Statistical Syllogistic, Part 1

Lawrence H. Powers

Title: Statistical Syllogistic, Part 1
Author: Lawrence H. Powers
Response to this paper by: Jonathan Adler

Statistical Syllogistic, Part 1

In this paper, I attempt to lay down a systematic basis for evaluating arguments like statistical syllogisms but having a more complex structure. I call this paper 'Part 1' for a simple reason. My efforts here will run afoul of various difficulties. Sometimes I won't have figured out the solution to a mathematical problem that arises. Sometimes I run into a philosophical problem but don't know how to solve it. So it makes me feel better to imagine that somewhere in my future is a Part 11, in which I will solve all these problems. Part 11 is at this point purely fictitious.

As an example of a statistical syllogism, let us consider the following. 90% of professors like Bach. Joan is a professor. Therefore Joan likes Bach. First of all, I have absolutely no idea how many professors actually like Bach. The statistics in my examples are just made up statistics.

Well, in this argument the first premise gives a proportion of a reference class, professors, which have a target property, liking Bach. This first premise states a frequency probability: the probability that a professor likes Bach is 90%. The second premise assigns an individual, Joan, to the reference class; Joan is a professor. And the conclusion asserts that that individual has the target property; Joan likes Bach.

The argument is, of course, <u>inductive</u> in the broad sense that the premises do not give certainty to the conclusion but only give a probability. Relative to the premises the conclusion is 90% probable. That is, there is a 90% probability, relative to the premises, that Joan likes Bach. Since Joan is not a class, this probability is an <u>epistemic</u> probability. So the effect of a statistical syllogism like the one about Joan is to transform a frequency probability for a reference class into an epistemic probability for an individual falling into that reference class.

I shall also say that the syllogism is 90% <u>valid</u>, meaning the conclusion is 90% probable relative to the premises.

The premise of our argument says that 90% of professors like Bach. This means that 10% of professors don't like Bach. Either they dislike his music, or are indifferent to it, or have never heard it, or perhaps even they heard it and liked it but didn't know who it was by, and so didn't appreciate Bach for having written it. In one way or another, they fail to like Bach. So our argument is, I shall say, 10% contravalid, valid to the opposite conclusion.

In this paper I will be interested in the validities of arguments. I will not be concerned with various "detachment" problems. For instance, what I call the first detachment problem asks: under what conditions will the 90% probability of the conclusion relative to the premises become a 90% probability for me personally? And a second detachment problem asks: if the conclusion has a certain probability for me personally, under what conditions do I believe it is <u>true</u>? And of course, other problems can be raised. For instance, under what conditions would an argument be useful for convincing someone other than myself.

I shall be interested only in questions of validity.

Earlier I said that a statistical syllogism is inductive in a broad sense. This is true, of course, but this broad sense of 'inductive' is somewhat misleading. Statistical syllogisms aren't really inductive at all. I shall say they aren't inductive in a real or Humean sense. They aren't really inductive; they are, so to speak, quasi deductive.

We begin to suggest this idea by looking at the conclusion of our statistical syllogism: Joan likes Bach. Well, this conclusion isn't a general law-like statement like all ravens are black, nor is it a general explanatory theory like atomic theory or relativity theory. It's just a statement that a particular individual, Joan, has a certain property, liking Bach.

Still this point is not decisive. "All examined ravens are black; therefore the next raven to be examined will be black," is a real inductive argument and it appears to state that a particular raven, the next one, will be black. One might reply that the conclusion here only appears to be particular; it really says that the next crow, whatever crow that may be, will be black. But, then, it may be that we know what raven is to be examined next. For instance, we may intend next to examine our friend Paul's pet raven Elmer. So we argue that Elmer will turn out to be black. And this conclusion is particular.

One might however try to push further about the conclusion of statistical syllogisms vs. real inductions. If we really know exactly what proportion of professors like Bach, someone must have surveyed all professors, including Joan, even though they didn't retain the information about the particular professors and only passed on the information about the proportion. But to know that all observed ravens so far have been black we don't need to have looked at Elmer, who isn't one of the so far observed ravens. The Elmer argument seems ampliative in a way that the Joan argument doesn't.

Still, I admit that at this point, we are straining at gnats. This discussion of the conclusion of a statistical syllogism is at best suggestive. My real reason for saying that statistical syllogisms aren't really inductive doesn't have to do with their conclusions; it has rather to do with Hume's problem of justifying induction.

Hume's problem isn't about statistical syllogisms. There isn't really any problem about justifying statistical syllogisms. Their strengths and weaknesses are transparent.

Consider an interesting idea about justifying induction. In a book and an article, Donald C. Williams[1] has suggested a justification of induction. He assumes that real inductions are sampling arguments. His justification is really therefore an attempted justification of sampling

---

[1] Donald C. Williams, The Ground of Induction Harvard University Press; Cambridge Massachusetts, 1947 and "The Probability of Induction," p. 395 ff., in Probability, Confirmation, and Simplicity, ed. M. Foster and M. Martin, Odyssey Press, N.Y., 1966

arguments. His justification actually reports to his fellow philosophers results which are well known in the mathematical foundations of sampling theory.

There is a theorem in mathematics that says that if you have a large population which is, say, 40% Republican and if you pick a sufficiently large sample size s, then you can prove that a very large percentage, say 95%, of all subsets of size s of that population have very close to 40% of their members Republican. More simply put: most possible samples match the population, match in having about the same proportion of Republicans.

