in *The Sensitivity Principle in Epistemology*, Cambridge University Press, 2012, 242-268.

Sensitivity and Closure

Sherrilyn Roush

(12,637 words)

I. Introduction

From the mid-1980's to the early 2000's the wide-ranging resources of the concept we now call *sensitivity*, which Robert Nozick used to give an analysis of the concepts of *knowledge* and *evidence*, went largely unappreciated in epistemology. This was in part because these resources were upstaged by a glamorous implication the condition has for skepticism, and in part because of loss of faith in the project of giving a theory of knowledge at all, due to the failure time and again to construct a theory without counterexamples. The sensitivity condition, or as Nozick called it the variation condition, which requires that were p to be false you wouldn't believe it, had its own apparent counterexamples. And while the implication of this condition for skepticism was elegant and principled – it is possible to know that there is a table in front of you *without* knowing you are not a brain in a vat – it had the price of denying closure of knowledge under known implication, that is, denying that knowing q and knowing that q implies p are together sufficient to make the belief in p that you have on that basis knowledge. Many felt this was too much to pay for what seemed to be the sensitivity condition's primary selling point.

However, the sensitivity condition need not prevent closure if it is not taken as a necessary condition for knowledge. Drawn to that move's exciting implications about skepticism, no one gave an argument for taking the condition as necessary. Conditions can figure in lots of different ways in the definition of a concept. Sensitivity was already imagined as only one of a set of necessary conditions for knowledge of p, each independent of the others, including belief in p, the truth of p (and in Nozick's theory adherence to p). Why should a definition have exactly this many but no more independent clauses? One might reply that closure should not be a property that is *independent* of the property of knowledge that goes beyond truth and belief; it should follow from that property. The weakness of this reply is that it is grounds for rejecting other theories too; few if any of the major theories of knowledge of recent times have that feature.[1]

Alvin Goldman was aware of this fact about his process reliabilist condition on justified belief (and thereby knowledge) from the inception of his theory. That one's true belief in q was formed by a reliable process, and that one's valid deduction of p from q was reliably formed, do not together imply that one's belief in p so formed was formed by a reliable process. This is because reliability of a process is not defined as requiring infallibility, even for deduction, and wherever one sets the threshold, two processes that are just above the threshold can concatenate to one process whose reliability dips below it, due to the presence of double the sources of error. For this reason Goldman did what any rational person who subscribes to the closure of knowledge or justified belief would;

he introduced a recursion clause, allowing that one has a justified belief if either one's belief was formed by a reliable process or it was formed by application of a conditionally reliable process to a justified belief (Goldman 2008, 340-341).

In internalist views of justification (and potentially, thereby, knowledge) the conditions imposed on the concept(s) are sometimes not defined explicitly enough to deductively imply a verdict one way or the other on closure, or are explicitly defined and do not have an implication one way or the other. (I will deal with Tony Brueckner's example of an internalist view that does imply closure below.) Intuitive arguments, of which there are many, can be taken to provide reasons to add closure as an independent requirement in an internalist view of knowledge, but those intuitions are available to externalists too. That sensitivity does not imply closure, and taken as a necessary condition easily brings failure of closure, is not a good argument for its being less adequate than other conditions.

In other places (Roush 2005, 2009, 2010a) I have developed, through examination of a number of issues**,** what I see as the extensive explanatory resources in the concept of sensitivity when it is combined with adherence. Here I will focus on the consequences of the move I made of combining sensitivity (and adherence) with closure via a recursion clause. My imposition of closure on a sensitivity-based view of knowledge has seemed to some unexplanatory and to lead to cheap knowledge. I will argue that as it stands the view is no less explanatory, and leads to no more cheap knowledge, than other views. However, my main objective here is to explain why these issues depend very much on the formulation of the closure clause and to present a new formulation. The new view does not have problems with explanation or cheap knowledge, and I will use it to explain how the problem of closure is entirely provoked and resolved by attention to the growth of potential error.

My theory of knowledge is a refiguring of Nozick's theory that uses probability rather than counterfactuals, and in which the sensitivity condition is neither a necessary nor a sufficient condition. Schematically, it is sufficient for subject S to know p, a contingent proposition, if p is true, she believes p, and:

1) $P(-b(p)/-p) > s$, where $s \leq 1$ and

2) $P(b(p)/p) > t$, where $t \leq 1$,

That is, the probabilities are high that she does not believe p given that p is false and that she does believe p given that it is true. s and t are thresholds determined by the disutility of false positive and false negative errors respectively for the one who is evaluating whether S knows that p, who may or may not be S herself. "≤" rather than "<" accommodates the possibility of an evaluator who has infallibilist utilities. For anyone living outside bizarro world, s and t will be greater than .5.[2]

There are several reasons why this formulation of sensitivity in terms of conditional probability is superior to the counterfactual version requiring that if p were false the subject wouldn't believe p. Typically a similarity relation is used to determine which set of possible worlds matters in the evaluation of a counterfactual; we want to know whether the consequent is true in the world(s) where the antecedent is true that are most similar to the actual world. With sensitivity this means we want to know whether the person believes p in the –p world(s) most similar to the actual world. Defining a similarity relation in a principled way so as to match the expectations of an epistemic sensitivity property has been elusive, the best accounts giving a list piecemeal, and even if one achieved a match it would be hard to avoid an air of arbitrariness, since everything is similar to everything else in some respect. However, with a conditional probability approach there is no restriction on which sets of values for other propositions that are compatible with –p, that is, which possible –p scenarios, are taken into account. All go into the weighted average that determines the value of the conditional probability in 1. Still these sets of –p-compatible values for other propositions only contribute to the evaluation in proportion to their probability; only the *probable* –p scenarios make a significant difference to the value of the conditional probability.

We thus eliminate the need for an extra similarity relation to carefully carve out the set of scenarios that matter; they all do. Yet the impossibility of knowledge does not follow, since not only do less probable scenarios make a smaller difference, but also because of the threshold in condition 1, –p scenarios whose probability is below 1-s, s<1, will not make any difference at all to whether one has knowledge. It is only if we are infallibilists (with s = 1) that no scenario's probability is too low to prevent its having an effect. Not only does this seem less artificial than trying to draw a line around a certain set of possibilities just so, but it seems to me that similarity is not the gist of what makes a possibility matter epistemically. Imagine you know there is a tree in front of you because you are looking at it. Now imagine a possible scenario in which a very special, tiny, cosmic ray hits the tree. It vaporizes the tree and stops to sit there in the space that is left behind. Moreover, this cosmic ray has properties that sometimes make it look much larger than it is. In fact, in the possible world we are considering it looks just like a full-sized tree when viewed from your distance.

The laws of physics would not need to be different, or breached, for this cosmic ray to be possible. The laws are probabilistic and allow for highly improbable occurrences. In fact the possible world described is only negligibly different from the actual world in the matters that go beyond p. It is surely much more similar to the actual world – not only in the eyes of a beholder of the tree, but also plain physically – than is a world in which the tree is not there because landscapers come with big equipment and a truck and cut it down, dig it up, and haul it away. I think we would expect you to refrain from believing there is a tree there in a world in which it has been hauled away. But I think we would also say the possibility of a cosmic ray that looks like a tree does not undermine your knowledge that there is a tree in front of you when you are looking at it. The reason, I think, is that the world described is too *improbable* to matter. In judging whether we have knowledge we seem to care more about what our belief behavior would be in those possible scenarios we have a good chance of bumping into, than about whether the possibilities are similar.

