfer is invalid, meets the condition for the validity of Threshold$_a$ Transfer. Conversely, the latter condition entails transitivity, so Threshold$_a$ Transfer is invalid on those non-transitive frames on which Comparative Transfer is valid (see 9.1).

We make the following definitions for a frame <W, R> ($\subset$ expresses proper subsethood):
<W, R> is *quasi-nested* iff for all $w, x, y, z \in$ W, if R*wx*, R*wy*, R*xz*, and R*yz*, then either R*xy* or R*yx*.
$R_0(w) = R(w) \cap \{u: R(u) = R(w)\}$
$R_1(w) = R(w) \cap \{u: R(u) \subset R(w)$ and for no $v \in R(w)$: $R(u) \subset R(v) \subset R(w)\}$

Observation 15.0: A quasi-nested serial frame is quasi-reflexive.
Proof: Put $x = y$ in the definition of 'quasi-nested'.

Proposition 15.1: Let <W, R> be a finite serial transitive quasi-nested frame. If $x, y \in R_1(w)$ then either R(x) = R(y) or $R(x) \cap R(y) = \{\}$.
Proof: Suppose that $R(x) \cap R(y) \neq \{\}$. Hence, for some z, R*xz* and R*yz*. Since $R_1(w) \subseteq R(w)$, R*wx* and R*wy*. Thus, since R is quasi-nested, either R*xy* or R*yx*. Without loss of generality, suppose that R*xy*. Thus, since R is transitive, $R(y) \subseteq R(x)$. Suppose that $R(x) \neq R(y)$. Hence $R(y) \subset R(x)$. But $x \in R_1(w)$, so $R(x) \subset R(w)$. Thus $R(y) \subset R(x) \subset R(w)$, contradicting the hypothesis that $y \in R_1(w)$. Hence R(x) = R(y).

Proposition 15.2: Let <W, R> be a finite serial transitive quasi-nested frame, $w \in$ W, and J consist of one member from each equivalence class of the equivalence relation defined by R(x) = R(y) over $R_1(w)$. Then $R_0(w)$ and R(j) for $j \in$ J partition R(w). Moreover, if j and k are distinct members of J then $R(j) \cap R(k) = R_0(w) \cap R(j) = \{\}$.
Proof: By definition, $R_0(w) \subseteq R(w)$ and $R(j) \subset R(w)$ for $j \in J \subseteq R_1(w)$, so $R_0(w) \cup (\cup_{j \in J} R(j)) \subseteq R(w)$. For the converse, let $x \in R(w)$. Since R is transitive, $R(x) \subseteq R(w)$. If R(x) = R(w) then $x \in R_0(w)$ and we are done, so suppose that $R(x) \subset R(w)$. Since W is finite, there is a sequence x_0, x_1, …, x_n all in R(w) of maximal length such that $x_0 = x$ and for $0 \leq i < n$, $R(x_i) \subset R(x_{i+1}) \subset R(w)$. So there is no $v \in R(w)$ such that $R(x_n) \subset R(v) \subset R(w)$, otherwise we could put $v = x_{n+1}$. Thus $x_n \in R_1(w)$, so for some $k \in$ J R(k) = R(x_n). Since $x \in R(w)$ and by 15.0 R is quasi-reflexive, $x \in R(x) = R(x_0) \subseteq R(x_n) = R(k)$, so $x \in \cup_{j \in J} R(j)$. Thus $R(w) = R_0(w) \cup (\cup_{j \in J} R(j))$. Moreover, if j and k are distinct members of J then $R(j) \neq R(k)$, so $R(j) \cap R(k) = \{\}$ by 15.1, and if $x \in R(j)$ then by transitivity $R(x) \subseteq R(j) \subset R(w)$, so $x \notin R_0(w)$.

<u>Proposition 15.3</u>: Let <W, R> be a finite serial transitive quasi-nested frame, $b \in [0, 1]$, $X \subseteq W$, $w \in W$, and Pr a regular probability distribution on <W, R>. Then $b\text{Pr}_w(\text{P}_{\geq b}[X]) \leq \text{Pr}_w(X)$.

<u>Proof</u>: We use induction on $|R(w)|$. The base case is vacuous because R is serial. Suppose that for all $x \in W$ such that $|R(x)| < |R(w)|$, $b\text{Pr}_x(\text{P}_{\geq b}[X]) \leq \text{Pr}_x(X)$. We have two cases to consider.

