Abstract: Deduction is important to scientific inquiry because it can extend knowledge efficiently, bypassing the need to investigate everything directly. The existence of closure failure—where one knows the premises and that the premises imply the conclusion but nevertheless does not know the conclusion—is a problem because it threatens this usage. It means that we cannot trust deduction for gaining new knowledge unless we can identify such cases ahead of time so as to avoid them. For philosophically engineered examples we have "inner alarm bells" to detect closure failure, but in scientific investigation we would want to use deduction for extension of our knowledge to matters we don't already know that we couldn't know. Through a quantitative treatment of how fast probabilistic sensitivity is lost over steps of deduction, I identify a condition that guarantees that the growth of potential error will be gradual; thus, dramatic closure failure is avoided. Whether the condition is fulfilled is often obvious, but sometimes it requires substantive investigation. I illustrate that not only safe deduction but the discovery of dramatic closure failures can lead to scientific advances.

Introduction

A great advantage of deductive reasoning is that it can extend our knowledge of a subject matter without our doing further substantive investigation, which makes inquiry more efficient, and sometimes even safer. Imagine a species with exactly three natural predators – tigers, hawks, and snakes – and everyone in a particular group of this species knows this. One of the group cries out as he is being killed so everyone knows that there is a predator around. Another of the group, who sits high in the trees, signals that there is no hawk. Another in the trees with a good view of the ground announces that it is not a tiger. If the members of the group can reason, then without another call, without taking chances stepping on the ground or inspecting every cave, without waiting to see the predator, everyone knows it is a snake.[1]

Reasoning is a knowledge-multiplier. However, if knowledge is not closed under known implication then this advantage is not present in all cases; valid deduction can take us from a belief that has the knowledge property to a belief that lacks the property. In some proposed cases of closure failure, this loss can even happen in one step. That a castle has been there for at least a hundred years implies that the world did not come into existence full-blown five minutes ago along with all the evidence that it had been around for a much longer time. However, the evidential traces that one would have thought give me knowledge of the premise proposition do not protect me at all from the possible falsehood of the inferred conclusion. That evidence would have been there even if the castle was five minutes old.

There have been efforts to characterize cases of closure failure in a principled way, but this has generally been seen as of theoretical rather than practical significance. Clearly if we are not able to identify cases of closure failure then it will be hazardous in every case to trust reasoning in the knowledge-multiplying role. However one might presume that in practice we will not easily be misled; we will know these cases when we see them because they raise "inner alarm bells." (Hawthorne 2005: 33, 36) We should only be confident of that, though, in cases that are philosophically engineered for the job of being unknown, such as propositions involving young castles incognito or cleverly disguised mules whose pretense is undetectable, or at least not detectable by the means that the imagined subject has used in coming to knowledge of her premise. In science the whole point is to find out things we are not already familiar

[1] This example is drawn from the Vervet monkey, which Brian Skyrms (2010, 22-24) uses to argue that it is possible for logic to have evolved via signaling behavior, with this case displaying a primitive version of the disjunctive syllogism.

with, but also would not be sure a priori we could not know; in some cases whether we can know them is itself a substantive question. Thus, unless we have a principled way of identifying the problematic cases, our epistemology has the implication that scientists should not have a blanket trust in results gained indirectly by deduction from something already known. If they do, then as far as anyone can tell they are always rolling the dice.

From this point of view, the interest of the question whether or to what extent knowledge is closed does not lie in matters concerning philosophical skepticism, such as whether there is a principled way of embracing the view that knowing that we have hands is not enough to give us knowledge that we are not brains in vats, while holding on to the view that deduction preserves the knowledge property in other cases. The interest of closure here lies in whether we can find a criterion to characterize a domain in which we can trust deduction wholesale, so that we can confidently pursue knowledge indirectly, in science and ordinary life.

Since in familiar cases of apparent closure failure deduction fails to preserve the knowledge property, it would be natural to think that to make the world safe for pursuing knowledge indirectly we should identify the distinction between cases where closure fails and cases where it does not. I will argue that this is not the distinction that identifies where deduction can be trusted as a knowledge-multiplier, by showing that there is a large domain of cases of closure failure that are not hazardous for the pursuit of knowledge indirectly by deduction. The aspect of familiar cases of apparent closure failure that threatens the pursuit of indirect knowledge is not that the conclusion belief cannot be labelled "knowledge," but more specifically that in these cases deduction leads to a conclusion belief with very large potential error. However the fact that it is bad to introduce large amounts of potential error does not mean that in order to trust deduction we need it to produce no further potential for error at all. If we are fallibilists then we accept that a small amount of potential error will attend most beliefs we call "knowledge," so we can hardly object to a little more, as long as we can be sure that it does not go beyond our tolerance. Thus the domain we need to identify is those cases where there are predictable bounds on the growth of potential error over steps of deduction. Such a domain can be defined, as I show below, and it includes an infinite number of cases where closure fails but the potential error the step of deduction introduces makes the potential error in the conclusion belief only slightly higher than one's requirements for knowledge will tolerate. In such a case deduction has taken us a long way toward the goal of knowledge, and someone with a slightly lower threshold would count it as knowledge.

Whether deductive inference extends our knowledge in a given case depends on whether the potential error it introduces is large or small. Thus characterizing the domain in which deduction is a knowledge multiplier requires a quantitative analysis. Here I argue that if we take failure of probabilistic sensitivity, defined below, as expressing the kind of error we want to avoid we can see three things. First, there is a general condition sufficient to rule out the cases in which deduction introduces maximal potential error. Second, the domain of cases that fulfill the condition can be shown to be well-behaved in the following sense. We can define degrees of loss of sensitivity over steps of deduction, and derive upper bounds for the accompanying growth of potential error over steps, that are fully predictable, quantitatively intuitive, and tolerable for science and life. These bounds can be derived from the strength of the knowledge of the premise belief and of one's reliance on that belief. Finally, we will see that maximal closure failure – where maximal potential error is reached in one step – is not only a phenomenon for scientists to avoid. Its possibility can raise questions of interest to the empirical science itself, and whether or not a case is one of closure failure can be far from obvious.

Why Think Knowledge is not Closed Under Known Implication?

Truth is preserved by logical implication – in other words the set of truths is closed under logical implication – but this does not settle the issue of whether knowledge is closed, because knowledge requires more than mere true belief. There are both intuitive and theoretical reasons for thinking that knowledge has cases of closure failure in which that extra part beyond true belief is not necessarily preserved by deduction. In the simplest kind of case, a single-premise, single-step deduction, the subject knows the premise, q, knows that the conclusion, p, follows from the premise, and bases her belief in the conclusion on her belief in the premise, but nevertheless appears not to know the conclusion.[2] In a common example,

q: That is a zebra
p: That is not a cleverly disguised mule

A subject, *S*, with normal eyesight and basic knowledge, who is looking at a zebra in plain view five feet away, perhaps at a zoo, would ordinarily be said to know that the animal is a zebra. Yet the grasp she has on its being a zebra, and her knowledge that this implies that it is not a mule, and so, in particular, not a cleverly disguised mule, and her basing her belief that it is not a cleverly disguised mule on her belief that it is a zebra, do not appear to be sufficient for her to know that the animal is not a cleverly disguised mule. Nothing she has done protects her against such unusually deceptive disguises.[3]

Valid deduction has apparently not preserved the knowledge property. One way to characterize what was had and what was lost on the way to the conclusion belief in this and many cases is via a property called *sensitivity*. In the probabilistic version of this property a subject has it when her belief in p is such that

$P(\text{-b}(p)/\text{-}p) > s$ $.5 \leq s \leq 1$ *Sensitivity*

The probability is greater than s that the subject doesn't believe *p* when *p* is false, where *s* is a threshold whose value is greater than .5, and how much greater depends on the utilities of the evaluator. *S* above fulfills this property for the proposition *q*, since the probability that *S* does not believe it's a zebra given that it isn't is high. In the probable ways for it to be something other than a zebra it would be an elephant or a giraffe, or a rock or a tree, all of which she can distinguish from a zebra. However *S* definitely does not fulfill the property for the proposition *p* above that follows from *q*: to be sensitive to *p* it must be probable that she is prone to detect a cleverly disguised mule and refrain from believing that it is not such a thing, but our *S* is not checking for this possibility. She is insensitive to *p*, and a definition of knowledge that requires sensitivity will imply that this is a case of closure failure of knowledge.