Neither of conditions 1 and 2 is necessary for knowledge because on the new tracking view defined in Roush (2005), one can know p not only by tracking p (fulfilling 1 and 2), but also by believing p and knowing something that one knows implies p (and basing one's belief in p on all of that). Thus, schematically, S knows p, a contingent proposition, if and only if S believes p, p is true, and either:

> S tracks p (fulfills 1 and 2)
> or
> there are $q_1, \ldots, q_n$, such that $q_1, \ldots, q_n$ together imply p, S knows that $q_1, \ldots, q_n$ imply p, and S knows $q_1, \ldots, q_n$.[3]

This is not yet the definition of knowledge, since the conditional probabilities in 1 and 2 are not determinate for reasons I will discuss below in connection with the generality problem (hence the persistent qualifier "schematically"), but it gives the idea for current purposes. I have also since changed the recursion clause in this definition, for important reasons I will discuss below. The discussion up to that point is independent of those changes.

II. Knowledge Of Logical Truth

Obviously a definition of knowledge of logical implication is needed to complete the recursion clause and the (schematic) definition of knowledge. The sensitivity condition has the well-known problem of being undefined for logical truths because their negations are impossible, and counterfactuals with impossible antecedents are prima facie undefined. The problem is hard to fix, and it does not go away in the re-formulation by conditional probability. However, the idea of sensitivity is broader than this. It expects that when we have knowledge we are being *responsive* to the way things are. Necessary truths are different from empirical truths – at least they are being treated as radically different if we take their negations to be impossible and their probabilities to be 0, the only circumstance in which this problem arises in the first place – so it should be no surprise if the kind of responsiveness we should expect for knowledge of logical truths and other necessary truths takes a different form.

The truth value of a logical truth is the same no matter what, namely true, and this has trivialized several otherwise promising theories of knowledge (see Roush 2005, 134-36). For example, the necessity of a logical truth means that you wouldn't, indeed couldn't, easily believe it and be wrong; a belief in it is automatically *safe*. However, a logical truth's being true no matter what does not imply that a belief in it is responsive to it, even at the intuitive level. One can believe a proposition without proper appreciation of its truth.

What could appreciation or responsiveness to a logical truth be when the truth value is the same no matter what? One might think that it would simply be believing it no matter what, and that is not a bad idea. Consider that this condition is not trivialized by the necessity of the logical truth. The facts that you actually believe p and that p is true no matter what do not imply that you believe it no matter what. The latter is requiring not just a belief but a disposition to believe, the fact that a certain belief would be present in a large set of possible scenarios. Believing no matter what is not my analysis of responsive to logical truths, but it is in the neighborhood.

What we need to appreciate about logical truths is the way that they impose relations between propositions, and thereby impose requirements on the relations between our beliefs. Your belief states toward propositions should comply with those relations even when they are not determined by them. For example, if contingent q implies contingent p and someone believes q, but were he to form a belief in q he would still not believe p, then we would have to say that he does not know that q implies p. If, were he to give up belief in p he would retain belief in q, he surely would also fail to know the implication. Implication is a relation, so to appreciate a logical truth that is an implication claim your belief states toward the propositions on either side of that relation must be responsive *to each other*. Thus, I say the key extra feature beyond true belief in knowing that $q_1, \ldots, q_n$ imply p is fulfillment of the following conditions:

3') $P((-b(q_1) \vee \ldots \vee -b(q_n))/ -b(p)) > u \leq 1$

4') $P(b(p)/(b(q_1) . \ldots . b(q_n)) > v \leq 1$

where u and v are thresholds that depend on the error tolerances of the evaluator, as above. If you know the implication, then by 3' the probability that you do not have one of the q_i beliefs given that you do not have a belief in p is high. Similarly, if you know the implication then the probability you believe p given that you have beliefs in all of q_1 through q_n is high.

Tony Brueckner objects to this definition on the grounds that a subject's fulfillment of these conditions merely "attest[s] to the firmness of his belief that the implication holds" (Brueckner, "Roush on Knowledge: Tracking Redux?" this volume, ##; hereafter 'RKTR'). But this is a false description of the conditions, for they are not properties of the subject's belief in the implication at all. In fact the belief in the implication and properties 3' and 4' are logically independent, for a person may behave as if the implication is true, withdrawing and according credence to some of the p's and q's according as his credences in others of them change, without having a belief in the proposition "q_1, …, q_n imply p," and vice versa. If a belief is a disposition to act – my preferred way of thinking about it – we could say that the subject's patterns of willingness or unwillingness to act on p when he is willing or unwilling to act on q_1, …, q_n are distinct from his willingness or unwillingness to act on the claim *that* q_1, …, q_n imply p. For a simple example, imagine someone who does not have the concept of logical implication.[4]

Brueckner's example of the supermarket tabloid trades on the same confusion. He imagines that I have what he calls a "firm" belief that the implication holds, that I got from a screaming *National Enquirer* headline, and takes firmness of that belief to be the same as satisfaction of 3' and 4'. But this cannot be so because 3' and 4' are not properties of the implication belief. "Firmness" intuitively suggests that I have a certain stubbornness about the implication claim, but I can have that without acting in accord with the relations among beliefs in q_1, …, q_n and p, respectively, that that implication dictates. In fact, getting a stubborn belief in the implication from the headline alone as suggested is a good way of imagining not properly acquiring the appreciation of the relations among beliefs that the implication imposes.

Brueckner says that the reason I do not know the implication claim when I acquire it from the headline is that "I fail to see that q implies p, in that I fail to see the logical connection between q and p" (RKTR, #). This assumes a requirement that the subject appreciate the relation that is claimed between q and p, with appreciation defined as seeing, presumably in the metaphorical sense. My view requires that the subject appreciate that relation, with appreciation defined as having dispositions to manage beliefs in p and q in accord with it. The only difference, then, is in how we understand the needed appreciation. Brueckner claims that there is nothing "epistemic" about the conditions I require for knowledge of logical implication, so his complaint must be based on the difference between "seeing" and differential dispositions to believe. That is, the objection must assume that counterfactual conditions on belief cannot capture what is epistemic. One is free to make that assumption, of course, but it would make the current dispute unnecessary.

Some may be uncomfortable with my externalist approach to knowledge of logical and other necessary truths, because none of these definitions require even access to reasons or understanding of why the logical truth is true. My definitions have sometimes been mis-cited as requiring the subject to be in a position to give an argument – perhaps because it is hard to believe anyone would leave that out of an account – and these intuitions may underlie Brueckner's dissatisfaction with my view. I mean what I say, though, and there are several advantages to my approach. First, though the kinds of logical and mathematical truths discovered by researchers require explicit thought and proof if they are to become known, the majority of logical implications that are known by human beings are simple ones known by the man on the street, who would not be able to give an argument for them, and might not be able to formulate the general logical rules of which his beliefs are instances. (Even the woman on the street might not be able to do these things.)[5] Yet truths such as “If A and B then A” and the independence of the continuum hypothesis are truths of the same type, and we should expect there to be something similar in what it is to know them. Leaving out from the requirements the ability of the subject to prove the claim makes that possible.