Case (i): $\text{Pr}_w(X) \geq b$. Hence $b\text{Pr}_w(\text{P}_{\geq b}[X]) \leq b \leq \text{Pr}_w(X)$, as required.

Case (ii): $\text{Pr}_w(X) < b$. Now for all $u \in R_0(w)$, $R(u) = R(w)$, so $\text{Pr}_u(X) = \text{Pr}_w(X) < b$, so $R_0(w) \cap \text{P}_{\geq b}[X] = \{\}$, so $\text{Pr}(\text{P}_{\geq b}[X]|R_0(w)) = 0$. By 15.2, for any $Y \subseteq W$:

(*) $\text{Pr}_w(Y) = \text{Pr}(Y|R(w)) = \text{Pr}(Y|R_0(w) \cup (\cup_{j \in J} R(j)))$
$= \text{Pr}(Y|R_0(w))\text{Pr}_w(R_0(w)) + \sum_{j \in J} \text{Pr}(Y|R(j))\text{Pr}_w(R(j))$
$= \text{Pr}(Y|R_0(w))\text{Pr}_w(R_0(w)) + \sum_{j \in J} \text{Pr}_j(Y)\text{Pr}_w(R(j))$

By (*) for $Y = \text{Pr}_{\geq b}[X]$:

$\text{Pr}_w(\text{P}_{\geq b}[X]) = \text{Pr}(\text{Pr}_{\geq b}[X]|R_0(w))\text{Pr}_w(R_0(w)) + \sum_{j \in J} \text{Pr}_j(\text{Pr}_{\geq b}[X])\text{Pr}_w(R(j))$
$= \sum_{j \in J} \text{Pr}_j(\text{Pr}_{\geq b}[X])\text{Pr}_w(R(j))$

But if $j \in J$ then $R(j) \subset R(w)$, so $|R(j)| < |R(w)|$ because W is finite, so by the induction hypothesis $b\text{Pr}_j(\text{P}_{\geq b}[X]) \leq \text{Pr}_j(X)$. Thus:

$b\text{Pr}_w(\text{P}_{\geq b}[X]) = \sum_{j \in J} b\text{Pr}_j(\text{Pr}_{\geq b}[X])\text{Pr}_w(R(j))$
$\leq \sum_{j \in J} \text{Pr}_j(X)\text{Pr}_w(R(j)))$
$\leq (\text{Pr}(X|R_0(w))\text{Pr}_w(R_0(w)) + \sum_{j \in J} \text{Pr}_j(X)\text{Pr}_w(R(j)))$
$= \text{Pr}_w(X)$

by (*) for Y = X, as required. This completes the induction.

<u>Corollary 15.4</u>: Let <W, R> be a finite serial transitive quasi-nested frame, $a, b \in [0, 1]$, $X \subseteq W$, and Pr a regular probability distribution on <W, R>. Then $\text{Pr}_{\geq a}[\text{Pr}_{\geq b}[X]] \subseteq \text{Pr}_{\geq ab}[X]]$.

Proof: For $w \in \text{Pr}_{\geq a}[\text{Pr}_{\geq b}[X]]$, $ab \leq b\text{Pr}_w(\text{Pr}_{\geq b}[X]) \leq \text{Pr}_w(X)$ by 15.3, so $w \in \text{Pr}_{\geq ab}[X]$, as required.

<u>Proposition 15.5</u>. The principle $\text{Pr}_{\geq a}[\text{Pr}_{\geq b}[X]] \subseteq \text{Pr}_{\geq ab}[X]]$ is valid on a finite serial frame for all $a, b \in [0, 1]$ if and only if the frame is transitive and quasi-nested.

<u>Proof</u>: Given 15.4, we need only prove the left-to-right direction. Suppose that $\text{Pr}_{\geq a}[\text{Pr}_{\geq b}[X]] \subseteq \text{Pr}_{\geq ab}[X]]$ is valid on a finite serial frame <W, R> for all $a, b \in [0, 1]$.

First, note that <W, R> is transitive by Proposition 1.0 for $a = b > 0$.