So the justification of sampling arguments goes as follows. We have a population with an unknown proportion of Republicans. We argue: most possible samples match the population. Our sample is one of the possible samples. So our sample matches the population. But our sample is, say, 60% Republican. So the population is about 60% Republican. End of justification.

Now the interesting point about this justification is that Williams assumes that the real inductions are sampling arguments and that real inductions and sampling arguments need justification. He also in effect assumes that a statistical syllogism <u>doesn't</u> need justification, for he justifies the inference "60% of the sample are Republicans/therefore about 60% of the population are Republicans," by citing a statistical syllogism: Most samples match the population; our sample is a sample; therefore our sample matches the population. And so Williams is in effect assuming what I'm saying here: statistical syllogisms aren't really inductive.

Williams assumes that real inductions are sampling arguments. I don't think this is right. It seems true because "all observed ravens are black; therefore all ravens are black' might apparently be re-stated as "100% of ravens in our sample are black: therefore 100% of all ravens are black."

Now of course sampling arguments are a lot more like real inductions than statistical syllogisms are. The conclusion, "60% of the population are Republicans" is clearly general and ampliative in a way that 'Joan likes Bach' isn't. Still I don't think sampling arguments are full-fledged real inductions.

First, the conclusion of a sampling argument isn't "100% of ravens…" It's rather "about 100% of ravens…" whereas the inductive argument concludes "<u>all</u> ravens."

Second, "100%" isn't really the same thing as 'all'. For infinite classes, there is a clear difference. Arguably, 0% of natural numbers are square numbers, because square numbers get rarer and rarer as we go through the natural numbers (for instance one trillionth of numbers from 1 to one-trillion-squared are square numbers), and so 100% of natural numbers are non-square. But clearly it is false that all natural numbers are non-square.

And finally, and most importantly, I agree with John Pollack's view that in real inductive generalization, the direct conclusion is not that all <u>actual</u> ravens are black ((x) (rx⊃bx)) but is rather a law like generalization with counterfactual content that all actual ravens are black and moreover if there had been further ravens, they would have been black also. In sampling arguments the conclusion relates only to the actually existing population and has no counterfactual content.

Still a justification of sampling is interesting. Possibly, thinking of actual ravens as a sample of the possible ravens which might have been, such a justification could be extended in some way as a justification of real induction.

Now despite my statement that statistical syllogisms don't really need justification, a justification has been suggested by Wilfred Sellars.[2]

Sellars' justification is simple. Consider the argument '90% of professors like Bach; Joan is a professor; therefore Joan likes Bach.' Now let 'x' range over all individuals and then for each individual x there will be an argument "90% of professors like Bach; x is a professor; therefore x likes Bach. An argument of this sort will (assuming the premises of the Joan argument are true) have true premises just in case x is a professor. And in 90% of such cases, the conclusion will be true also, since 90% of professors like Bach. So the Joan argument is of a sort, such that 90% of true-premised arguments of that sort have a true conclusion. In other words, this sort of argument is 90% truth preserving. And that is Sellars' justification.

One may well object here that Sellars uses the term 'argument' in a rather unusual way. An individual x will normally have more than one name n and may be referred to by many definite descriptions d. Then for that individual there would be many premises "n is a professor" and "d is a professor" rather than only one "x is a professor." Well, Sellars has extentionalized the concept of 'argument.' Perhaps some other term would be better, such as 'Premise situations and envisaged conclusion situation' or whatever. I shall hereafter speak of 'arguments' in Sellars' way when doing Sellars-type reasoning.

Now Sellars' justification of statistical syllogisms is similar to a standard justification of deductive reasoning. Given a true-premised deductive argument, say, "it's sunny; if it's sunny, it's spring, therefore it's spring," the standard justification is that such an argument is deductively valid and this means that if the premises are true, the conclusion must be true. So the argument is truth preserving.

Such justifications are convincing, I think. And this is so despite the fact that there is at least an air of circularity both in the Sellars' justification for statistical syllogisms and in the standard justification for deduction.

In the Sellars case, one might complain that Sellars is justifying the 90% syllogism about Joan by giving another 90% syllogism of exactly the same form, a syllogism on the meta-level

[2] Wilfred Sellars, Essays in Philosophy and its History, D. Reidel Publishing, Dordrecht, Holland, "Induction as Vindication" p 401.

and about arguments. This second 90% syllogism is: 90% of true-premise arguments of the x sort have true conclusions; the Joan argument is a true-premised argument of the x sort (with x = Joan); therefore the Joan argument has a true conclusion.

Similarly, in the deductive case, the modus ponens about sunshine and spring is being backed up by a meta-level modus ponens about premises and conclusions: If the premises (of the sun-spring argument) are true, the conclusion must be true; the premises are true; therefore the conclusion must be true.

If the air of circularity in these two justifications doesn't cause much alarm, perhaps it is because we never really felt arguments of these sorts needed justification in the first place.

At any rate, I don't think Hume would have thought they needed justification. Hume famously argued that real inductive arguments do need justification.

Part of his argument was that, in his sort of inductive argument, you can not prove that if the premises are true the conclusion must be true. Here Hume assumes that if you could prove that, then the inductive argument would be justified. In other words, Hume assumed that deductive inference is justified, and that one reason real inductive arguments need justification is that they aren't deductive arguments.