[2] We assume that the subject knows the implication, and bases her belief in the conclusion on her belief in the premise, because otherwise we should not expect her knowledge of the premise to act as a source of knowledge of that conclusion.

[3] A subject could believe there is no clever disguise on the basis of background knowledge that tells her it is very unlikely. But such a subject is appealing to further premises beyond *q*, so the case is not relevant to the potential closure failure in this case, where the question is whether her knowledge of *q* is sufficient basis.

Cases of purported closure failure that are typically discussed involve loss of sensitivity. Thus if we are concerned about closure failure, we could do worse than to use the sensitivity condition to map its behavior. I will add another condition to this:

$P(b(p)/p) > t$ $.5 \leq t \leq 1$ *Adherence*

which requires that the subject be likely to believe *p* given that *p* is true, even if other matters vary. She is not diverted from belief in *p* when she should not be. Together sensitivity and adherence are the probabilistic *tracking* conditions and I use a definition of knowledge in which they are necessary and sufficient conditions for knowledge of empirical truths[4, 5](Cf. Roush 2005, Chs. 1-3), though in one section below, as indicated there, I will drop the adherence requirement for ease of exposition. In the discussion below I will only be following the fate of the sensitivity property under known implication, but I have introduced the adherence condition because in the case of multiple-premise closure derivation of bounds on the loss of sensitivity over steps of deductive inference depends on the subject fulfilling both tracking conditions for each of the premises.

Loss of sensitivity in beliefs can be seen as growth in potential error because the less sensitive one is to *p* the more prone one is to the error of a false positive, believing p when p is false. Formally, the rate of false positives on *p* is

$P(b(p)/-p)$

the probability that you believe *p* when it is false. Thus, the level of sensitivity is 1 minus the false positive rate:

$P(-b(p)/-p) = 1 - P(b(p)/-p)$

As the level of sensitivity to a proposition goes down, the rate of false positive error on that proposition goes up. Here we are not concerned with growth of potential error with regard to a single proposition, but I will use the locution "growth of potential error" for a situation where in a step of deduction one goes from a given level sensitivity toward one proposition to a lower level of sensitivity toward another, logically implied, proposition. The growing error is only potential because in the situation that concerns us what is different with each step of deduction is a rate of error, a conditional probability. The absolute probability of the conclusion of the deductive step is the same as the probability of the premise. We assume premise q is true and because q implies p, that p is also true, and more generally that the absolute probability of the propositions is the same because $P(q) = x$ and $P(p/q) = 1$ together imply that $P(p) = x$.

---

[4] Roush (2005) defines knowledge disjunctively, imposing closure by brute force using a recursion clause. With enough steps of deduction such a clause allows growth of potential error that is as large as one likes while the conclusion belief still counts as knowledge. The current definition without the recursion clause removes that problem and the results in the rest of this paper suggest that we do not lose the advantages that the recursion clause was intended to give. Thus I now use the concept of knowledge with probabilistic sensitivity and adherence as both necessary and jointly sufficient conditions.

[5] The conditional probabilities listed for the tracking conditions are not sufficiently modal to support the intended interpretation of them as dispositions. However these conditional probabilities are merely shorthand for a universal quantification over a set of probability functions that corresponds to a set of possible situations. How to define which situations among all possible -p situations are the ones the subject is responsible for is a long-standing question which I answer via further probabilistic conditions (Roush 2005, Ch. 3).

Given that the step of deduction does not change the truth value or absolute probability of the proposition believed one might wonder why conditional probabilities concerning one's conclusion belief really matter. If the purpose of striving for sensitivity is the achievement of true beliefs, then since we are assumed to have achieved both for the premise, *q*, why should we also worry about sensitivity to the conclusion *p*? Achieving a true belief in *q* guarantees that a belief in *p* is also true, so what further work does sensitivity to *p* have to do? The question this asks is related to the question what the value is of knowledge (here sensitive, adherent true belief) over mere true belief. This form of question has been pressed particularly against externalist views of knowledge whose extra condition on knowledge beyond true belief tends to make a belief more likely to be true. But for sensitivity and adherence we can answer the question of why they are valuable even if one already has an actually true belief, and it is that such beliefs allow one to maintain a fitting belief state towards *p* over time and changing circumstances; they form the only Evolutionarily Stable Strategy in any game where false belief and failure to believe a truth are of disutility. The property of being derived from a merely true belief does not bring this robustness. (Roush 2010, Roush in press)

The zebra case is of a kind that I will call freefall, or maximal, closure failure, because in one step of deduction the subject goes from a belief that is tracking the truth well, to one where she has the maximum possible potential error. It is not just that she is a little more prone to be wrong about *p* than she was about *q*, but that she has nothing whatsoever that guards against the possibility of her believing *p* when it is false. We need a name for this case because we will see that there are cases of closure failure where sensitivity is lost over steps of deduction, but only relatively little per step. These are of interest because they mean that the existence of closure failure per se is not an obstacle to trust in deduction to give us knowledge indirectly. If your threshold for knowledge is 95% then if your conclusion belief tracks at, say, 90% your deduction will not have gotten you knowledge, but tracking at 90% is not terribly worse than tracking at 95%; an evaluator whose threshold was a little lower would have counted it as knowledge. We should be satisfied if the degree of loss of sensitivity were only 5% per step, as long as we could rely on it, and be sure it would never be 50% in one step. Because, as we will see below, we can derive bounds on the growth of error, we can keep track of our potential error by counting our steps and the loss of sensitivity will be manageable. For example, if you do not want to end up with less than 77% sensitivity, then do not trust beyond four steps of deduction from a premise with sensitivity 95%. Relying on deduction is only hazardous if there are cases where the loss of sensitivity is dramatic – where the conclusion belief has little or no sensitivity – and we do not know how to identify them.

In what follows I will identify a probabilistic condition that is violated in cases of freefall closure failure, and show that when the condition holds upper bounds on the loss of sensitivity, and accompanying growth of potential error, over steps of inference can be derived, if we assume that the subject's knowledge of the premise includes sensitivity. John Hawthorne (2005, 38) holds "out very little hope for a plausible restriction on closure that allows with the skeptic that we are in no position to know we are not brains in vats but that allows deduction to extend knowledge in the normal case." However the attempts at this so far have supposed that the distinction that secures a domain for safely extending knowledge should be made between cases of closure failure and cases of closure fulfillment. Since, as I will show, closure failure per se does not line up with hazardous uses of deduction, we may hope that the quantitative analysis of the growth of potential error succeeds in defining a zone where deductive inference can be trusted to extend knowledge.

Definitions: Knowledge, Knowledge of Logical Implication, Basing

Knowledge is *closed*, for our purposes, if and only if for knowledge of *p* it is sufficient that a subject knows some premise *q*, knows that *q* logically implies *p*, and has a belief in *p* that she bases on her belief in *q*. In this paper I will focus on closure of empirical knowledge, the only kind of knowledge for which sensitivity is an appropriate requirement, in my view. (Roush 2005, 134-147) To derive quantitative results about closure failure and deterioration of sensitivity, we need in addition to the tracking conditions defining empirical knowledge, above, a set of necessary conditions for *knowledge of logical implication*, and necessary conditions for *basing belief in p on belief in q*. As it happens we need not commit ourselves to full definitions of these latter two properties since a few necessary conditions will be sufficient to derive the properties of interest.