One might object that such externalist conditions as I have proposed cannot be nearly enough to explicate what it is for a mathematician to know a theorem. Such knowledge does require the proof, and the understanding that she acquires in that process. This claim appears correct but there is an ambiguity in the notion of requiring. Acquisition of knowledge requires proof in these sophisticated cases, but I am defining what it is to know, not what is required to *attain* that status. I do not take the process of coming to a belief as related in a necessary way to its status as knowledge. (Counterfactual conditions in general do not; they are current-time-slice-views, not historical views. The claim that I should have a historical view would require an argument on different grounds.) That process is obviously related in a contingent way to the status – it happens to be a psychological fact that some logical knowledge cannot be acquired by us except through explicit work – and this is enough to explain the intuition that it should be a requirement for having the knowledge. The objection is also right that proofs can yield deep understanding, but the state of understanding is different from the state of knowing.

To take the availability of argument as a requirement on all knowledge of necessary truths would leave the man on the street little if any logical knowledge, and to split the two categories – ordinary vs. sophisticated – would fail to explain their similarity. It would also have awkward consequences for classification. The prooflessness of the mathematical wonder Srinivasa Ramanujan has apparently been exaggerated – the editors of his journals from the years in India say it is clear that he had proofs, and his not writing them down was due to the need to conserve paper. But even if all of the great mathematicians eventually, or quickly, were able to give proofs of the necessary truths they discovered, a view that requires that ability for knowledge is committed to there never being any moment at which a mathematician knows a theorem without that articulation ability, never, ironically, a pure “seeing” moment. The journals aside, there is a famous anecdote about Ramanujan, that when G. H. Hardy remarked that the number on the taxicab they were sharing was uninteresting he immediately informed his cabmate that the number 1729 was the smallest natural number representable in two different ways as a sum of two cubes. Of course he subsequently was able to prove the claim, and generalizations of it have become theorems about the “taxicab numbers,” but unless we stretch the notion of ability in an ad hoc way we are liable to be denying too much knowledge.

Not all logical truths are implication claims, of course, but the relation of implication can guide us through the rest because of the distinctive fact about a logical truth (in monotonic logics), that it is implied by every proposition in the language. To know non-implicational logical truth, r, then, I say that what we need to be responsive to in our belief behavior is this special relation r has to all other propositions. One does not have to know *that* r is implied by every proposition in order to know r – most people do not even have that belief. Rather it must be that there is no proposition such that one fails to appreciate the implication *relation* it has to r. It must not be that one might believe it and *not* believe r. That is, the proper way to be responsive to a (non-implicational) logical truth is to have a disposition to believe it come what may among one's beliefs. Schematically, the conditions for this are:

c') For all $q_1, \ldots, q_n$, $P((-b(q_1) \vee \ldots \vee -b(q_n)) / -b(r)) > w \leq 1$
d') For all $q_1, \ldots, q_n$, $P(b(r)) /(b(q_1) . \ldots . b(q_n)) > x \leq 1$,[6]

for w and x thresholds of error tolerance possessed by the evaluator. (See Roush 2005 for more detail.) If a subject knows r, then in a case where the subject did not believe r he would not believe anything, and when he does believe anything, he also believes r.

This is not a requirement for responsiveness to the fact that r is true no matter what, but responsiveness to the relation r has to every proposition one might believe, and thus to the constraints on belief behavior the logical facts impose. The difference between a "come what may" phrasing and a "come what may among one's beliefs" phrasing, is that the latter does not require you to be someone who would believe r even if you believed nothing else; the former does. The difference is not merely technical. It indicates that what being responsive to logical propositions requires is not principally a relation to how the world is as regards logic, but how whatever that truth is imposes requirements on the relations among one's beliefs. This makes my view compatible with a non-realist view that logical truths are not so in virtue of the way the world is, but in virtue of something like convention, or the rules of language or the contours of human concepts. Whatever logical truths are, they impose constraints on belief of the sort I have described.

The phrase "come what may" and the implicit invocation of Quine's web of belief might prompt one to wonder whether the view can plausibly accommodate the fact, or view, that revision of logic could be justified. If in order to know the Law of Excluded Middle (LEM) one must have a disposition to believe it no matter what else one believes, then does it not follow that one could not count as knowing it if one might seriously consider an argument from L. E. J. Brouwer that this principle must be rejected?

This does not follow, for interesting reasons. The dispositions that are required of one are expressed as conditional probabilities whose values must be above a certain level, w or x. Unless the threshold chosen is 1, knowing LEM allows for a probability of 1-x that you would give up that belief given other beliefs. It is legitimate to demand that the probability of your revising, 1-x, is small if we are to count you as knowing, and there are at least two ways of achieving that. If one were a researcher who dealt with arguments about revision of logic on a regular basis then there is a high probability that one would *consider* changing one's mind, but we would expect such a person not to be a push-over for arguments for radical views, and that keeps the probability of an actual change down. A person in the street not versed in logic might be someone inclined to naïve enthusiasm when presented with radical views, and so would be a pushover when presented with an alternative logic. But he can know LEM, or instances of it, even so, provided he is unlikely to come across or understand accounts of radical logics. As one becomes more and more disposed to changing one's mind about LEM, perhaps due to increasing exposure to alternatives, one's knowledge of it slips away too, on this view, even when one's belief does not change. But this does not seem to conflict with intuitions because it is a process in which the robustness of one's commitment erodes.

A less realist view of logic, where logical truths are rather viewed roughly on the model of conventions, would bring an additional feature here. On such a view, if you are considering changing your belief in LEM, you are considering changing your language, for changing your logic at least implies this even if it is not identical to it. However, those circumstances of not believing LEM do not violate the conditions on knowing LEM. They would be scenarios in which you don't believe LEM but you don't have other beliefs either in the sense that concerns the conditional probability conditions on your current belief in LEM; what a belief is in that sense is conditioned by your current language. So c' is fulfilled. And in all of the possible beliefs as defined for your current language, including logic, you do have the belief in LEM, so d' is fulfilled. This is another place where the difference between belief come what may and belief come what may among your beliefs shows itself.

Knowledge of necessary truths is fallible—you always might have made a mistake—but it has been difficult to incorporate fallibility in other theories of this kind of knowledge (Roush 2005, 134-6) because the proposition does not have the potential to be false. The view here avoids this problem because the responsiveness one must have is not to the difference between the truth and the impossible falsity of the logically true proposition, but to the relations each of these truths has to other propositions and to the relations they impose on groups of other propositions. The required dispositions are dispositions *among your beliefs* only – the question is not what one would do if the logically true proposition were false, but what one would do with one belief given one's other beliefs or lack thereof. Most of us will take the parameters w and x to be less than 1 in the conditions above – when we take them as 1 we are infallibilists and the challenge of creating a fallibilist theory does not arise – which means that the disposition that makes you count as appreciative or responsive need not guarantee that you avoid error in all possible circumstances.

III. Closure and the Growth of Error

Achieving closure via a recursion clause, such as the one above, where the requirements for knowledge of the implication are not infallibilist but have thresholds, has a problem with the growth of error that I did not deal with in my book. Suppose that I track q to degrees s and t and am responsive to degrees u and v to the fact that q implies q'. Let s, t, u and v each be less than 1. It follows from the definition that I know q' although, because the thresholds are below 1, my potential error in my belief in q' is higher than the potential error I have in my belief in q. The growth over one step may not be large enough for worry. (We will calculate actual values below.) However, now that I know q' I can do the same thing for any q'' that I know q' implies, and this sequence can be continued indefinitely. It does not need to be carried on very far for me to be counted as knowing things for which my sensitivity is only 50%. The same is true of Goldman's recursion clause.