But Hume argued that in the case of his sort of inductive arguments, there is not even a probable reasoning that would lead from the premises to the conclusion, for such reasoning would involve knowledge of cause and effect. (Inquiry IV part II) Here Hume seems to say that if we have a sort of argument where you can prove the conclusion is probable given the premises -- like a true-premised 90% statistical syllogism – there is no problem of justification. In other words, statistical syllogisms are justified, and one reason real induction needs justification is because real inductions aren't statistical syllogisms.

It might seem that we can give a Sellars-like or Williams-like justification of one sort of Humean induction, the next raven sort. This justification will actually be a fallacious pseudo-justification. Let me state it here and then explain why it is fallacious.

Suppose we examine ravens in order, one at a time, the first, the second, and so forth, and now we are arguing "We have examined 1000 ravens and all examined ravens have been black; therefore the next (1001st) raven will be black."

But previously we gave this sort of argument each time we examined a raven. We looked at the first raven, found it black and argued "We have examined 1 raven and all ravens examined so far are black; therefore the next (2nd) raven will be black." And then we found the second raven to be black and gave a second argument. "We have examined 2 ravens…, and the next (3rd) will be black." And so now we are giving the 1000th argument of this sort. But the first 999 arguments had true premises and turned out to have a true conclusion. Now our present argument will or won't also have a true conclusion.

If it has a true conclusion, then there is another argument of this sort with true premises and a true conclusion, a 1000 altogether, and we go on to consider the 1001st such argument.

If our 1000th argument has a false conclusion, then 999 of the 1000 arguments had true conclusions and the 1000th argument is the last true premised argument of this sort, for thereafter it will never again be true that all examined ravens are black. So we can prove, before looking at the 1001st raven that at least 999/1000 of true-premised arguments of this sort have a true conclusion! So it seems there is a probability of at least 999/1000 that the 1000th argument will have a true conclusion! An amazingly high probability, indeed!

But this justification is too good to be true. And indeed, let us suppose that exactly 999/1000 of the true-premised arguments of that sort have true conclusions. This assumption entails that the 1000th argument will have a false conclusion, not that there is a probability of 999/1000 that it will have a true one.

As an analogy, suppose there is a bag containing 100 balls and you pick out 90 of them and lay them aside. Suppose all 90 are red. You can then prove that at least 90% of the balls originally in the bag were red. But you can't prove that any of the remaining balls still in the bag are red. So there is no provable probability that the next ball will be red.

Now in the rest of this paper, I will not be concerned with justification problems. I will simply assume that statistical syllogisms are a legitimate form of reasoning and I will run with that assumption as far as it will go.

I shall have further interest in Sellars' method of validating arguments, but I shall be thinking of this validation method not as a justification but as a method of assigning validities. In a generalized form it will play a role in statistical syllogistic similar to the role that formal semantics plays in deductive logic.

In particular, I shall try to assign validities to two forms of argument, which at first sight may seem to have no particular validity. I shall be considering the chain argument and the convergence argument.

Where 'this' refers to a particular individual and A, B, and C are classes, the chain argument is: 90% of A is B; 90% of B is C; this is A; therefore this is C. Intuitively this argument is equivalent to a chain of two 90% valid arguments: 90% of A is B; this is A; therefore this is B; and 90% of B is C: this is B; therefore this is C. Intuitively we think that the validity of the chain of two arguments and therefore of the chain argument should be 90% of 90% or 81%. I shall try to use Sellars' method to attach some such validity to the chain argument.

In the convergence argument, we will want to consider different percentages. So let X% be some specific percentage and Y% also. Then the convergence argument is: X% of A are C; Y% of B are C; this is both A and B; therefore this is C. Intuitively, the convergence argument is the convergence of two simpler arguments to the same conclusion -- this is C. One is the X% valid argument: X% of A are C; this is A; therefore is C. The other is the Y% valid argument: Y% of B is C; this is B; therefore this is C. In the case of the convergence argument, we do not have any obvious intuition as to how valid the argument should be. Usually, when discussing such an argument, the main comment is that it would be nice if we had another premise, a premise saying that Z% of things that are both A and B are C. Of course such an additional premise would, so to speak, "swamp out" the X% and Y% premises and the validity of the argument would be Z%. I shall try to assess the validity of the argument <u>without</u> the extra premise.

Let us begin with the chain argument. 90% A are B and 90% B are C. If we could deduce that 81% of A are C, we could then easily give the argument a Sellars' validity of 81%, for we would consider the various cases Ax and 81% of these would be ones where Cx.

Unfortunately, we cannot deduce that 81% of A is C. In fact, we cannot determine anything about what proportion of A's are C's. So we can't validate the argument in this way.

It is consistent with the two 90% premises that 0% of A is C or that 100% of A is C or anything in between. Imagine a room. In the room there are 9 ducks. And there is another duck out in the hallway. And also in the room is one goose. So 90% of the ducks are in the room and 90% of things in the room are ducks. But we do not conclude that 81% of the ducks are ducks. Obviously 100% of the ducks are ducks.

Again, same room. Again, 9 ducks in the room and one duck in the hallway. But now 81 geese are in the room with the 9 ducks. So 90% of the ducks are in the room and 90% of things in the room are geese. But we don't say 81% of the ducks are geese. Obviously 0% of the ducks are geese.

So since we cannot validate the chain argument in the simple way that Sellars validated the Joan argument, we shall need to look for some more subtle procedure.

To get a clue as to what this should be, we turn our attention momentarily to the convergence argument. If in the convergence argument we add the extra premise that Z% of things both A and B are C, the validity will be Z%. In that case there are three specifications concerning the individual <u>this</u>. This is A, this is B, and this is A and B. Of these the third is most specific. So its frequency probability Z% becomes the epistemic probability that this individual is C.