The tracking definition of empirical knowledge is guided by the idea that to count as knowledge a belief should be responsive to the way the world is. Logical (and other necessary) truths are different from empirical truths[6], so it can be expected that the appropriate kind of responsiveness would also be different. In particular, in my view to qualify as knowledge of a logical truth a belief should be responsive to the special place that logical truths have *among* propositions in our language. We see this special place by first defining what it is to know that one proposition logically implies another. Then, taking a cue from the fact that logical truths are implied by all propositions, knowledge of logical truth p is the responsiveness of your belief in p to the fact that such a truth is implied by every proposition in your language (Roush 2005, 134-147, Roush 2012, 246-253).

For current purposes we need only the first step. Thus, *S knows that q logically implies p* only if (Roush 2005, 134-147, 2012, 246-253) *q* does logically imply *p* and:

3') $P(\text{-b}(q)/\text{-b}(p)) > u,\ .5 < u \leq 1$

4') $P(\text{b}(p)/\text{b}(q)) > v,\ .5 < v \leq 1$

3' says the probability *S* does not believe *q* given that she does not believe *p* is greater than *u*. *S* is responsive to the fact that logical implication supports modus tollens. 4' says the probability she believes *p* given that she believes *q* is greater than *v*. She is responsive to the fact that logical implication supports modus ponens. Logical implication is a relation among propositions; for a subject's beliefs to be

[6] Thus I assume an analytic/synthetic distinction. We can, with Quine, allow revisions of our logic, but we will, with Carnap, count those as changes to our language (Carnap 1950, Quine 1951).

responsive to that relation requires that her beliefs in the propositions relate dispositionally to each other in the same way as the propositions do.[7, 8 9]

A necessary condition for one belief to be based on another can be gotten from the simple thought that a person's belief in *p* is not based on belief in *q* if she would believe *p* even if she did not believe *q*. In that case there must be something other than the belief in *q* supporting belief in *p*. So, suppose *S* believes *p* and believes *q*. *S*'s *belief in p is based on her belief in q* only if

5) $P(\text{-b}(p)/\text{-b}(q)) > z, \quad .5 \leq z \leq 1$ *Basing* I

6) $P(\text{-b}(p)/(\text{-b}(q) \wedge q')) = P(\text{-b}(p)/\text{-b}(q))$ for all q' *Basing* II

That is, according to 5, the probability is high that *S* would not believe *p* given that she does not believe *q*, and, according to 6, no other fact can change this, substituting itself for the missing belief in *q*. These conditions are clearly too weak to capture all of what basing means but they express a clear sense in which belief in *p* "depends on" belief in *q*, and the remarkable thing is that they are sufficient to derive the error bounds we need.[10]

The basing conditions conform to intuitions in cases. For example if

q: There is a cookie in the jar
p: The jar is not empty

it is sensible that we should count a belief in *p* as based on a belief in *q* only if, in colloquial terms, the subject wouldn't believe the jar was not empty if she didn't believe there was a cookie in the jar; there isn't any other matter, *q*', a brownie being in the jar for example, that would make her believe the jar was not empty if she didn't believe there was a cookie in the jar. More rigorously, the probability she

---

[7] A further natural condition on responsiveness to logical implication would be appreciation of monotonicity. Thus, where *q* implies *p*, for all *q*',

$$P(b(p)/(b(q) \wedge q')) \geq P(b(p)/b(q))$$

that is, the probability that the subject believes *p* given that she believes *q* is not diminished by any other matter or belief. This means in particular that she is responsive to the fact that adding a premise does not undermine the logical implication.

[8] As above with the tracking conditions, each of these conditional probabilities is shorthand for a universal quantification over a set of probability functions, which set is identified using a further probabilistic condition (Roush 2005, Ch. 4).

[9] Sensitivity to a logical or other necessary truth would involve the negation of a necessary truth in the condition of a conditional probability, which is not defined (without a good deal of extra sophistication). The fact that in this account knowledge of a logical truth is not a responsiveness to that logical truth – and so does not include sensitivity as defined above – means that we avoid that problem.

[10] One might think that for basing a condition should be imposed that the subject has a forward commitment to *p* that comes from her belief in *q*. Perhaps so, but that condition would naturally be expressed by $P(b(p)/b(q)) > v$, $.5 \leq v \leq 1$, which we need not impose on basing here since it is already in condition 4' defining appreciation of the modus ponens direction of the logical implication. Notably that condition makes no difference to the derivation of error bounds for single premise cases as we will see below. We might also want to rule out counting it as basing if the subject wouldn't believe *p* when she didn't believe *q* because she wouldn't believe *p* regardless of not believing *q*. For this we could add the condition $P(\text{-b}(p)/\text{-b}(q)) > P(\text{-b}(p))$. This issue does not affect the bounds on potential error derived below.

does not believe the jar is not empty given that she does not believe there is a cookie in the jar is high, and that is so regardless of anything else.

What if there *were* a brownie in the jar? Should she really refrain from believing the jar is not empty? The question is not what she should believe about the jar but what it takes for her belief that the jar is not empty to be based on her belief about the cookie. If were she not to believe there is a cookie she would still believe the jar is not empty, say because of the brownie, then her belief that the jar is not empty is not based solely on her belief that there is a cookie.[11] The conditions for basing aim to capture the notion of believing *p* "only because" of belief in q, since if something else contributes to the subject's believing *p* then we are not evaluating whether knowledge of *q* is sufficient.

Though, as with 3, 4, 3', and 4', the intended conditions in 5 and 6 are dispositional[12], the dependence is not causal, another possible aspect of the basing relation that is not captured in these conditions. However this weak condition is sufficient for the role that basing plays in deriving the upper bounds on deterioration of sensitivity that I discuss below. As with knowledge of logical implication above, the basing conditions put requirements on the dispositions of the subject's beliefs with respect to *each other*, and not their responsiveness to the world.

Good Behavior

The question whether or how much sensitivity is lost over steps of inference can now be formulated rigorously. In the case of an inference with one premise, *q*, and one step of valid inference[13], from *q* to *p*, assume that the subject believes *q* and tracks *q* at levels s and t for the sensitivity and adherence conditions respectively. Assume also that *q* logically implies *p*, and that the subject believes *p* and is responsive at levels *u* and *v*, respectively, to the modus tollens and modus ponens directions of logical implication between *q* and *p*. Suppose too that she bases her belief in *p* on her belief in *q* at level *z* and fulfills the second basing condition.

To answer whether or how far knowledge is closed we need to ask whether or to what degree the belief that our subject has in the conclusion *p* is sensitive to *p*. Using the definitions given above, and under a condition to be identified below, the assumptions about the subject are sufficient to determine her level of sensitivity to *p*. By the rule of total probability[14], the sensitivity of the conclusion-belief in *p* is equal to an average of P(-b(*p*)/-b(*q*)) and P(-b(*p*)/b(*q*)) weighted by P(-b(*q*)) and P(b(*q*)) respectively, all given -*p*:

$P(\text{-b}(p)/\text{-}p) = P(\text{-b}(q)/\text{-}p)P(\text{-b}(p)/\text{-b}(q)\wedge\text{-}p) + P(\text{b}(q)/\text{-}p)P(\text{-b}(p)/\text{b}(q)\wedge\text{-}p)$

[11] In this case the basing conditions would be fulfilled if we took *q* to be the disjunction of "There is a cookie" and "There is a brownie" or if we took *q* to be "There is a cookie and a brownie". We could also model it as a multiple-premise implication with "There is a cookie" and "There is a brownie" as premises. The multiple-premise basing conditions are given in Appendix 2.

[12] See footnote 5.

[13] For purposes of exposition my wording "step(s) of inference" and "step(s) of deduction" depicts the subject as doing explicit reasoning, but this is only a special case. The closure property has been defined as involving known implication rather than valid inference, and one can fulfill the necessary conditions I have put on known implication and on basing without explicit reasoning, with only dispositions to respond to some beliefs by forming or dissolving others.