What causes the runaway multiplication of potential error is not the allowance of knowledge by known implication but the recursion. However, we can reformulate the relevant clause of the view (see above, #) without the recursion by changing one word:

> S tracks p (fulfills (1) and (2))
> or
> there are $q_1, \ldots, q_n$, such that $q_1, \ldots, q_n$ together imply p, S knows that $q_1, \ldots, q_n$ imply p, and S tracks $q_1, \ldots, q_n$.

The difference occurs in changing the second "knows" to "tracks" in the second clause.[7] It is now not enough to know that p follows from something that you know follows from something that you know … follows from something that you track. You must instead be no more than one implicational step away from a proposition that you track if you are to count as knowing. This avoids the problem of counting a belief with any old degree of potential error knowledge, because it turns out that any designation of the number of allowed steps of known implication from a tracking belief puts strict limits on the growth of potential error, and implies a degree of preservation of sensitivity.

Calculating error over known implications in the way that will yield such results requires explicit definition of an assumption that the subject's belief in p is *based on* her belief in q.[8] This is a standard assumption in any closure clause because otherwise her belief in p is only accidentally related to the belief whose knowledge status closure says is supposed to give her a right to it. In that case there is no good reason to think that her knowledge of q is giving her knowledge of p. I define the basing relation using the following necessary condition: The belief in p is *based on* the belief in q only if

$P(-b(p)/ -b(q))$ is high,

that is, only if you are unlikely to believe p given that you do not believe q. Fully generally,

5) $P(-b(p)/(-b(q_1) \vee -b(q_2) \vee \ldots. \vee -b(q_n))) > z \leq 1$.

If this conditional probability were as high as 1 then your beliefs in $q_1, \ldots, q_n$ would be solely responsible for your belief in p.

Note that this basing claim is not part of what is required to know that $q_1, \ldots, q_n$ imply p. Knowing that $q_1, \ldots, q_n$ imply p does not prevent you from also believing or being disposed to believe p on the basis of other beliefs that imply p. Rather, this is a claim of basing needed to formulate the closure requirements. In the closure question we are assuming that the only thing you have that might get you knowledge of p is your knowledge of $q_1, \ldots, q_n$, because the question is whether the latter is sufficient.

To summarize the whole view, again schematically with regard to conditions 1, 2, 3', and 4', S's true belief in p is knowledge if and only if

1) $P(-b(p)/-p) > s$, where $s \leq 1$, (sensitivity to p)
and
2) $P(b(p)/p) > t$, where $t \leq 1$, (adherence to p)

or
there are $q_1, \ldots, q_n$ such that $q_1, \ldots, q_n$ together imply p, S believes that they do, and for every $q \in \{q_1, \ldots, q_n\}$ S fulfills 1 and 2, that is,

$P(-b(q)/-q) > s$, (sensitivity to every premise)
and
$P(b(q)/q) > t$, (adherence to every premise)
and
S fulfills 3' and 4' for $q_1, \ldots, q_n$, that is,

3') $P((-b(q_1) \vee \ldots \vee -b(q_n))/-b(p)) > u \leq 1$,
(implication sensitivity)
and
4') $P(b(p)/(b(q_1) . \ldots . b(q_n))) > v \leq 1$
(implication adherence)

and
S fulfills 5 for $q_1, \ldots, q_n$, that is,

5) $P(-b(p)/(-b(q_1) \vee -b(q_2) \vee \ldots. \vee -b(q_n))) > z \leq 1$
(basing $b(p)$ on $b(q_1), \ldots, b(q_n)$)

or the number of steps of known basing implication is $m > 1$.

For reasons I discuss below, the number of premises and number of steps of known implication also affect the growth of potential error, so m and n also depend on the error tolerances of the evaluator in a way that we will see.

If I fulfill all of the clauses beyond "or" in the definition just given, what can we say about the sensitivity of my belief in p? What is the probability I do not believe p given that p is false? Information sufficient to determine this is given by the fact that q_1, …, q_n together imply p, my sensitivity to the q_1, …, q_n, and the basing of my belief in p on my beliefs in q_1, …, q_n. First using a single premise q for simplicity, by total probability,

$$P(–b(p)/ –p) = P(–b(q)/ –p)P(–b(p)/ –b(q). –p) + P(b(q)/ –p)P(–b(p)/b(q).–p).$$

By the fact that –p implies –q, the right hand side is greater than or equal to

$$P(–b(q)/ –q)P(–b(p)/ –b(q). –q) + P(b(q)/ –q)P(–b(p)/b(q). –q).$$

The terms here involve sensitivity to q, basing, and implication adherence, thus:

$$(>s)(>z) + (< (1-s))(< (1-v)).$$

Taking all thresholds as .95, this becomes:

$$(>.95)(>.95) + (<.05)(<.05) > .90.$$

In general, for one-step, one-premise known implications, $P(–b(p)/ –p) > s\cdot z$. Although level of adherence to the implication, v, occurs in the equation it has no impact on our question of a minimum because the second term in which v occurs is composed of maxima. If one's level of sensitivity to q is s, then one's level of sensitivity to some p that one knows follows from q is diminished from s only by the level of deviation from perfect basing. If your belief in q is truly the only basis for your belief in p, then no potential error is introduced at all; sensitivity is fully preserved over the known implication. In general, for m-step, one-premise known implications, the preserved sensitivity level, P(–b(p)/ –p), is $s\cdot z^m$. Thus, for 2 steps and all thresholds set at .95, P(–b(p)/ –p) is .86, for three steps .81.

Multiple-premise closure brings further error considerations. Each premise known fallibly contributes potential error, so the potential error in the resulting belief in p is greater than it would have been with fewer premises. The question is how much greater. Taking the next case, two-premise implication, we want to know the minimum level of sensitivity one will have to p if one is counted as knowing it by tracking each of q_1 and q_2, knowing that they together imply p, and basing one's belief in p on them. This can be calculated as follows. Because –p implies –q, which implies $-q_1 \vee -q_2$, the latter, call it A, becomes the condition in all of the conditional probabilities:

$$P(-b(p)/-p) = P((-b(q_1).-b(p).-b(q_2))/A) + P((-b(q_1).-b(p).b(q_2))/A) + P((b(q_1).-b(p).-b(q_2))/A) + P((b(q_1).-b(p).b(q_2))/A).$$

The last term is 0 since $P(b(q_1).b(q_2)/(-b(q_1) \vee -b(q_2)))$ is 0. Rewriting each remaining summand,

$$P(-b(p)/-p) = P(-b(p)/(-b(q_1).-b(q_2)))P(-b(q_1).-b(q_2)/A) + P(-b(p)/-b(q_1).b(q_2))P(-b(q_1).b(q_2)/A) + P(-b(p)/b(q_1).-b(q_2))P(b(q_1).-b(q_2)/A)$$

For the moment I will assume that when $-q_1 \vee -q_2$ is the case it is equally likely to be $(-q_1.-q_2)$, $(-q_1.q_2)$, or $(q_1.-q_2)$. The first term is determined by basing and sensitivity to each of q_1 and q_2 and equals:

$z \cdot s^2/3 = (.95)(.95)(.95)/3 = .286$[9]

The second term is determined by basing, sensitivity to q_1, and adherence to q_2, and equals:

$z \cdot s \cdot t/3 = (.95)(.95)(.95)/3 = .286$

The third term is symmetric to this one and also equals .286. Thus we have that the sensitivity to p that results from knowing it in the way imagined is at least .86. That amounts to 4% less fidelity than with one-premise (one-step) implication.