Suppose we don't have the Z% premise. Then there are two equally specific specifications this is A and this is B for which we have frequency probabilities. No obvious validity results. Suppose we also don't have the Y% premise. Then all we have is that X% of A is C and that this is A and it is B. The information that this is B however, does no work, since we have no frequency probabilities for B's or for AB's. So the validity is X%, from the frequency probability for A's, the most specific specification about <u>this</u> for which we have a frequency probability.

Now we return to the chain argument.

We tried to validate the particular argument about the individual <u>this</u> by referring it to a class of arguments about all the individuals x that were in the given class A. But the chosen reference class of x arguments did not have a known frequency probability of having true conclusions. So our attempt failed.

So we try a less specific reference class for which, hopefully, we will find a known frequency probability. We shall abstract not only on the individual but also on the given class A. Now we let the letter 'A' be a variable ranging over all classes A which satisfy 90% A are B, for the given B and we consider all pairs <x, A> where 90% A are B and x is in A.

It is convenient at this point, instead of working directly on the chain argument itself, to first look at three similar but simpler arguments, where the problems will arise in a simpler form. These three arguments are arguments 2, 3, and 4 below.

The Chain Argument	Argument 2	Argument 3	Argument 4	Argument 4`
90% A are B	All A are B	90% A are B	90% A are AB	90% A are B
90% B are C	90% B are C	All B are C	90% AB are ABC	90% AB are C
This is A	this is A	this is A	this is A	this is A
This is C	this is C	this is C	this is ABC	this is BC

We begin with argument 4. I call this the inclusion chain argument, since class A includes class AB which includes class ABC. This argument is unproblematically 81% valid. A 90% part of class A is AB and a 90% part of AB is ABC. So ABC is 81% of A. So the Sellars validation is immediate.

We can simplify argument 4 (into argument 4') so as to make it look more like the original chain argument and so as to bring out more clearly the real differences. In the first premise, it is the same to say an A is an AB or to say an A is a B, so the second 'A' in the first premise may be omitted. Similarly, in the second premise a second 'AB' may be omitted. Finally, since the third premise says 'this is A,' the A in the conclusion may be omitted.

So 4' is just 4 re-written. It is 81% valid and 81% of A's are BC's.

The problems about the original chain argument correspond to the BC in the conclusion of 4' and the AB in the second premise of 4'. The BC problem is about A's that aren't B's. It's a non-B problem. 81% of A's are B that are C's and 9% are B's that aren't C's. But no, some, or all of A's that aren't B may be C. So anywhere from 81% to 91% of A's may be C's.

The AB problem is the transitivity problem. In the first premise some A's are B's. In the second some of these same B's are C's. But if the B's that are C's weren't the B's that are A's, then the term B wouldn't really be connecting the A's and the C's and so it might be that no A's are C's. The problem is like that of the undistributed middle term in categorical syllogistic.

These two problems occur separately in arguments 2 and 3.

In argument 3 the BC problem appears in pure form. Intuitively the argument should be 90% valid since it breaks down to a chain of a 90% valid argument and a 100% valid argument. However the premises say that 90% of A is B and all B is C. So it follows that 90% of A is BC and at least 90% of A is C. But the A's that aren't B's may either none, or some, or all be C's. So the percentage of A's that are C could be anywhere from 90% to 100%. So what validity should we attach to the argument?

I shall propose two different answers to this question. One, which will assign this argument a 90% validity, I call the nice solution. Unfortunately, the nice solution doesn't really fit with my general methods in this paper. Therefore I shall regretfully be calling the other not-so-nice solution the official solution.

The nice solution borrows a very appealing idea suggested by Glen Shafer in his book, Mathematical Theory of Evidence,(Princeton, 1976). Schafer suggests that epistemic probabilities do not obey all the laws of standard probability theory, which is the theory of frequency probabilities.

In particular if we have no evidence whatsoever as to whether P or ~ P, standard probability theory says the probability of P and the probability of ~ P must add up to 1, or 100%. So since we have no reason to prefer either, we assign the ignorance probability of 50% to each.

But Shafer thinks this is silly. If we have no evidence whatsoever in favor of P, we should assign it an epistemic probability of 0%, meaning 'no support' and if we also have no evidence whatsoever for ~P, we should assign it also a probability of 0%, leaving the remaining 100% probability free-floating and unassigned. Then as evidence comes in for P, say witnesses testify that it's true, the probability for P will edge up, and as evidence comes in for ~P, say witnesses testify for it, then ~P's probability will edge up. The probability for P plus that for ~ P plus the free-floating probability always add up to 100%, not just the first two together.

I find this idea very appealing, and it suggests the nice solution to the problem of argument 3.

Since we know that at least 90% of A are C and don't know more, we assign a 90% epistemic probability to the conclusion and the argument is 90% valid. Since we don't know any A's are non-C, we assign 0% contra validity to the conclusion that this individual is non-C.

Another remaining 10% validity is left free-floating. I like this solution, but unfortunately it doesn't really fit with my general methods in this paper. The official solution, or not-so-nice solution, is just what you would now expect: the 10% free floating probability of the nice solution is split into two 5% parts and these are assigned as extra validities for this being C and this being not C. So the validity of the argument is 95% and the contra validity is 5%.