Second, we show that <W, R> is quasi-reflexive. For suppose that Rwx but not Rxx, so $w \neq x$. Let $|W| = n \geq 2$. Define a probability distribution Pr over W thus:

$\text{Pr}(\{x\}) = 2/3$

$\text{Pr}(\{y\}) = 1/3(n-1)$ for $y \neq x$

Then $\text{Pr}_x(W - \{x\}) = 1$, so $x \in \text{Pr}_{\geq 1}[W - \{x\}]$, so $\text{Pr}_w(\text{Pr}_{\geq 1}[W - \{x\}]) \geq \text{Pr}_w(\{x\}) \geq 2/3$, so $w \in \text{Pr}_{\geq 2/3}[\text{Pr}_{\geq 1}[W - \{x\}]]$, but $\text{Pr}_w(W - \{x\}) = 1 - \text{Pr}_w(\{x\}) \leq 1 - \text{Pr}(\{x\}) = 1/3$, so $w \notin \text{Pr}_{\geq 2/3}(W - \{x\})$.

Finally, we show that <W, R> is quasi-nested. For suppose otherwise. Then for some $w, x, y, z \in W$, if Rwx, Rwy, Rxz, Ryz, but neither Rxy nor Ryx; thus $w \neq x$, $w \neq y$, $x \neq z$, and $y \neq z$. Since R is quasi-reflexive, Rxx, Ryy, and Rzz; thus $x \neq y$. Since R is transitive, Rwz but neither Rzx (otherwise Ryx) nor Rzy (otherwise Rxy); thus $w \neq z$. Hence $|W| = n \geq 4$. Define a probability distribution Pr over W thus:

$\Pr(\{x\}) = \Pr(\{y\}) = 6/25$
$\Pr(\{z\}) = 12/25$
$\Pr(\{u\}) = 1/25(n - 3)$ for $u \notin \{x, y, z\}$

Since $\{z\} \subseteq R(z) \subseteq W - \{x, y\}$, $\Pr_z(\{z\}) \geq 12/13 > 12/19$. Since $\{x, z\} \subseteq R(x) \subseteq W - \{y\}$:
$\Pr_x(\{z\}) = \Pr(\{z\})/\Pr(R(x)) \geq \Pr(\{z\})/\Pr(W - \{y\}) = 12/19$
Similarly, $\Pr_y(\{z\}) \geq 12/19$. Hence $\{x, y, z\} \subseteq P_{\geq 12/19}[\{z\}]$. Thus:
$\Pr_w(P_{\geq 12/19}[\{z\}]) \geq \Pr_w(\{x, y, z\}) \geq \Pr(\{x, y, z\}) = 24/25$. Hence $w \in P_{\geq 24/25}[P_{\geq 12/19}[\{z\}]]$. But $\{x, y, z\} \subseteq R(w)$, so $\Pr_w(\{z\}) = \Pr(\{z\})/\Pr(R(w) \leq \Pr(\{z\})/\Pr(\{x, y, z\}) = ½ < (24/25)(12/19) = 288/475$, so $w \notin \Pr_{\geq 288/475}(\{z\})$. Thus the principle $\Pr_{\geq 24/25}[\Pr_{\geq 12/19}[X]] \subseteq \Pr_{\geq 288/475}[X]]$ is not valid on the frame, contrary to hypothesis.

Notes

1 There is a subtlety here, for if E=K is true, why should it follow that the agent knows E=K? If one is an agent who doesn't know E=K, couldn't one know that one knows X without knowing that X is part of one's evidence, or *vice versa*? In practice, however, the putative counterexamples to positive introspection for knowledge do not depend on whether the agent knows E=K, so one can stipulate that the agent does know it. Thus positive introspection principles for knowledge and for evidence stand or fall together, given E=K, even though, if they fall together, the nonempty class of counterexamples to one differs slightly from the nonempty class of counterexamples to the other. Epistemic logic in any case tends not to treat such subtle differences in mode of presentation as differences in what is known, since it treats knowledge as closed under truth-functional entailment.

2 In this context a normal modal logic, identified with the set of its theorems, is a set of formulas of a standard language for monomodal propositional logic (with $\Box$ as the only primitive modal operator) containing all truth-functional tautologies and the K axiom $\Box(p \supset q) \supset (\Box p \supset \Box q)$ and closed under uniform substitution, modus ponens, and the rule of necessitation. The correspondence result can be proved by the standard method of canonical models (see e.g. Hughes and Cresswell 1984: 22-7).