In general, for one-step, two-premise closure, for the conclusion-belief in p,

$P(-b(p)/-p) = a \cdot z \cdot s^2 + b \cdot z \cdot s \cdot t + cz \cdot t \cdot s$

where $a = P(-q_1.-q_2/(-q_1 \vee -q_2))$, $b = P(-q_1.q_2/(-q_1 \vee -q_2))$, and $c = P(q_1.-q_2/(-q_1 \vee -q_2))$. In this calculation I made the assumption that a = b = c. But it turns out that the values of a, b, and c make very little difference if z, s, and t are all high (as they will be in the cases of interest), because the terms by which a, b, and c are multiplied are then close to equal, and a + b + c = 1. For example, if in an attempt to minimize the lower bound on sensitivity of our belief in p we take a = .90, b = .05, c = .05, then as long as we keep s, t, and z the same at .95, $P(-b(p)/-p) = z \cdot s^2 = .86$, the same as when a = b = c. Thus, in the cases where if at least one of q_1 or q_2 is false one of them is vastly more likely than the other to be the false proposition, it need not make any difference to the growth of error.

For the general case of n premises and one step of known implication, and if s is approximately equal to t, the preserved sensitivity in the conclusion, p, P(–b(p)/ –p), is greater than approximately $z \cdot s^n$.[10]

What will be a problem for some is the fact that getting lower bounds on sensitivity of the conclusion-belief in the multiple-premise implication case depends crucially on t, the subject's adherence to q_1 and to q_2, as it did not in the one-premise case. The renaissance of interest in the sensitivity condition has not been accompanied by enthusiasm for Nozick's adherence condition (except in my case), but it makes a difference in minimizing the growth of error. If one endorsed sensitivity but not adherence, then one would not have general error-control with more than one-premise closure.

IV. Explanation Of Closure

Brueckner objects to the recursion clause that I used in the book to impose closure, that it does not explain *why* knowledge would be preserved under known implication. It just stipulates that it is. However, judging by the standard Brueckner sets with the examples of views he thinks do bring with them an explanation of closure, I have to disagree. He has us consider an evidentialist view:

> I have good evidence e for q, and I correctly deduce p from q. On what basis do I then know p? An evidentialist can answer that if one has adequate justifying evidence E for φ, and φ implies ψ, then E is also adequate justifying evidence for ψ. This explains how it is that I can know p on the basis of my deduction of p from q: my evidence e for q puts me in a position to know the implied p. (RKTR, #)

If E is evidence for φ, and φ implies ψ, then E is also evidence for ψ. Why should we think this? What is the understanding of evidence that would make it true? The question is not trivial since the definition of evidential support as probabilistic relevance makes it false. The claim is called the special consequence condition, and it is arguably false even on an intuitive notion of evidential support, as illustrated by examples of φ that are conjunctions with independent conjuncts.[11] Even if there is a notion of evidence that makes this claim true, though, it is not obviously so. Citation of a claim very much in need of defense and explanation is hardly an explanation.

I agree with Brueckner that Goldman's reliabilist view is explanatory of closure. Probably the reason Brueckner thinks this is that the notion of a reliable process figures in the recursion clause. Although deduction does not preserve reliability level it is itself a (conditionally) reliable process, so even with the addition of the recursion clause justified belief remains predicated upon the concept of reliability. Where I disagree with Brueckner is his implied denial that my view has the same explanatory structure. Just as conditional reliability of a process is not the same as reliability of the process but is conceptually analogous, responsiveness to logical implication is not the same as but is conceptually analogous to sensitivity as defined for empirical truths. They are both responsiveness to truths in the way that the nature of the truth in question calls for. Why should two concatenated responsiveness relations yield a relation we should take as making a belief knowledge when the level of responsiveness of the concatenation will be somewhat degraded? Why should we have taken the analogous concatenation of reliable and conditionally reliable processes to preserve knowledge, as Brueckner seems to think we should, when the level of reliability was degraded?[12]

Thus, my original view had no more problem explaining closure than other views have. It had no more of a problem of endorsing cheap knowledge either. It counted me as knowing I am not a brain in a vat provided only that I tracked my hand and knew that this implied I was not envatted. But the principle Brueckner gives us that ensures evidentialist justified belief is closed has the same problem: It implies that having evidence that one has a hand combined with valid deduction of unenvattedness from the evidence gives one evidence that one is not envatted, because according to the principle, the evidence that you have a hand *is* evidence for what follows from your handedness. It does not seem that any of us has evidence that he is unenvatted.[13]

All that said, my view, like the others, did have problems in these areas. The claim I want to defend now is that the new closure clause resolves both of them in one stroke. My original closure clause (and Goldman's) suffered from the uncontrolled error problem, and this is the source, I think, of the sense that closure has been imposed but not explained. But an uncontrolled error problem is also the source of puzzlement about how knowledge itself could be closed. Why else do we pause at the idea that we know it is not a cleverly disguised mule merely by knowing that it is a zebra, except that what gives us apparent knowledge that it is a zebra does nothing to protect us against the error entailed by the presence of a cleverly disguised mule? Similarly with cheap knowledge; why else would we be surprised that we could know we are not envatted by knowing that we have a hand than that what makes us count as knowing the latter does nothing to protect against the error that envattedness would represent?

Because my new template for closure clauses addresses the error problem, it also relieves us of the other two problems in a novel way. Clauses of the new sort will make a theory closed in one sense, but not in another. Starting with the latter, the change from "knows" to "is sensitive to" means that the following standard formulation of closure is not fulfilled:

> If S knows q, knows that q implies p, and believes p on the basis of these things, then S knows p.

If we accept one-step closure, though, with thresholds matching our error-tolerances for knowledge, the following weaker kind of claim is true:

> If S is sensitive to q, knows that q implies p, and believes p on the basis of these things, then S knows p.

The stronger claim can be achieved by a recursion clause only on pain of the snowballing error problem. However, when intuitive arguments are given in favor of closure, the supporting examples involve no more than one known implication and are not direct evidence of a recursion. This and the fact that the growth of potential error is controlled in the new view, makes this new kind of closure clause seem to me to capture well what is intuitively right about the claim that knowledge is closed. It is a bonus that we now also know how to define clauses allowing m steps and multiple premises in a way that maintains a specified maximum of potential error and minimum of sensitivity in the conclusion belief.

The consequences of this shift in attention are distinctive. I have not given a closed theory of knowledge with sensitivity as a necessary condition. However, it follows from what we have just seen that for every level of sensitivity, y, there is a disjunctive theory with a sensitivity condition using a threshold greater than y and a one-step, single-premise closure clause, that will ensure that no belief whose sensitivity is below y will count as knowledge. The condition that tells you how to find such a theory is: $s \cdot z > y$. There are many such theories, with varying values for s and z, and there are also all of the theories for multiple-step, single premise known implications where $s \cdot z^m > y$, and similarly for multiple-premise known implications where $z \cdot s^n > y$. For every theory fulfilling my conditions with sensitivity as a necessary condition, and so with knowledge not closed, there is a theory fulfilling my conditions that is at least as strong in the sensitivity requirement and according to which knowledge *is* closed in the weaker sense stated above.