[14] Total probability says, for the simplest case: $P(A/C) = P(A/B\wedge C)P(B/C) + P(A/\text{-}B\wedge C)P(\text{-}B/C)$

Under our assumptions and the further condition I will identify below as the *No-Freefall Condition*, this implies[15]:

$P(\text{-b}(p)/\text{-}p) \geq P(\text{-b}(q)/\text{-}q)P(\text{-b}(p)/\text{-b}(q)) + P(\text{b}(q)/\text{-}q)P(\text{-b}(p)/\text{b}(q))$

This in turn implies:

$P(\text{-b}(p)/\text{-}p) \geq sz + (1\text{-}s)(1\text{-}v)$

$\geq sz$

The subject's level of sensitivity to a *p* her belief in which is based on her belief in a *q* that logically implies *p* and which is itself sensitive to degree *s*, is at a minimum *s* multiplied by the level of her basing of her belief in *p* on her belief in *q*. In other words, her sensitivity to the conclusion is diminished from her sensitivity to the premise only by the extent to which she does not exclusively rely on her belief in the premise.[16] If $P(\text{-b}(q)/\text{-}q) > s$, then $P(\text{-b}(p)/\text{-}p) \geq sz$. If her belief in *p* were based solely on her belief in *q* then *z* would be one and her sensitivity to the conclusion *p* would be just as high as her sensitivity to the premise *q*. The more the subject depends on her belief in *q* for her belief in *p*, the more she gets for her belief in *p* of what her belief in *q* has to offer.

In a one-premise, one-step inference in which the thresholds are at $s = .95$, $z = .95$, the subject's sensitivity to the conclusion, *p*, will be .90. This makes sense of the many cases that intuitively seem to have no closure failure at all. They may be cases where we are successfully imagining perfect basing, and so should not expect to have any loss of sensitivity at all, or they may be cases where the basing level is so high that they are intuitively indistinguishable from the perfect case. For *m*-step, one-premise implications, the preserved sensitivity level, $P(\text{-b}(p)/\text{-}p)$, is sz^m. Thus, for two steps and all thresholds set at .95, $P(\text{-b}(p)/\text{-}p)$ is .86, for three steps .81. This matches our intuitive sense that our conclusions do become more tenuous the more steps of deduction we use to get to them, as we can see by comparing how much we would trust an inference from *q* to *p* that was one step and one that involved ten steps.[17]

Bounds on the loss of sensitivity can also be derived for multiple-premise inferences, with the definitions of basing and knowledge of logical implication, and the No-Freefall Condition, suitably generalized.[18] For *n*-premises and a single step of deduction, preservation of sensitivity depends on adherence to the premises, as it does not in the single-premise case, but assuming $s = t$, the sensitivity of the conclusion-belief is zs^n. This analysis makes sense, in particular, of the way that knowledge appears to be lost when deduction takes the form of conjoining premises into a conjunction. It seems that a subject may know of *X* that he will come to the meeting, and of *Y* that she will come to the meeting, but the longer the list of such beliefs gets the more implausible that she knows them all in a conjunction. The longer the list, the more likely that *somebody* will not make it to the meeting. Conformably, in this case even if the basing is perfect, and the sensitivity and adherence of each premise belief are .95, the sensitivity of the conclusion will diminish with every additional premise used. With

[15] See Appendix 1.

[16] In the single-premise case the preservation of sensitivity does not depend on the level of adherence, *t*, of the premise belief, nor on the subject's knowledge of the implication, though it depends on the implication itself.

[17] Of course the possibility of a reasoning error that accompanies every step of deduction also contributes to our intuition that the standing of our beliefs is more tenuous with more steps, and in a given case that might be the dominant explanation for the intuition.

[18] See Appendix 2. I have not investigated the case of multiple-premise, multiple step inference, i.e., where both *m* and *n* are greater than one.

two premises the sensitivity of the conclusion belief will be .86, with three it will be .81, with four .77, with five .74. With eleven premises one will have a conclusion belief with no sensitivity at all. That is in the uncommon case when one has .95 sensitivity and .95 adherence to each of the claims that a person will come to the meeting. With .90 sensitivity and adherence to each premise, one will have virtually no sensitivity (.53) to the conjunction if the group has as few as five members. If one has sensitivity .85 to each of *X*, *Y*, etc. attending, then one is down to .50 sensitivity to the proposition that they will all come if the group has as few as four members. But it is no surprise that if you need a quorum of *n*, you should get reassurances of attendance from some number greater than *n*, and that how far greater depends on the *n* and the reliability of the reassurances. The quantity of loss with increasing numbers of premises parallels intuitions, as does the fact that in these cases sensitivity erodes gradually without a dramatic change at a particular number or for a particular type of premise.

It may seem odd that preservation of sensitivity should depend so strongly on basing, and apparently nothing else. But it is right that basing should end up with a prominent role in preservation of a knowledge property, since it is what tells us when to expect that any of what the premise belief has could transmit to the conclusion belief; the less basing there is, the less receptive the conclusion believing has been to what the premise belief has to offer. It may seem that preservation of sensitivity is too easy on this account, requiring only perfect basing. However perfect basing alone is not sufficient, because for the well-behaved cases just discussed we are assuming fulfillment of the No-Freefall Condition discussed below. The easy preservation of sensitivity doesn't come from a mysterious power of basing but from the way that the No-Freefall Condition excludes funny business.

Important as basing is for preserving sensitivity, it is also easy and common not to fulfill it. One is not basing perfectly if one has any chance of making a certain kind of error in the deduction; though you actually have a belief in the premise *q*, the error might be of a sort that would give you a belief in *p* even if you did not have belief in *q*. One also does not base belief in *p* fully on belief in *q* if one is relying in addition on background knowledge, consciously or unconsciously. The human mind makes prodigious use of background knowledge, and when it does, then the argument from *q* to *p* is an enthymeme with tacit premises. Consider that in the zebra case our intuition that the subject does not know the animal is not cleverly disguised mule recedes when we reflect that such a disguise would have no immediately obvious incentive and require a lot of resources, that zoo conspiracies of this sort are rarely if ever heard of, and that the subject also knows all of these things. But when we have these thoughts, we should quickly draw ourselves back to the topic, which is not whether she knows the conclusion but whether she knows it on the basis merely of her knowledge of the premise that the animal is a zebra that was gleaned from taking a look at it.

Often intuitions tell us clearly what the basis of a belief is, but sometimes they do not, and stipulations can be untrustworthy. Consider the case where

q: There is a dog.
p: There is an animal.

Much is often made of the fact that we can imagine a world in which if there weren't a dog there would be a convincing enough fake cat that the subject would not believe there was a dog, but would still believe there was an animal. If so, then the subject is sensitive to *q* but not to *p*, even though *q* implies *p* and we can suppose she knows this. This is a case where sensitivity has been lost over inference, but one where intuition apparently says there should not be closure failure; surely her knowledge that it is a

dog is enough to give her knowledge that it is an animal. This case seems to show that sensitivity is too much to expect from knowledge in the first place.

However, to make the example work as a case where a sensitivity requirement implies implausible closure failure we must assume the subject's belief that there is an animal is based on her belief that there is a dog, and in this case stipulation is not sufficient to insure this. The unusual feature we have introduced in which the nearest alternative to there being a dog is there being a fake cat implies that the subject violates the first basing condition. The most likely situation given that the subject does not believe there is a dog, -b(*q*), is that there is a fake cat, but in that case she is still likely to believe there is an animal. That is, basing requires

$P(\text{-b}(p)/\text{-b}(q)) \geq z, \ .5 < z \leq 1$

but under the stipulations of the case, the value of this term is actually low. She is likely to believe there is an animal, even if she does not believe there is a dog.[19]

This result is not a quirk of the formalism but corresponds to epistemologically relevant features. If the most probable alternative to there being a dog was there being nothing, then basing would be possible. For the likely ways for her to not believe there is a dog would then be for her to believe there is nothing, and so, to not believe there is an animal. Her belief that there is a dog would have a sensitivity that does not take for granted that there is an animal, so the belief that there is a dog can be the basis of her knowledge that there is an animal.