Of course the extra 5%'s do not really seem to flow from the premises but come from background considerations. As such they remind us of irrelevant validities such as those in deductive logic, like the validity by which a necessary truth follows from premises totally unrelated to it.

Further consideration of the issue of these two different solutions is deferred to Part II.

I said above that the official solution is more in accord with my "general methods" in this paper. Since these have not yet emerged in the interesting cases, the reader may wish to re-read later or defer until later the following detailed explanation.

My general method considers all quadruples <x, A, B, C,> where x is an individual; A, B, and C are sets and $x \varepsilon A$, 90% A are B, all B are C, and it asks: in what proportion of these quadruples is the x in the C? This proportion is the Sellars validity of argument 3.

Consider all such quadruples where A and B are fixed so that 90% A are B and where BC is fixed as being B and where C-hood and non C-hood are assigned in any fixed way to individuals outside of AUB. So everything is fixed except the C-hood and non C-hood of individuals, which are AB*. (Here B* is non-B. In general if a single letter stands for a set the letter followed by an asterisk stands for the complement.)

In half the ways of assigning C-hood in that area, a given individual in that area will be C, in the other half C*. So the probability of an AB* being C is 50% in that way of fixing B, A, BC, and A*B*. Since this is true for every way of fixing those items, the 50% probability prevails throughout the quadruples. So the probability of an x, which is A, but not B being C is 50%. So the probability of an A being not B and C is 5%. And that's the official solution.

Now let us turn to argument 2 where the transitivity or AB problem arises in pure form.

Intuitively the argument should be 90% valid since it breaks down to a 100% valid inference followed by a 90% valid inference.

But, similarly to the situation in the original chain argument, we cannot deduce that any particular proportion of A's are C's. Here come those ducks and geese again! We can't deduce that 90% of A's are C's.

There are 9 ducks in the room but now there is no duck out in the hallway. So all the ducks are in the room. There is additionally either 1 goose or 81 geese, either 90% of things in the room are ducks or 90% are geese. But 100% of ducks are ducks, not 90%, and 0% of the ducks are geese, again not 90%.

My solution to this problem, as I suggested earlier, is to abstract on both this individual and the class A. Consider all classes A which are subsets of B and for each such A all pairs <x, A> where x is an A. What proportion of all such pairs for all the A's has x in C?

Well first, if x is a member of A for some A which is a subset of B, then x is a member of B. And if x is a member of B, then x is a member of some subset A of B, and 90% of such x's, the members of B, are C.

Further if x1 and x2 are both members of B, they are members of the same number of subsets of B as each other. Therefore 90% of pairs <x, A) have x in C. Therefore the Sellars validity of argument 2 is 90%, when we abstract on both this and A.

And this is my solution to the transitivity problem.

First, the math about subsets; if B has N members it has 2^N subsets. If we subtract {x}, then B- {x} has 2^{n-1} subsets A'. Then the 2^{n-1} sets A = A' + {x} are the subsets containing x.

The reader might wonder why in the above, I was not satisfied to simply show 90% of the x's for which there was an A were C's and instead insisted on proving that 90% of the pairs <x, A> had x in C. The short answer is that I am trying to show that 90% of the "arguments," one for each pair <x, A> in which the premises are true have a true conclusion x is C. If one x that wasn't C had millions of A's and the x's that were C had only one A each, then we might have most of the arguments with false conclusions even if most of the x's were C. That's the short answer.

Here's the long answer. I present a paradox which arises when you think of things the wrong way and then show how thinking of things the right way solves the paradox.

Suppose that yesterday I heard the following sound argument: Joan is a music lover; if Joan is a music lover than Joan likes Bach; therefore Joan likes Bach. So I knew yesterday that Joan likes Bach.

Suppose today I remember that I heard that argument yesterday, and I can't remember what the antecedent (Joan is a music lover) was. I just remember that there was such a true antecedent. Surely if I remember that there is a sound argument that Joan likes Bach, even if I don't remember what that argument was, I still know that Joan likes Bach.

And in fact I can today know that Joan likes Bach by the following valid 'indirect' argument.

$\exists$p((p$\supset$Joan likes Bach) & p)

Joan likes Bach

And this 'indirect' argument is just as valid and sound as the original argument.

Now suppose yesterday I heard the true premised argument.

90% of Professors like Bach

Joan is a Professor

Joan likes Bach,

And on this basis, having no other relevant information about Joan, I assigned 90% probability to Joan liking Bach.

Today I remember hearing this argument, but I have forgotten the reference class (professors) and just remember that there was a reference class X. Do I have grounds today for thinking that there is a 90% probability that Joan likes Bach? Surely knowing that there is evidence for that probability is enough, even if I don't remember exactly what that evidence is.

So I try to indirectly justify that probability by considering.

($\exists$ X) (90% X likes Bach, Joan is an X)

Joan likes Bach.

And I hope that this indirect argument is 90% valid. But it has no significant validity at all. The premise is true if and only if there are at least 9 Bach likers in the universe and at least one non-Bach liker. For in that case the 9 and the 1 are a class X' such that 90%X' like Bach. And we can then form a class X of which Joan is a member by replacing Joan for one of the members of X', replacing the 1 by Joan if she doesn't like Bach or one of the 9 if she does. Conversely, if the above premise is true, then X has 9_n members who like Bach and n who don't, for some n. At a minimum therefore there are 9 Bach likers and 1 non-Bach liker.

But surely the argument

There are 9 Bach likers in the universe at least, and at least 1 non Bach liker

Joan likes Bach

is not significantly valid.