One can make knowledge as expensive as one likes with regard to sensitivity, and still maintain a kind of closure. But have we have trivialized what it is for a theory to take knowledge as closed? The point of adding a closure clause to a theory by brute force was to make the knowledge concept more permissive. But if a definition requires a belief to have sensitivity > s in order to count as knowledge, then we add nothing to the extension of the concept if we disjoin this condition with one describing a belief that the subject inferred from a belief that has sensitivity $> r = s/z$. That belief by inference counts as knowledge because we chose the threshold r so that the resulting belief would achieve sensitivity >s! The ability to define and calculate error control over steps of inference does not need to change the extension of the concept of knowledge; what the bounds on the growth of error allow us to do is identify a multiplicity of types of inferences –

permutations of values of s, t, m, n, and z – that will yield knowledge of a given sensitivity.

This analysis does somewhat undercut the interest of the closure problem in the form of a yes-no question – closure or not? – but that is a side-effect of addressing the problem that I think makes the yes-no question persist. Examples can be found where known implication does seem to preserve knowledge and where it disastrously fails, and we are at a loss to see a principled way of identifying the former without stepping into the abyss of the latter. Both sorts of examples are logical implications, after all. I say the principled way to draw lines is in terms of potential error, and I have shown how to define closure clauses so that they won't yield the unpleasant surprise of a forced choice between cheap knowledge and skepticism. If that problem disappears though, one might wonder, could these *really* be closure clauses? I say they are because they explain to us how and why we can accept knowledge by inference not piecemeal but across the (error-defined) board of our choice. They give us what is essential to the idea of closure, namely, that we can accept a case as knowledge by known implication from a sensitive belief while knowing nothing more than that error thresholds were met on the initial belief and the basing relation. And we won't have any surprises about the sensitivity of the conclusion-belief.

V. Logical Implication: Tough, Not Cheap

Brueckner's complaint about the bruteness of my imposition of closure occurs within an objection to the consequences of my view for brain-in-a-vat skepticism. The premise-belief of the skeptical argument – I have a hand, say – is sensitive, but the conclusion-belief – I am not a brain in a vat – is not. Why should we think that the latter is knowledge? With the results I just derived about error the skeptical case becomes even more puzzling. How can it be that sensitivity is largely preserved in one-step deductive inferences, as I just argued a priori, when in the skeptical case we get in one step from a premise and an inference with as much sensitivity as you like to a conclusion with maximum potential error?

In fact neither my original recursive closure clause nor my new closure template implies that it is possible to know you are not a brain in a vat just in virtue of knowing you have a hand and knowing that unenvattedness follows. This is because in this skeptical case the familiar implication claim is false; false claims cannot be known so the premises of the closure clauses are not fulfilled.

That I have a hand does not imply that I am not a brain in a vat because I could be a brain in a vat with a hand. One might think this is cute – smart undergraduates say it regularly – but merely a verbal trick. Of course one might still be a brain in a vat with a hand. The point was that one could not be a *handless* brain in a vat if one has a hand. However, though that does give us an implication it is not one that holds any surprises that a skeptic could exploit. If you know that you have a hand then you know that you are not a handless anything. I claim that any attempt to repair the lack of implication between these two claims is doomed to fail, and have argued this elsewhere (Roush 2010b). Strengthening the antecedent sufficiently to get the implication means that supposing we know it is supposing we know a lot, plenty enough, on intuitive grounds independent of commitments about closure, to know we are not brains in vats. Weakening the conclusion to where it is implied by having a hand similarly gives us no surprise that a skeptic could take advantage of; knowing what the deduction lets you know in that case is not knowing much. The lack of implication in the skeptical case means that my closure clauses do not imply something that they cannot explain.

Logical implication is mercifully unforgiving, something I noticed in the context of skepticism only after writing the book, but had noticed in the book in the context of lottery propositions. That I will not have large sums of money tomorrow – something people think I know – doesn't imply that I won't win the lottery today. I could win and promptly be robbed, in whatever sense is needed to relieve me of money in an electronic age – perhaps the girl with the dragon tattoo hacks into the lottery company's account. The usual response to this point is a casual remark to the effect that an example that did involve implication could be constructed. However, I argued in the book that this is not so, because of the same kind of trade-off just described for the skeptical case. Brueckner agrees about the trade-off for the lottery case, so I will leave it at that.

Brueckner wrongly thinks there is an asymmetry, though, between my treatment of literal lotteries, and the kind of lotteries discussed by Jonathan Vogel (1990). Surely I know that my car is parked in the F-lot – I parked it – but I do not know it wasn't stolen while I was sitting in my office for these many hours, even though the latter follows. The same point is true with this case as with the lottery: The implication does not hold because, for example, my car could still be in the F-lot while having been stolen, re-stolen, and returned to its original place. I don't have to know my car wasn't stolen in order to know it is parked in the F-lot, because the former does not follow from the latter. Similarly, I don't have to know I won't win the lottery in order to know what I do seem to know, that I won't have large sums of money tomorrow, because that I won't have lots of money tomorrow doesn't imply that I won't win the lottery today. I do know that my car was *probably* not stolen, though, because this I track. The probable scenarios in which it was probably stolen – such as that the theft rate is very high in the F-lot – are also scenarios where I wouldn't believe it was probably not stolen. I claim the analogous thing in the literal lottery. Intuitively, I do know that I will probably not win the lottery, and the sensitivity condition explains this: If I were not probably going to lose I would not believe I was probably going to lose.

Logical implication is tough. It must be that there is no possible circumstance in which the premises are true and the conclusion false. It is distinct from any induction with however so many instances one might induce from, and, as I made clear in the book, I do not allow a closure clause for induction. Brueckner finds the implications of this for the bootstrapping argument strange. I find his objections unpersuasive.

The familiar bootstrapping argument (Vogel 2000) has me sitting in the driver's seat of my car with no reason to trust my gas gauge. I engage in the follow procedure: I look at my gauge and form the belief that it has a certain reading, say "F", and on the basis of what it says I form the belief that my gas tank has the amount of gas the reading indicated. I conjoin those two beliefs. I do this, and nothing else, n times, and now I have a list of beliefs in conjunctions. As Brueckner writes it:

1. My gauge says 'F' & the tank is full.
2. My gauge says '1/2 F' & the tank is half-full.
3. My gauge says 'F' & the tank is full.
4. My gauge says '1/4 F' & the tank is one-quarter-full.

.
.
.

n. My gauge says '3/4 F' & the tank is three-quarters-full.

From 1 I infer that my gauge was accurate on occasion 1, since a match between what it says and what is the case is what accuracy amounts to. I do this for each of 1 through n and now have many beliefs of the form "My gauge was accurate on occasion m," for many distinct m. I infer from all of these instances that:

C: My gauge is reliable.

Clearly this inference does not yield knowledge of its conclusion. I claim that I do not know C in this way because I do not track it: If my gauge were not reliable I would still believe it is because of the silly procedure I am using. If the gauge reading often did not match the gas level in the tank, that is, if it were not reliable, I would still believe it is reliable. A lack of sensitivity is what is wrong with the fact that I am not checking the accuracy of the gauge except by consulting the gauge.