Though the fakeness of the imagined cat does prevent sensitivity to the conclusion *p*, it is not what prevents basing of belief in *p* on belief in *q*. If the most probable alternative to a dog was a real cat then her not believing there was a dog also would not deter her from believing there is an animal. The most probable alternatives to there being a dog are what determines whether belief that there is a dog can be the basis of belief that there is an animal because it is that stipulation of probable alternatives that determines what is and is not still up for grabs when the subject takes the option to believe or not believe there is a dog. It tells us what, given the way the world is, the subject does and does not need to be sensitive to in order to be sensitive to there being a dog. If the most probable alternative is a cat, then there being an animal there is already settled, and we should not expect her belief that there is a dog to be what gives her knowledge of that.[20] Whether there is closure failure depends on which alternative scenarios are specified as relevant because whether there is basing depends on them, and whether there is closure failure depends on basing.

[19] Arguably the same failure of basing is present in the zebra case, because a normal subject is likely to believe it's not a cleverly disguised mule even if she does not believe it's a zebra, due to the improbability of such a thing for the reasons noted. I will not press the case because even if the subject has that disposition she will likely not retain it if we ask her directly whether she knows it is not a cleverly disguised mule. However it is interesting because the lack of transfer of sensitivity would not be due to something about sensitivity but something about what is required for basing.

[20] This case appears analogous to failure of transmission of warrant in Wright's (2003) sense, where transmission fails in cases for which in order for the premise to be warranted the conclusion must already have been warranted. See below in the main text.

Pursuing Knowledge Indirectly

If knowledge requires sensitivity then knowledge is not closed under known, basing, logical implication. Sensitivity can deteriorate from the premise to the conclusion, and this is true not only in freefall cases but also in the well-behaved domain just discussed. But though knowledge is not closed, we can identify a property that it has in the well-behaved domain that is as good as closure for the purposes of supporting indirect inquiry by logical deduction. *S* knowing *q*, knowing that *q* implies *p*, and basing her belief in *p* on her belief in *q* are not sufficient to guarantee she knows *p*. However, on the current account, knowledge in the well-behaved domain has a related property. If for ease of exposition in this section we take knowledge to require only sensitivity and not adherence, we have:

if
q implies *p*
S knows$_s$ *q* (sensitive at > *s*), and
S believes p on the basis of her belief in *q* (> *z*),
then

S knows$_{sz}$ *p* (sensitive at > *sz*)[21]

I call this property *closure enough* because it gives us the most important thing that closure gives for the purposes of inquiry, namely, assurance that we can safely bypass the work of investigating the conclusion directly, because knowledge of the premise and a logical inference will be sufficient for the conclusion belief to have a specifiable level of sensitivity. In the domain where the No-Freefall Condition is fulfilled, we can accept a belief as knowledge by knowing *nothing more than* that appropriate error thresholds were met on a belief it was validly inferred from, and fulfillment of the basing conditions. As with cases of closure fulfillment, so in cases of fulfillment of closure enough, we will not have any surprises or explosions in the potential error of the conclusion belief.

The only difference is that due to the fact that the conclusion belief could have a lower sensitivity than the premise belief, for purposes of making use of closure enough in investigation it helps to work backwards beginning with the sensitivity level, *s*, that the evaluator requires for a belief to count as knowledge. Knowing that is the sensitivity we want for the conclusion belief, we can use its value to identify options for achieving it indirectly, via the many permutations of possible values for *z*, *m*, *n*, and the *s* and *t* of the premise(s) that will yield a sensitivity at least that high in the conclusion. For example, if your desired sensitivity level for the conclusion is > *r*, then you could choose a single-premise, single step implication – that is, $n = 1$, $m = 1$ – and choose an *s* for the premise and a *z* for basing in any way that makes $sz > r$ true; if $r = .95$ then $s = .97$ and $z = .98$ will do it, but so will $s = .98$ and $z = .97$, and an infinite number of other options. You will simply aim for a degree of sensitivity in the premise belief(s) that will be sufficient to bring the degree of sensitivity you want for the conclusion belief under the assumption of a level of basing.

Generally speaking in this domain, increasing the sensitivity of the premise-belief(s), other things equal, will always increase that of the conclusion belief, and increasing basing, other things equal, will always increase sensitivity of the conclusion belief. Fewer premises and fewer steps of implication between premises and conclusion, other things equal, will also always increase the sensitivity of the conclusion

[21] Obviously this formulation is for the special case of one-premise, one-step inferences, but it could be generalized.

belief. In the well-behaved domain the sensitivities of premise and conclusion beliefs are related in a deterministic way, making the rate of sensitivity loss completely predictable. The set of cases where knowledge is closed is just the special case of this where the relation of the strength of sensitivity of premise belief and conclusion belief is always identity. Closure is a special case of closure enough.

The No-Freefall Condition

For the practical purposes of inquiry, we do not need to identify a domain in which knowledge is closed, but only to identify a domain in which the degree of loss of sensitivity is fully predictable. We get this in the domain of well-behaved cases discussed above, which is identified by fulfillment of the following condition. In the single-premise, single-step case, with q the premise, and p the conclusion of the implication, the requirement is:

P(-b(*q*)/-*q*.-*p*) ≥ P(-b(*q*)/-*q*) *No-Freefall Condition* (NFC)

The probability you don't believe *q* given that *q* is false and *p* is false is not less than the probability you don't believe *q* given that *q* is false with no stipulation about *p*. In other words, your sensitivity to the premise of the inference is not lower when the conclusion is false; the falsity of *p* does not reduce your sensitivity to *q*.[22] A case will be well-behaved as long as the falsity of *p* does not interfere with the relation between the fact of -*q* and refraining from belief in *q*. This condition classifies the zebra case above as one in which we should not trust deduction from "That is a zebra" to give us knowledge of *p* indirectly, because when *p* is false, meaning that it *is* a cleverly disguised mule, the subject is less likely to refrain from believing that it is a zebra despite the fact, in such a case, that it is not a zebra.

A belief that it is not a cleverly disguised mule that is based on background knowledge, such as the absence of an obvious motive for disguising a mule, difficulty of preparing the costume, and implausibility of elaborate zoo conspiracies, seems intuitively to have some purchase on the matter, more at least than we get from viewing the distinctive appearance of the animal and declaring "Zebra!" without a concern about the possibility of disguises. Conformably the No-Freefall Condition is not violated in this case of believing on the basis of background knowledge. The animal's being a cleverly disguised mule would not reduce your sensitivities to the facts that zoo conspiracies are unlikely and very convincing fakes are hard to pull off.

In accord with the intuition that the tree people get to know that the predator is a snake by process of elimination and without further investigation, the No-Freefall Condition is not violated in that case either:

q_1: The predator is either a tiger, or a hawk, or a snake
q_2: The predator is not a hawk
q_3: The predator is not a tiger

[22] For ease of exposition I am eliding the distinction between cases where the falsity of *p* reduces the sensitivity of the subject's belief in *q* and cases where the falsity of *p* eliminates the sensitivity altogether. It is the latter cases that are maximal closure failures. The degree of loss of sensitivity in the former cases will be predictable, but will depend on the quantities in this formula rather than on the degree of basing.

p: The predator is a snake

The probability that they would not believe it was not a hawk given that it was a hawk is unaffected by a stipulation that the predator is not a snake. The tree people's sensitivity to the premises does not depend on the conclusion being true.

It is useful to contrast the No-Freefall Condition with an informal idea that has been discussed as a characterization of cases of closure failure. According to this idea closure fails in cases where the conclusion is *heavyweight*, that is, a proposition, *p*, for which there can be no perceptual reasons for believing *p*.[23] (Dretske 2005, Hawthorne 2005) Many cases that are taken to be closure failures fulfill this criterion. The proposition that I am a brain in a vat is constructed to be unknowable by perceptual means, as is the proposition that a castle came into existence five minutes ago along with all of the evidence that it is much older. No surprise, then, that what the subject did to come to know that she has a hand or that that castle is old would not be enough to give her knowledge that she is not a brain in a vat or that the world didn't come into existence five minutes ago with a perfect disguise.