In other words, it is not true that any significant proportion of x's (like Joan) for which there is some X or other are therefore Bach likers.

What is true, as Sellars' original validation brought out, that for each X, 90% of the conclusions of the arguments involving X have true conclusions.

Better put, the following indirect

argument is 90% valid:

90% of <x, X> such that 90% X like Bach and x is in X are such that the x likes Bach.

But <Joan, the X of yesterday> is such an <x, X>

Joan likes Bach.

Now, finally, we return to the original chain argument. The solution is now just what you would expect. Abstracting on this and A, we find 81% nice validity that this is BC, 9% nice validity to this is B but not C, 5% not nice validity to this isn't B but its C anyway, and 5% not nice validity to this isn't B but it isn't C either. So putting the validities together, the argument is 86% valid and 14% contravalid.

Here we need to prove the 81% validity to this is BC.

The following argument is valid because its premises are inclusion chain premises.

90% of <x, A> have $x \in B$.

90% of <x,A> with $x \in B$ have $x \in B$ and $x \in C$

81% of <x,A> have $x \in B$ and $x \in C$

And each premise is true

Premise 1

Because for each A, we have 90% of its x's in B, so 90% of the pairs for this A have x in B. This is true for every A, so 90% of all pairs have x in B.

Premise 2

Because each $x \in B$ has the same number of A's, namely

$$\binom{B^*}{n}\binom{B-1}{9n-1} \text{ for each suitable n,}$$

or

$$\sum_{9n<B} \binom{B^*}{n}\binom{B-1}{9n-1}$$

n≤ B*

My discussion of the chain argument is concluded. However before going on to the convergence argument one point has arisen in that discussion which is no doubt bothering many readers and which I need to consider more fully.

It will have been noticed that I assigned a validity of 86% to the chain argument form. But the duck and geese examples bring out that some instances of that argument form are not 86% valid. For if in the general form, we take C to be A, the inference <u>this is A therefore this is A</u> is 100% valid and if we take C to be A*, the inference <u>this is A therefore this is A*</u> is 0% valid.

This at first seems shocking, for we are used to the idea in deductive logic that a valid argument form has all its instances valid. But it shouldn't shock us, for even in deductive logic we are used to the idea that an <u>invalid</u> argument form like affirming the consequent (p⊃q, q:.p) may have valid instances (p⊃p, p:.p). The reason is that .instantiation may add new information, for instance that q=p, to the premises and this new information may change an invalid into a valid argument.

But it is a familiar fact that adding new information to an inductively valid argument may undermine its validity or may enhance its validity. Therefore instances of valid arguments with a certain validity should not be expected to have that same validity.

The grue bleen argument is an instance of this, but I confess I myself learned the general fact here from a very interesting example in a paper by George Bowles that I heard at an informal logic conference.

At any rate I now give a series of simple examples illustrating this point:

- All observed ravens are <u>observed</u>. Therefore all ravens are observed.
- 100% of voters in my sample are in my sample; therefore 100% of all voters are in my sample.
- 90% of people in this room are smarter than I am but <u>I</u> am in this room. Therefore <u>I</u> am smarter than I am.
- 90% of voters in my sample are (it turns out) members of a small fringe political party. Therefore 90% of all voters are members of a small fringe political party.
- All the ravens I've observed turned out to have a disease very rare among ravens. Therefore all ravens will turn out to have a disease very rare among ravens.

I now turn to the convergence argument.

In the convergence argument, if either X% or Y% is 100% the argument is deductively valid. If either is 0%, deductively contravalid. If both 0 and 100 occur, the premises are contradictory.

Here I am supposing, to avoid conceptual difficulties involved in the talk about 'proportions' in infinite classes, that A, B, C, and the universe of discourse of the argument are all finite[3]

But this statement is true of the usual <u>sequence</u> of natural numbers and does not state a well-defined property of the <u>set</u> of natural numbers, for if this same set is presented in a different sequence, a different proportion, or none at all, may result.

For instance in the sequence: 1st square, 1st non-square, 2nd square, 2nd non square, … the proportion of squares is 50%, though the set presented is exactly the natural numbers.

Other sequences can give any other proportion, or an oscillation and no resulting proportion.

In assigning a Sellars validity to this argument, we abstract on everything. We take all quadruples from a given universe of discourse which are of the form <x, A, B, C> and where x is a member of A and B and where X% A are C and Y% B are C. Looking at small universes we find that the proportion of such quadruples which have $x \in C$ varies depending on the size of the universe. So as to allow all possible sizes of A, B, and C to arise, we let the size of the universe go to infinity and take the limiting proportion as the validity.

But what is this limiting proportion? Unfortunately here the math is a bit heavy for me and I haven't yet proven a result. However I believe and here conjecture that it is given for X% and Y% by a formula I call F:

$$F = \frac{XY}{XY + X^*Y^*}, \qquad \text{where } X^* = 100 - X \text{ and } Y^* = 100 - Y$$

Though I don't have a proof that F is right, I do have some plausibility arguments.

First let us look at the behavior of F. If X% is 50%, it does no work, and the validity is given by Y%. If X and Y are both above 50, they cooperate together to give validity higher than either alone. If they are both lower than 50, the validity is lower than either alone. And if they are on opposite sides of 50, the validity is somewhere between them.

The table shows some selected validities, rounded off in some cases. This behavior is at any rate plausible.