3 To derive quasi-truth from negative introspection, consider two cases: (i) p truth-functionally entails $\Box p \supset p$, so by normality $\Box p$ entails $\Box(\Box p \supset p)$; (ii) $\neg\Box p$ truth-functionally entails $\Box p \supset p$, so by normality $\Box\neg\Box p$ entails $\Box(\Box p \supset p)$; but by negative introspection and normality, $\neg\Box p$ entails $\Box\neg\Box p$.

4 Such frames also have the convergence property that if Rwx and Rwy then for some z both Rxz and Ryz, which corresponds to the modal axiom $\Diamond\Box p \supset \Box\Diamond p$. Since they are reflexive and transitive too, they validate the modal system proposed as a logic of knowledge in Stalnaker 2006.

5 One objection to Positive Introspection for knowledge is that a simple creature with no grasp of the distinction between knowing and not knowing may still know truths about its environment, but without being able to so much as entertain the truth that it knows them. This objection does not generalize against the principle that if you know that you know, then you know that you know that you know, for if you know that you know then in the required sense you do grasp the distinction between knowing and not knowing. However, such objections are too fine-grained for the present setting, where logical omniscience is assumed: if you know one truth but cannot grasp another you presumably cannot grasp their disjunction either, which follows from the truth you do know. For present purposes we are accepting the principle $\Box p \supset \Box(p \vee q)$, which is valid in all standard epistemic models.

6 Compare the models in the appendix to Salow 2015, used for a related purpose. The subsets of any given set ordered by the subset relation also have the convergence property. As a result, the frame in the text validates the logic S4.2; see footnote 4.

7 An earlier version of this material was presented at Oxford and a British Academy-funded conference on epistemology in Cambridge; I thank participants for useful suggestions, and Bernhard Salow, Kevin Dorst, and John Hawthorne for valuable discussion and correspondence.

References

Blackburn, Patrick, Maarten de Rijke, and Yde Venema 2001: *Modal Logic*, Cambridge: Cambridge University Press.

Briggs, Rachael. 2009: 'Distorted reflection', *Philosophical Review*, 118: 59-85.

Dorst, Kevin. 2016: 'Evidence: a guide for the uncertain', MS.

Feldman, Richard. 2007: 'Reasonable religious disagreements', in Louise Antony (ed.), *Philosophers Without God: Meditations on Atheism and the Secular Life*: 194-214. Oxford: Oxford University Press.

Feldman, Richard 2014: 'Evidence of evidence is evidence', in Jonathan Matheson and Rico Vitz (eds.), *The Ethics of Belief*: 284-300. Oxford: Oxford University Press.

Fitelson, Branden. 2012: 'Evidence of evidence is not (necessarily) evidence', *Analysis*, 72: 85-88.

Geanakoplos, John. 1989. 'Game theory without partitions, and applications to speculation and consensus', *Cowles Foundation Discussion Paper*, 914, Yale University.

Greco, Daniel. 2014: 'Could KK be OK?', *Journal of Philosophy*, 111: 169-197.

Hawthorne, John, and Ofra Magidor. 2009. 'Assertion, context, and epistemic accessibility', *Mind*, 118: 377-397.

Hawthorne, John, and Ofra Magidor. 2011. 'Assertion and epistemic opacity', 119: 1087-1105.

Hintikka, Jaakko. 1962: *Knowledge and Belief*. Ithaca, N.Y.: Cornell University Press.

Hughes, George, and Max Cresswell. 1984: *A Companion to Modal Logic*. London: Methuen.

Roche, William. 2014: 'Evidence of evidence is evidence under screening-off', *Episteme*, 11: 119-124.

Salow, Bernhard. 2015: 'Elusive externalism', MS.

Stalnaker, Robert. 2006: 'On logics of knowledge and belief', *Philosophical Studies*, 128: 169-199.

Stalnaker, Robert. 2009: 'On Hawthorne and Magidor on assertion, context, and epistemic accessibility', *Mind*, 118: 399-409.

Tal, Eyal, and Juan Comesaña. 201X: 'Is evidence of evidence evidence?', *Noûs*, forthcoming.
Weisberg, Jonathan. 2007: 'Conditionalization, reflection, and self-knowledge', *Philosophical Studies*, 135: 179-197.

White, Roger. 2006: 'Problems for dogmatism', *Philosophical Studies*, 131: 525-557.

Williamson, Timothy. 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy. 2014: 'Very improbable knowing', *Erkenntnis*, 79: 971-999.