However while being heavyweight may be a necessary condition for skeptical cases it cannot capture the domain of closure failure since some cases of purported closure failure have, or can be made to have, lightweight conclusion beliefs. For example, in the zebra case we could have taken the conclusion of the inference to be not unknowable by perceptual means, but merely unknown by the means the subject used to come to knowledge of the premise. Perhaps the disguise is clever enough to fool someone four feet away, but if the subject had gotten closer and inspected the fur the fakery would have been noticeable. It would still be the case that she could have gotten perceptual knowledge of the premise – that is a zebra – without having the goods for knowledge of the implied conclusion that it is not a cleverly disguised mule. It could still be a case of closure failure.

The goal of defining a criterion of knowledge that lines up skeptical cases with closure failure cases may be hopeless, but if our interest is in showing how to avoid hazardous deductions in the course of pursuing knowledge indirectly, then that alignment is not relevant. The important issue is how much potential error a given step of deduction introduces, and whether we can know that without doing the work of investigating the conclusion independently. So the important property of the conclusion that the animal is not a cleverly disguised mule is not whether it is unknowable but whether the one step of deduction that led to that belief introduced massive potential error. If avoiding uncontrolled growth of potential error is the point of making a distinction among cases of deduction from apparently known premises, and if potential error is measured by lack of probabilistic sensitivity, then the No-Freefall Condition gives a condition sufficient to avoid the hazardous deductions.[24]

A key difference between characterizing closure failure as involving heavyweight propositions and characterizing the maximal closure failures as violations of the No-Freefall Condition is that the first

[23] Hawthorne defines the heavyweightedness of a proposition using the condition that perception or perception aided by logic cannot give reasons to believe it. I have left out the further clause about logic because with it heavyweightedness would not be well-defined for determining of a case of purported closure failure whether the conclusion of the inference was heavyweight. Someone who independently thought the inference was not a case of closure failure could identify the conclusion as a lightweight proposition on grounds that it is known by that very inference.

[24] I have only shown that NFC is a sufficient condition for avoiding uncontrolled growth of error, not that it is necessary, which means that there may be cases that fail NFC but have bounds on the loss of sensitivity.

identifies the problem in terms of the content of the conclusion proposition. Inferences from ordinary observational beliefs to the conclusion that the world did not come into existence five minutes ago with all the signs of old age will fail to bring knowledge of the latter because the latter has a content that prevents its being known by perceptual means. In contrast the No-Freefall Condition sees the problem as hanging on the relation between the premise and conclusion. The case is a problem because if the world came into existence five minutes ago with all of the signs of old age that it currently has, then you would not be sensitive to your premise that the castle is at least one hundred years old, or to many other observational propositions. The big hoax would mean that you would believe that castle has been there for a very long time even if it hasn't. These are cases where the falsity of the conclusion undermines the sensitivity of your belief in the premise of your deduction. This difference is why the No-Freefall Condition correctly classifies the version of the zebra case where the conclusion is perceptually knowable, yet not known by the means used to know the premise.

The avoidance of untrustworthy cases via the No-Freefall Condition is analogous to Crispin Wright's characterization of failure of transmission of warrant. (Wright 2003) Though sensitivity and warrant are quite different properties a belief might have, and neither implies the other, Wright's criterion, like the No-Freefall Condition, focuses on the relation between premise and conclusion. In cases of transmission failure Wright says you will find that the subject can be warranted in believing the premise only by being antecedently warranted in believing the conclusion. In cases of failure of the No-Freefall Condition the subject is only sensitive to the premise if the conclusion is true. The No-Freefall Condition agrees with Wright's characterization in some clear cases. For example,

q_1: That girl looks just like Jessica
q_2: That girl is actually Jessica

p: That girl is not Jocelyn

where Jessica and Jocelyn are identical twins. Arguably, one needs to be warranted in believing the girl is not Jocelyn in order to be warranted in thinking she is Jessica judging only by her looks. The No-Freefall Condition gives the analogous verdict on this case. The girl's being Jocelyn makes the subject less sensitive to the girl's being Jessica – less likely to refrain from believing it is Jessica given that it isn't – because Jocelyn looks exactly like Jessica, and that will send him into a freefall of potential error if he believes she is not Jocelyn on the basis of a belief so formed that it is Jessica. In both analyses what prevents epistemic good standing of the conclusion is not that it is unknowable, but that something about its relation to the premise prevents the latter's epistemic good standing.

The No-Freefall Condition will be useful in empirical investigation for identifying a wholesome domain where deduction is a safe knowledge-extender. Its violations may seem uninteresting for science, serving merely the negative purpose of illustrating what sort of thing to avoid. This appears to be supported by the fact that the condition is so easy to evaluate in familiar cases. In the freefall closure failures we have seen here the condition is decidable a priori, or using only obvious background knowledge. A young world's perfect disguise as an old one will prevent you from being sensitive to any old-age claim about its constituents. The disguised mule will have you believing it is a zebra when it is not, as long as your inspection for zebrahood is not close enough to discern the fakery.

There is no investigative work required to see these things, and the content of these violations also appears to be of no interest to empirical science. One might think that even if determining whether the

NFC holds in a given case does require substantive investigation, the kind of undermining involved is too recherché to play a significant role in scientific investigation. Concern over whether evidence is misleading is a daily occurrence – Is the cell structure really that small or is the microscope showing an artifact of the dehydration technique of preparation? But in the kind of possibly misleading evidence involved in freefall closure failure, whether evidence for q is misleading depends on the truth value of what we want to infer from q. It would be as if you were hoping to infer the cells' fast reproduction rate because it follows from small cell structure, but whether the dehydration preparation technique produces an artefactual appearance of small cell structure depends on whether there is fast reproduction. Surely such unlucky interdependences are uncommon.

They may not be common, but there are cases of freefall closure failure in science that are difficult to verify, unsafe to ignore, and impossible to banish from the science itself. They can also be significant. For in cases where verifying the freefall closure failure requires substantive investigation, we may remain unaware of it for a long time and for all that time mistakenly suppose that knowing *q* – as we suppose we do – is sufficient for knowing a *p* that is implied by *q*. Identifying freefall closure failure within science can expose an unexamined assumption, *p*, our justification for which depends on its truth.

For example, a great insight that Copernicus had about the heavens was a geometric argument to the effect that it is possible that the stars would appear to us to be moving even if they are not. (Copernicus 1939, Book I, Section 5) In this possibility, an unmoving earth with stars rotating around it and a rotating earth with stars fixed would make the stars appear the same in the view of an observer fixed to the earth: either way, the stars would appear to move. If so, then the fact that they appear to move would not be good evidence that they do move.