X	100	---	0	----	50	X	60	40	90	10	90	90	X
Y	__	100	__	0	Y	50	60	41	60	40	40	30	X*
F	100	100	0	0	Y	X	69	31	93	7	87	76	50

A more compelling plausibility argument is by analogy. Our convergence argument is intuitively the convergence of two arguments, one with X% validity and the other with Y% validity. Let us consider the convergence of two such arguments in another field: argument from authority.

Abe and Ben are two authorities who always know the correct answer. We ask

---

[3] For an instance of the difficulties about infinite sets, return to the squares among the natural numbers. I said earlier that 0% of natural numbers are squares.

them yes-no questions. Abe tells the truth X% or 90% of the time and randomly lies the other 10% or X*% of the time. Ben tells the truth 70% or Y% of the time and lies 30% or Y*% of the time.

The facts are summarized in the four-part box.

		Ben	
		Truth Y%	lie Y*%
Abe	Truth X%	XY	XY*
	lie X*%	X*Y	X*Y*

So both tell the truth XY%% of the time. Both lie X*Y*%% of the time and then they disagree in the other cases.

So suppose they both agree on an answer. What is the probability that that answer is true? That is the probability that they are both telling the truth given that they agree.

And that is $\dfrac{XY}{XY+X^*Y^*}$, or F

If we try to use this same reasoning to validate the convergence argument itself, we do not succeed. Instead we find ourselves validating an argument similar to the convergence argument. I call it the 'similar argument.' It is

X% A is C

Y% B is C

x is A and y is B

$\underline{(Cx \equiv Cy)}$

x is C (equivalently, y is C).

This argument also has validity F. So our plausibility arguments are as follows. The Abe and Ben argument is analogous to the convergence argument and it has validity F. The Abe and Ben argument is actually an instance of the convergence argument and it has validity F.

The similar argument is similar to the convergence argument and it, the former, has validity F. The convergence argument is a special case (where x = y) of the similar argument and the latter has validity F.

But of course, plausibility arguments don't prove anything in mathematics!

At any event, the validity of the convergence argument is given by F -- or by some other formula. A problem begins to develop as we look at more concrete examples.

Suppose Joan is a member of the XYZ club, a local club in her community. Perhaps this is a club for people who both like to ski and like to sew, a sort of sewing skiers club.

And suppose we are given the premises of the following:

90% of professors like Bach

40% of XYZers like Bach

<u>Joan is a professor and an XYZer</u>

Joan likes Bach.

Now F says the validity of this argument is somewhere between 90% and 40%, each pulling in the opposite direction. (Actually it says the validity is 83%.)

But as I look at this argument, I find myself thinking the validity ought to be <u>higher</u> than 90%, because I find myself thinking that the XYZ club is a veritable hot-bed of Bach likers!

Now perhaps it's because I'm not really a music lover and have only the vaguest idea who Bach is. I heard somewhere that music is for savage beasts and haven't paid much attention to it since.

But isn't liking Bach a rather sophisticated achievement? Surely, I find myself thinking, only about 5% or 10% of people really like Bach!

But if <u>40%</u> of XYZers like Bach and XYZers are much more likely to like Bach then the ordinary person, then perhaps Joan's being in that club makes her <u>more</u> likely than the average <u>professor</u> to like Bach.

Now, as I make this pint, I realize that I have heard this very point before. Many years ago, a fellow read a paper at Wayne and reported that he and a colleague were working on a more dynamic statistical syllogism. The main point of his presentation was that in a mixed syllogism like our convergence argument, the significance of the 90% and the 40% should be assessed, not in terms of their positions relative to 50%, but in terms of their position relative to the background probability of something's having the target property C, for instance, being a Bach liker.

So I will call this point the dynamic syllogism point.

Well, then, suppose we take U to be the universe of discourse of our convergence argument and add an extra premise saying that the set C is, say, C% of U. (I use C slightly ambiguously here.) Then we will want to re-assess the validity of the argument taking C% into account.

The new validity will be

$$F(C) = \frac{C^* X Y}{C^*XY + CX^*Y^*} .$$

Though verifying F is giving me trouble, the correctness of F (C) is easy to prove.

F (C)'s behavior is just what the dynamic syllogism point asks for. It behaves just like F, except that things are now centered around C% instead of around 50%.

And here is another plausibility argument for F, for F = F (50).

Math

First let me describe the situation with respect to trying to prove F as the general validity. To prove that F is the general validity, it suffices to prove that F(50) is the general validity. I have proven, and it is a well known result anyway, that as the size of the universe of discourse U goes to infinity, a greater and greater proportion of its subsets are closer and closer to being 50% of U, and indeed in the limit practically all subsets are practically exactly 50% of U. (This result is actually equivalent to a result well known even to non-mathematicians: if you throw a fair coin a very large number of times, it will be almost certain that the number of heads will be almost exactly 50% of the number of tosses.)

So therefore almost all possible C's are 50% of U. So, since F (50) gives the proportion of quadruples with true conclusions for each such C, it is very plausible that F (50) prevails as the general validity.

Unfortunately, there is a countertendency. If say, X% = Y% =70%, a C at 70% of U will have more A's and B's compatible with it and therefore more quadruples than a C at 50%. In other words there are more C's at 50, but more quadruples per C at 70. So maybe F(70) rather than F(50) will prevail. I don't know.

Anyway, there is no such problem verifying the formula F (C).

Here is the proof.

Suppose as a percentage of U, the size of A, B, and C are known and suppose we are given X% A are C and Y% B are C. Suppose that the various sizes are compatible, so, say, X% of A fits into C and X*% of A fits into C*, and similarly for B. The size of A is A, of B is B, of C is C. (Slightly ambiguous.)