Brueckner finds this strange, apparently because I am sensitive to the conjuncts in the premises[14] and we can make n as large as we like; what then is the reason for distinguishing this inference over many instances from a bona fide deduction and applying the closure clause? He compares the argument here that I say does not yield knowledge of the conclusion to an argument of a neighbor that could be written down with the same sentences but in which he gets his beliefs by checking the actual level of gas in the tank when he comes to believe the second conjunct in each premise. Brueckner finds it strange that there are two arguments with exactly the same conclusion beliefs and premise beliefs and that it follows from my view that my argument does not yield knowledge whereas my neighbor's does yield knowledge. I would have thought that is exactly the consequence one would want a view to have about these cases, since the neighbor *does* know and I do not. It is evident from the example itself, and not some quirk in my view, that the reason for this difference will not be found in the sentences written on the paper.

Intuitively, the reason the neighbor knows that the gauge is reliable is that he has checked the level of the tank through another means than the gauge. This feature of his procedure also has the implication that he is *sensitive* to the reliability of the gauge. In case the gauge is not reliable chances are good – and better the higher the n – that mismatches will show up between what the gauge says and what the neighbor comes to believe is the level of the tank, because he is investigating the latter by an additional method. The neighbor knows because he tracks. I don't know because I don't track.

It is very neat when the consequences of a theory match intuitions, but am I arbitrarily distinguishing an induction of n instances – n very, very large – from a deduction? What is the big difference between deduction and a very strong induction? I think a closure clause for induction must be rejected because of the same issues about growth of error that I dealt with above in re-formulating the closure clause for deduction. There we were able to derive strict upper bounds on the amount of error that taking a merely known-to-be-implied belief as knowledge would give us. With induction we will not be able to derive a reasonable upper bound on error.

One might think that the only difference between the deductive and inductive cases is the strength of the support of the conclusion by the premises. Every term in the evaluation above depended on the fact that q implies p in order to change the condition –p in P(–b(p)/ –p) to the condition –q, so that the sensitivity to q could do its work. But we would expect of a good induction that its evidence, q, made its conclusion, p, probable, say P(p/q) > .95, so would we not just change a 1 for a .95 in every term in that move from –p to –q? Changing 1 to .95 even in four terms would not make a devastating difference. However, this is not the only change that would have to happen. The problem is that an induction must use many pieces of evidence if it is to establish a strong support relation, and although we saw that having two premises rather than one reduced the sensitivity we got for p only by 3-4 percent with the deductive closure clause, 30 pieces of data would make for 30 premises and bring us to no sensitivity at all. And 30 is a small data set.

But suppose for a moment that our thresholds were high enough that a 30-premise implication would not lose more than 5% sensitivity, and suppose that it was possible for human beings to be squeaky clean enough to fulfill those thresholds with the right kind of effort. An inductive closure clause would still not count the inference to the reliability of the gas gauge as giving knowledge. A key part of the reason that sensitivity transmits to a degree over known deductive implication was the sensitivity the subject is required to have to each of the premises. The inductive closure clause would have this requirement too. However, the premises of the inference to the reliability of the gauge are claims of accuracy of the gauge, and the subject is not sensitive to such claims. If the gauge were not accurate on a given occasion, that is, if the reading the subject saw did not match the actual gas level, the subject would still believe it was accurate, since she is not checking the actual level. This is true for all of the occasions, and so, premises.[15]

An inductive closure clause would be either pointless – because one or two premises rarely give strong support – or disastrous – because a sufficient number of premises, and feasible thresholds, would leave us without any sensitivity to the conclusion proposition. And it wouldn't help to undercut my conclusions about the gas gauge in any case.

There is another apparent problem of cheap knowledge by known implication that my view does not fall prey to. Conjunctions appear problematic because the conjuncts may have properties as different as you like. In a case where if the conjunction were to be false the most similar worlds all have the one false and the other true, and the one that would be false the subject tracks but the other conjunct she does not track, then she counts as sensitive to the conjunction without being sensitive to one of the conjuncts. If this is combined with a closure clause, then it looks like she can know the conjunct she is not sensitive to at all by knowing it is implied by the conjunction (Roush 2005, 110-112). The same kind of argument can be imagined against conditional-probability tracking, so it is important to see how my views avoid this.

Obviously, it follows from the results I explained in the previous section that cheap knowledge cannot be had by known implication and basing from a belief that is sensitive, so how does my theory manage to avoid the conjunction problem? First, my theory (old or new) does not allow that a person is sensitive to many conjunctions of this type, because of conditions that are responsible for the qualifier "schematically" above. The tracking conditions in my theory are quantifications over the class of all probability functions on the language of evaluation. The familiar conditional probabilities (1 and 2) must be true in a subclass of that class, a subclass constrained by conditions that tell us which propositions are to have their probability values fixed in the evaluation (Roush 2005, 76-93). They give an answer, for every sentence, whether the value for its corresponding proposition is to be fixed or variable. They are designed to let vary, for example for sensitivity, those matters that are more affected by –p than –p is affected by them because what we expect the subject to be able to pick up on is the difference that –p makes to the world.[16] It turns out, though I had not anticipated it, that these conditions restrict the possibility of fulfilling the sensitivity condition for a conjunction where one conjunct, A, is more likely than the other, B, to be false if the conjunction is false. In many of those cases the logical properties of conjunctions force the conjuncts to get fixed and the sensitivity condition for the conjunction is undefined.[17]

Fortunately the sieve is not too restrictive, in that it does allow the possibility of sensitivity to a conjunction in cases where $P(–A/–(A.B)) \approx P(–B/–(A.B))$. This is a good thing since it should be possible to be sensitive to generalizations with lots of instances of the same type and to propositions like "It is a black house." But in that case sensitivity to the conjunction requires tracking of each conjunct. One has to be sensitive to each conjunct at a somewhat higher level than the set threshold in order to give the conjunction-belief threshold sensitivity, however not at as high a level as would be needed if one tracked each conjunct and did a multiple-premise deductive inference to the conjunction. Also, unlike preservation of sensitivity over multiple-premise deduction, no fulfillment of the adherence condition is required for this sensitivity to the conjunction.

The fixing conditions would seem to be right to allow some exceptions where it is possible to track a conjunction without tracking a conjunct that is very unlikely to be false if the conjunction is false, by the same fallibilism that says we don't have to be disposed to respond properly to every logically possible scenario in which p is false in order to be counted as knowing p, but only to the probable –p scenarios. It would be arbitrary to rule out the one entirely while being committed to the other. But if there are exceptions to the fixing conditions, then how do I avoid the problem of the cheap knowledge of a conjunct one is not sensitive to that can be derived from a conjunction one is sensitive to? Obviously, it follows from the results I just explained that knowledge of a proposition one is grossly insensitive to cannot be had by known implication and basing from a belief that is sensitive, so how does my theory manage to avoid the conjunction problem?

Interestingly, it is because the basing condition that my type of clause requires for closure is not fulfilled for the subject's belief in the proposition she is insensitive to, *because* of that insensitivity. This can be proved from the definitions since in order for the subject's belief in, say, B, to be based on her belief in the conjunction A.B, it must be that P(–b(B)/ –b(A.B)) is above threshold z.[18] But she must be sensitive to B to fulfill this, and she is not. Indeed the less sensitive she is to B, the less she fulfills basing, and the intuitively strong counterexamples involve gross insensitivity. Thus, one cannot cheaply know a conjunct one is insensitive to via the new type of closure clause discussed above.