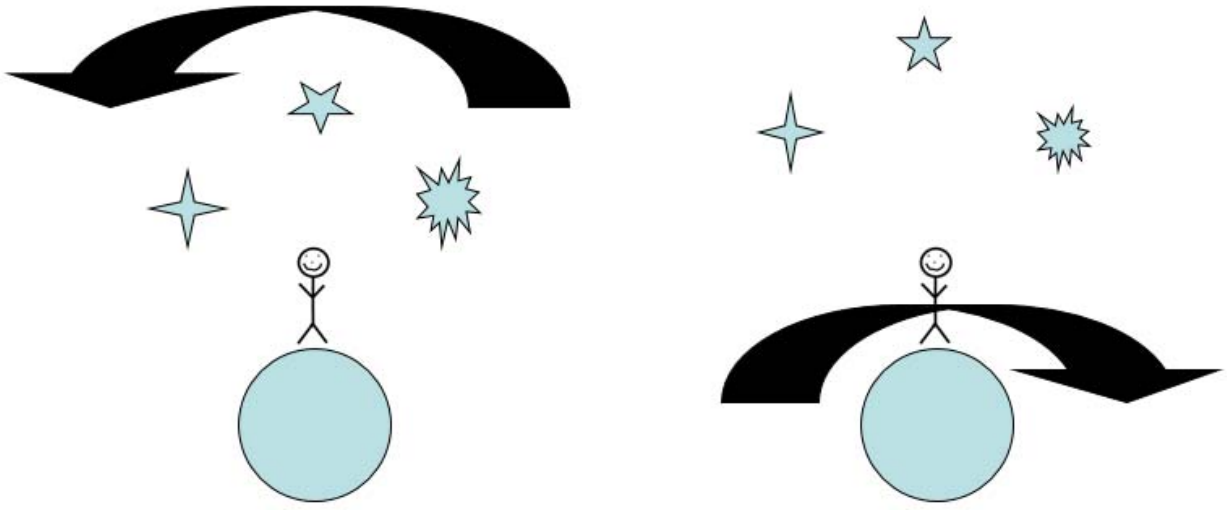

Consider:

q: The stars move, earth at center (Ptolemy)
q': The stars appear to move when viewed from the earth.

p: The earth does not move (rotate).

From the stars appearing to move as they do to the left-hand person, and from their actually moving so, and assuming the earth at the center of the world, we can infer that the earth is at rotational rest; *q*' and *q* together imply *p*. If the earth indeed does not rotate, *p*, then we are in the rotational rest frame, so we can tell whether the stars move or not, *q*; their actual movements are just as they appear to us to be. Copernicus's suspicion was that if the earth *did* rotate then we would not be able to tell that the stars were at rest even if they were. So his suspicion was that if the earth were not at rest our observations could not tell us not to believe that the stars were moving, because all of the evidence would look the same. (Copernicus 1939, Book 1, Section 5) This is, then, the suspicion that:

$P(\text{-b}(q)/\text{-}q.\text{-}p) < P(\text{-b}(q)/\text{-}q)$ Failure of No-Freefall Condition

The probability we do not believe the stars move given that they do not, is lower when the earth rotates than when it does not. Our ability to discern that the stars do not move would be hampered by the earth rotating. If so then a belief that the stars move that is supported by the best evidence available at the time is not a route by which one can come to know that the earth is at rest. That the earth is at rest was a framework proposition, assumed without active investigation, and the fact that it was implied by a proposition with apparently obvious evidence gave cover for it to retain that status. Only by identifying this as a freefall case could one see that even in the best case where *q* were true and one knew it, that would not give one knowledge of *p*, that is, that deduction from *q* was not a way to know *p*.

We didn't evolve with inner alarm bells to detect this case of closure failure, so it took some investigation to discover it. I call this thought of Copernicus's a suspicion because geometry alone is not sufficient to defend the claim. Copernicus's contemporaries and many for centuries afterwards were not out of place asking how it could be that we remain fixed to the earth through this gigantic movement, and are not carried off in the wind or hanging on for dear life. If the earth rotates or revolves, then why would we not feel it? There was nothing in the Aristotelian physics of the day that could explain how we would stay planted in the same spot, with no indication that we were on a moving earth.

The faith of some that Copernicus had the actual geometry of motions right, that not only did the moon revolve around the earth, but the earth also revolved around the sun and rotated on its axis, eventually led to a physics that could explain why we do not get left behind in the earth's movement. The theory of gravitational attraction not only put the heliocentric view on a solid physical footing in general, but also finished the argument that the key evidence for a framing assumption of the geocentric view had been misleading. The suspicion that there might be maximal closure failure, that rotation of the earth would make the apparent motion of the stars into misleading evidence of their motion, led eventually to a great advance in science. But the difficulty of even verifying that this was a freefall case – which required a new physics – protected the assumption that the earth is at rest for a long time.

Conclusion

The threat that closure failure presents to extending our knowledge by reasoning only lies in cases with large and unpredictable growth in potential error over steps. Thus the important distinction for these purposes is not between cases where closure fails and cases where it does not, but between maximal closure failure and gradual growth of potential error over steps of reasoning that has predictable upper

bounds. If loss of sensitivity is a kind of error we want to avoid, then in single-premise cases we can rely on fulfillment of the No-Freefall Condition (and other conditions for multiple premises) to indicate cases where we can trust deduction for expanding our knowledge. In such cases loss of sensitivity comes in degrees, gradually and predictably. However, cases where deduction can play its knowledge-multiplying role are not the only cases of interest to empirical science, because cases of freefall closure failure don't only occur in radical skeptical philosophy. Those cases that are relevant to empirical science are more hazardous than those in skeptical philosophy, because, as illustrated in the Copernican example, we generally lack inner alarm bells for cases we did not contrive ourselves. The No-Freefall Condition is useful because it tells us what question to ask to reassure ourselves that deduction can be trusted. In most cases it will be obvious that the NFC is fulfilled, but any case where it is not obvious may be the seed of a revolution.

References

Carnap, Rudolf (1950), "Empiricism, Semantics, and Ontology," *Revue Internationale de Philosophie* 4: 20-40.

Copernicus, Nicolaus (1939), *On the Revolutions of the Heavenly Spheres*. In *Ptolemy*, *Copernicus*, *Kepler*, Volume 16 of *Great Books of the Western World*, *Encyclopedia Brittanica*. Chicago: University of Chicago Press.

Dretske, Fred (2005), "Is Knowledge Closed under Known Entailment? The Case Against Closure," in *Contemporary Debates in Epistemology*, Matthias Steup and Ernest Sosa eds. Oxford: Blackwell.

Hawthorne, John (2005), "Is Knowledge Closed under Known Entailment? The Case For Closure," in *Contemporary Debates in Epistemology*, Matthias Steup and Ernest Sosa eds. Oxford: Blackwell.

Malinas, Gary and Bigelow, John, "Simpson's Paradox," The Stanford Encyclopedia of Philosophy (Fall 2016 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/fall2016/entries/paradox-simpson/>.

Quine, W.V.O. (1951), "Two Dogmas of Empiricism," *The Philosophical Review* 60: 20–43.

Roush, Sherrilyn (2005). *Tracking Truth*: *Knowledge*, *Evidence*, *and Science*. Oxford: Oxford University Press.

------- (2010), "The Value of Knowledge and the Pursuit of Survival," *Metaphilosophy* 41,3: 255-78.

------- (2012), "Sensitivity and Closure," Kelly Becker and Tim Black eds. *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press.

------- (in press), "The Difference between Knowledge and Understanding," *Explaining Knowledge*: *New Essays on the Gettier Problem*, Oxford University Press.

Skyrms, Brian (2010). *Signals*: *Evolution, Learning, and Information*. Oxford: Oxford University Press.

Wright, C. (2003). "Some Reflections on the Acquisition of Warrant by Inference." In S. Nuccetelli Ed., *New Essays on Semantic Externalism and Self-Knowledge*, 57–77. Cambridge, Mass.: MIT Press.

Appendix 1

1. $P(\text{-b}(p)/\text{-}p) =$

   $P(\text{-b}(q)/\text{-}p)P(\text{-b}(p)/\text{-b}(q).\text{-}p) + P(\text{b}(q)/\text{-}p)P(\text{-b}(p)/\text{b}(q).\text{-}p)$

   Total Probability

2. $P(\text{-}q/\text{-}p) = 1$ — *q* logically implies *p*

3. $P(\text{-b}(q)/\text{-}q.\text{-}p) \geq P(\text{-b}(q)/\text{-}q)$ — No-Freefall Condition

4. $P(\text{-b}(p)/\text{-b}(q)) \geq z \quad .5 < z \leq 1$ — Basing I

5. $P(\text{-b}(p)/\text{-b}(q).\text{-}p) = P(\text{-b}(p)/\text{-b}(q))$ — Basing II

6. $P(\text{-b}(q)/\text{-}q) \geq s \quad .5 < s \leq 1$ — *S* is sensitive to *q*

7. For all *A*, $0 \leq P(A)$ — Definition of probability

8. $P(\text{-b}(q)/\text{-}q.\text{-}p) = P(\text{-b}(q)/\text{-}p)$ — 2

9. $P(\text{-b}(q)/\text{-}p) \geq P(\text{-b}(q)/\text{-}q)$ — 3, 8

10. $P(\text{-b}(p)/\text{-}p) \geq$

    $sP(\text{-b}(p)/\text{-b}(q).\text{-}p) + P(\text{b}(q)/\text{-}p)P(\text{-b}(p)/\text{b}(q).\text{-}p)$ — 1, 9

11. $P(\text{-b}(p)/\text{-}p) \geq$

    $sz + P(\text{b}(q)/\text{-}p)P(\text{-b}(p)/\text{b}(q).\text{-}p)$ — 5, 10

12. $P(\text{-b}(p)/\text{-}p) \geq sz$ — 7, 11

Appendix 2

For the multiple-premise case, the subject's belief in *p* is *based* on her beliefs in $q_1, \ldots, q_n$ only if

$P(\text{-b}(p)/ (\text{-b}(q_1)\vee\ldots\vee\text{-b}(q_n))) > z$ *Basing* I

$P(\text{-b}(p)/(\text{-b}(q_1)\vee\ldots\vee\text{-b}(q_n))\wedge q')) = P(\text{-b}(p)/(\text{-b}(q_1)\vee\ldots\vee\text{-b}(q_n)))$ *Basing* II

Here we will work through the specific case of two premises, where a subject bases her belief in *p* on her beliefs in q_1 and q_2, the latter together logically imply the former, and she knows that. We assume that she is sensitive and adherent to the premises q_1 and q_2 because in the multiple-premise case sensitivity to the conclusion can't be insured without adherence to the premises. By the total probability rule her sensitivity to the conclusion can be written:

$P(\text{-b}(p)/\text{-}p) = P(\text{-b}(p)/(\text{-b}(q_1)\vee\text{-b}(q_2))\wedge\text{-}p)P((\text{-b}(q_1)\vee\text{-b}(q_2))/\text{-}p) + P(\text{-b}(p)/(\text{b}(q_1)\wedge\text{b}(q_2))\wedge\text{-}p)P((\text{b}(q_1)\wedge\text{b}(q_2))/\text{-}p)$

The second summand will not add much for putting a lower bound on sensitivity because it will go roughly as (1-*v*)(1-*s*)(1-*s*), but the term doesn't hurt either since it is not negative. Focusing on the first term, we can expand it, by total probability applied to its second term, P((-b(q_1)∨-b(q_2))/-*p*), to get:

P(-b(*p*)/-*p*) ≥ P(-b(*p*)/(-b(q_1)∨-b(q_2))∧-*p*)[P((-b(q_1)∨-b(q_2))/(-q_1∨-q_2))P((-q_1∨-q_2)/-*p*) + P((-b(q_1∨-b(q_2))/(q_1∧q_2))P((q_1∧q_2)/-*p*)]

P((q_1∧q_2)/-*p*) is zero, leaving:

P(-b(*p*)/-*p*) ≥ P(-b(*p*)/(-b(q_1)∨-b(q_2))∧-*p*)P((-b(q_1)∨-b(q_2))/(-q_1∨-q_2))P((-q_1∨-q_2)/-*p*)

P((-q_1∨-q_2)/-*p*) is 1, leaving:

P(-b(*p*)/-*p*) ≥ P(-b(*p*)/(-b(q_1)∨-b(q_2))∧-*p*)P((-b(q_1)∨-b(q_2))/(-q_1∨-q_2)) *

We have succeeded in reducing the question about sensitivity to p to questions about sensitivity to q_1 and to q_2, adherence to q_1 and to q_2, and the basing relations between beliefs in q_1 and *p*, and q_2 and *p*.[25] The subject's sensitivity to *p* depends, as we should expect, on which of q_1 or q_2 is more likely to be false if *p* is false, and on how likely she is to not believe q_1 and not believe q_2 when each of them is false and when each of them is true, and on how likely she is not to believe *p* when she does not believe q_1 and when she does not believe q_2.

Both conditional probabilities in * have disjunctions in their conditions, which makes it easy to be misled. We know that in deductive logic $A \supset C$ and $B \supset C$ together imply (A or B) $\supset C$. If *A* yields *C* and *B* yields *C* then if you have at least one of *A* and *B* you'll have *C*. But the analog does not hold in probability. The following set of conditions

P(*C*/*A*) > *x*

P(*C*/*B*) > *x*

P(*C*/(*A* ∨ *B*)) < *x*

is a coherent assignment of probabilities, a phenomenon called Yule-Simpson Reversal. (Malinas and Bigelow 2016) We can avoid such reversals in our case by making the linking conditional probabilities that go to deciding the values of the conditional probabilities with disjunctive conditions, the same for all possible branches. That is, sensitivity and adherence to q_1 is assumed to have the same minimum as sensitivity and adherence to q_2, respectively, and if in addition we set *s* = *t* then

P(-b(q_1)/-q_1) > *s*

P(-b(q_2)/-q_2) > *s*

P(b(q_1)/q_1) > *t*

P(b(q_2)/q_2) > *t*

[25] The role of adherence is less obvious but it factors in evaluating P((-b(q_1)∨-b(q_2))/(-q_1∨-q_2)) because -q_1∨-q_2 may be true when q_1 is true or when q_2 is true.

together imply that the minimum sensitivity that $-q_1 \vee -q_2$ yields does not depend on which of its disjuncts is more likely to be making it true. Similarly the right hand sides of these expressions being equal

$P(-b(p)/-b(q_1) > z$

$P(-b(p)/-b(q_2) > z$

implies that the minimum sensitivity to p that $-b(q_1)$ or $-b(q_2)$ yields does not depend on which of its disjuncts is more likely to be making it true.

$P(-b(p)/-p) = P(-b(p)/(-b(q_1) \vee -b(q_2)) \wedge -p)P((-b(q_1) \vee -b(q_2))/(-q_1 \vee -q_2))$ *

There are three possible ways for $-q_1 \vee -q_2$ to be true, and for each of those, four possible permutations of the truth values of the conjuncts in the conjunction $b(q_1) \wedge b(q_2)$. If $s = t = z = .95$, then greater than or equal to $st = .90$ of the twelve possibilities have either $-b(q_1)$ or $-b(q_2)$ true. Of these, at least .95, i.e. z, have $-b(p)$ true because $b(q_1)$ and $b(q_2)$ are bases for $b(p)$. In all, $P(-b(p)/-p) \geq stz = (.95)(.95)(.95) = .86$. If $s = t$, then n-premise one-step inferences have conclusion beliefs with sensitivity $\geq zs^n$.

This inequality also depends on assuming that $b(q_1)$ and $b(q_2)$ are independent given $-q_1$, given $-q_2$, given q_1,and given q_2 That is, for this way of avoiding freefall in multiple-premise cases, we need to add the conditions:

$P(-b(q_1)/-q_1 \wedge \pm b(q_2)) \geq P(-b(q_1)/-q_1)$ and

$P(-b(q_2)/-q_2 \wedge \pm b(q_1)) \geq P(-b(q_2)/-q_2)$

$P(b(q_1)/q_1 \wedge \pm b(q_2)) \geq P(b(q_1)/q_1)$ and

$P(b(q_2)/q_2 \wedge \pm b(q_1)) \geq P(b(q_2)/q_2)$

Your believing or not believing one of the premises should not reduce your sensitivity or adherence to the other. These conditions, together with the condition above that $s = t = z$, are the *Multiple-Premise No-Freefall Conditions* (MPNFC) for a single step of deduction. The multiple-premise, multiple-step case is left for further work.