VI. Taking Stock

Any time one uses a sensitivity condition that is not infallibilist, there will be varying strengths of sensitivity. Since the level of sensitivity makes a difference to how epistemically sound the belief is, I keep track of the expected degrees via variables for thresholds. (It is an advantage of conditional probability over counterfactuals that these degrees can be expressed so directly.) Once this is done, then because probability is an axiom system the behavior of error over known implication can be calculated. The results of these calculations allow us to know the error consequences of any level of closure clause we might accept, and are surprisingly reassuring about how quickly (or rather, slowly) error grows. This in turn allows us to avoid all sorts of cheap knowledge problems, and even brain-in-a-vat skepticism. The closure problem is really an error problem.

Brueckner, Anthony (this volume), “Roush on Knowledge: Tracking Redux?” this volume.

Goldman, A. 2008. “What is Justified Belief?” in E. Sosa, J. Kim, J. Fantl, and M. McGrath, eds. *Epistemology: An Anthology* (Blackwell Publishing), 333-347.

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

Roush (2009), “Précis of *Tracking Truth,*" and “Replies to Critics,” *Philosophy and Phenomenological Research* 79(1), 213-222, 240-247.

Roush (2010a), “The Value of Knowledge and the Pursuit of Survival,” *Metaphilosophy* Vol. 41, No. 3, 255-278.

Roush (2010b), “Closure on Skepticism,” *Journal of Philosophy*, Volume CVII, No. 5: 243-256.

Roush (2012), “Skepticism about Reasoning,” *New Waves in Philosophy of Logic*, Greg Restall and Gillian Russell eds. Palgrave-Macmillan.

Sober, Elliott, “The Special Consequence Condition of Confirmation,” Handout

Vogel, Jonathan (1990), “Are There Counterexamples to the Closure Principle?” in *Doubting*, M.D. Roth and G. Ross eds., 13-27. The Netherlands: Kluwer.

Vogel, Jonathan (2000), “Reliabilism Leveled,” *Journal of Philosophy*, 97/11: 602-23.

[1] I think that safety as a necessary condition does achieve closure automatically, without a recursion clause, despite recent arguments to the contrary. When the subject's belief in the conclusion is *based* on her belief in the premise, the worlds in which the subject believes the conclusion of the implication are a natural subset of the worlds in which she believes the premise. If so then the conclusion belief is as true as the premise belief that implies it. This automatic closure leads to a cheap knowledge problem, though.

[2] In what follows I will largely ignore the adherence condition or, when it makes no difference, let it tag along by using "tracking" interchangeably with "sensitivity," but the adherence condition remains part of my theory. It does make an important difference, below, in preservation of sensitivity in multiple-premise closure clauses, and in Roush 2010a.

[3] I can no longer see why I required that none of the $q_1, \ldots, q_n$ be equivalent to q in my original formulation of this condition, but of course I might have been right there and wrong here.

[4] For this reason I expressed doubt about whether we should even have the belief clause in the definition of knowledge of implication (Roush 2005, 47).

[5] Even logicians have a hard time giving a non-circular proof of the simplest logical truths.

[6] Though I have used the same labels, c' and d', as I used in the book, these are much simplified, and therefore not fully correct, versions of those conditions.

[7] A similar change is made for knowledge of necessary truths known by known implication from knowledge of other necessary truths. The beliefs in the premises of such an argument would have to fulfill the "sensitivity" requirements for knowledge of non-implicational necessary truths, c' and d' above.

[8] It is such an obvious assumption that I actually left it out in my definition in the book. It was not clear to me then how crucial it is to the calculation of growth of potential error.

[9] To simplify the presentation I have made use of the fact that $P(–b(q_1). –bq_2/A) = P(–b(q_1). –b(q_2)/ –q_1. –q_2)P(–q_1. –q_2/A) = P(–b(q_1). –b(q_2)/ –q_1. –q_2)(1/3)$ and similarly for the other terms. Note that we are given in the assumptions only that we track one q_n at a time, so we have two sensitivity contributions in the term just cited.

[10] Precisely, the sensitivity for three-premise one-step deduction $azs^3 + bzst^2 + \ldots + gzs^2t$, where $a + b + \ldots + g = 1$. Ignoring the small contribution of differences between a, b, c, …, g in case s and t are different, it is $zs^3/7 + 3zst^2/7 + 3zs^2t/7$. For n premises, the precise expression is $a_1zs^n + a_2zs^{n-1}t + \ldots + a_nzst^{n-1}$.

[11] An example from Elliott Sober (handout): The hypothesis that an overturned card is the jack of hearts entails that the card will be a jack, but if you learn that the card is red that does not provide support to the claim that it is a jack.

[12] I take the safety condition to be explanatory for the same reason, that deduction is a safe process. But I think that safety achieves closure without a closure clause. This makes closure natural and adds further explanatory force. However, it comes at a price since the view thereby counts as knowledge cases with vast potential error, e.g., a belief that one is not a brain in a vat. The closure is achieved automatically because of a complete disregard for false positive error.

[13] Safety's automatic closure also yields a cheap knowledge problem since if my hand-belief is safe then it follows automatically that my belief I am not envatted is safe. (And this latter conforms with intuitions under the assumption of the safety of the premise.) We balk at the idea that that is enough for knowledge, though, illustrating that safety, as technically defined, is not enough for knowledge.

[14] I track the second conjuncts because the gauge is assumed actually to be reliable. The issue is whether I know that it is.

[15] This feature did not come up in my original presentation in the book because the closure clause was recursive.

[16] The conditions also determine whether methods are fixed; so, methods are not always fixed, but one common situation where the conditions imply they are is where whether that method was used is independent of the truth value of p. Nozick's case of a father of an accused criminal does not, which I say is a good thing (Roush 2005, 68-71, 112-113). The conditions also imply that the random fabricator's method is fixed at what it is on the occasions when she actually uses her eyes, as I did not notice in the book (Roush 2005, 127-8).

[17] The theory (Roush 2005, 76-93) says that if a proposition, A, satisfies the following two conditions, together called "*", then it is fixed in the evaluation of sensitivity:

$$|P(A/-p) - P(A/p)| \leq |P(-p/A) - P(-p/-A)| \text{ and}$$

$$|P(-A/-p) - P(-A/p)| \leq |P(-p/A) - P(-p/-A)|,$$

where p is the proposition that may or may not be known. Now, if p is a conjunction A.B, then the conditions for a conjunct, A, being fixed are:

$$|P(A/-(A.B)) - P(A/(A.B))| \leq |P(-(A.B)/A) - P(-(A.B)/-A)| \text{ and}$$

$$|P(-A/-(A.B)) - P(-A/(A.B))| \leq |P(-(A.B)/A) - P(-(A.B)/-A|.$$

Due to the logical relations between conjunctions and their conjuncts both conditions become:

$P(–A/–(A.B)) \leq P(B)$

If both A and B (switching places with A) get fixed in this way then the sensitivity condition, with –(A.B) in the condition, is undefined. There are cases where one or another of these is violated, for example when either A or B, or both, have very low probability. However, if either conjunct has very high probability, e.g., "I will not win the lottery tomorrow," then they are fulfilled, so the theory does not let one know that without tracking it.

[18] $P(–b(B)/ –b(A.B)) = P(–b(B)/ –B)P(–B/–b(A.B))$