Then the probability of a point in C being AC is $\dfrac{X\%A}{C}$

And the probability of a point being B is $\dfrac{Y\%B}{C}.$

These are assigned independently so the probability of a point in C being A and B is $\dfrac{X\%Y\%AB}{CC}.$

Similarly the probability of a point that is in C* being both A and B is

$\underline{X^*\%Y^*\%AB}$

C*C*

So the probability of a point that is both A and B being C is

$$\frac{\frac{C}{C+C^*}\,\frac{XYAB}{CC}}{\frac{c}{c+c^*}\,\frac{XYAB}{CC} + \frac{c^*}{c+c^*}\,\frac{X^*Y^*AB}{C^*C^*}} =$$

$$\frac{\frac{XY}{C}}{\frac{XY}{C} + \frac{X^*Y^*}{C^*}} = \frac{C^*XY}{C^*XY+CX^*Y^*} = F(C).$$

As we have seen, if the background probability of something's being C is known, the convergence argument will be evaluated with validly F (C) rather than with validity F.

One problem, which I firmly defer to Part 11, is: the concept of background probability depends that of the universe of discourse of the argument. Is there some regular way of determining what that universe is? In deductive logic, it is the range of the variables. But in our argument, there is no individual variable until <u>this</u> or 'Joan' is replaced by 'x.' Presumably U should include A and B and C (or perhaps they would be considered relative to U as AU, BU, CU), but even if U should contain them, U would normally not be just the union of A and B and C but some "natural" super class of those three.

At any event, if a background probability for C can influence the validity of the convergence argument, what about background probabilities for A and B?

The derivation of F(C) shows that, if C's background probability is known, and if those of A and B are compatible with it, the validity is F (C); A and B do not change the validity.

If C's is not known, the main effect of A's and B's is to restrict the possibilities for C.

Let's look at an example.

Suppose we don't know Joan is a professor but find out (surprise) that she's a woman and are given:

70% of women like Bach

40% of XYZers like Bach

<u>Joan is a woman and an XYZer</u>

Joan likes Bach.

Now here it seems natural to take adult human beings as the universe of discourse. Professors, women, XYZers and Bach likers seem, mostly anyway, to be adult human beings.

But then women seem to be about 50% of U. While the XYZ club, being a small local club in Joan's town, would be a very small percentage of all adults.

If we make these assumptions, my earlier antimusical guess that only about 5% or 10% of adult people like Bach must be wrong.

For if 70% of W like Bach and 50% of U are W, then at least 35% of U like Bach, even if no men do. On the other hand, since 30% of W fails to like Bach, at least 15% of U must so fail. So, since the percentage Y% for the very small XYZ club puts no restrictions of C's size, we conclude that C is between 35% and 85% of U.

I conjecture that the Sellars validity would then be determined by looking at the range of possible C percentages. If 50% is in this range, as here, the validity is F (50) or F. If, say, the range were from 70% to 90%, the validity would be F (70), the validity for the permitted percentage closest to 50%. This conjecture will probably stand or fall with the main F conjecture.

Working out F in our particular case gives our argument a Sellars validity of 14/24 or (approx.) 61%. However, here a problem arises, for a different sort of analysis suggests a different validity.

Suppose we give a sampling analysis. 50% of U's are W's. Whereas XYZ is an insignificantly small part of U. So W's are a very hefty sample of U's. So, probably, the reason 70% of women like Bach is that 70% of people like Bach, well, adult people. But then we can replace the term 'woman' by the term 'adult person' in our argument. 70% of adult people like Bach, and 40% of XYZers, and Joan is a person and an XYZer.

But then being an XYZer is a more specific specification of Joan then being an adult person is, for we are assuming the latter is U. Therefore the validity of the argument should be 40%.

Now here the conflict between the Sellars validation method as I have developed it and the sampling analysis is not quite clear, since I am just conjecturing the former.

However a simpler argument will bring this out more clearly. For in the sampling analysis above, I in effect gave an analysis of the following:

the 'sampling' argument

50% of U are W

70% of W are C

this is U

this is C

By the sampling analysis, W is a very good sample of U, so 70% of U should be C. So the validity is 70%.[4]

But this 'sampling' argument (those are scare quotes) is our old friend, the chain argument. The only addition, which makes no difference to our previous analysis, is that the first term U is the universe of discourse. So, as a chain argument, this argument has a 'nice' validity of 35%, a bad of another 25%, and a total of 60% - not 70%.

The difference between the two approaches is simple. Both agree that 70% of U's that are W are C. But the Sellars approach thinks of the U's that aren't W as having an unknown proportion of C's, so it assigns the ignorance probability of 50% to a member of UW*'s being C. While the sampling analysis thinks the U's that aren't W will be 70% C like the U's that are W. Ergo, a different evaluation.

This conflict between the Sellars validity and the sampling analysis is a problem.

Although I was not expecting this particular problem, I was expecting, ever since starting this project that a problem like this would sooner or later show up. The reason is that the general method involved both in the standard justification of sampling and in my use of Sellars validation is a method in which a particular argument is referred to a class of arguments for which there is a theorem stating that a certain proportion of arguments in that class have a certain property. But what if a particular argument can be referred to two different classes each of which has a different theorem applying a different probability to the same or a logically related target property? Well, then we may have on the meta-level a convergence of two different statistical syllogisms assigning different validities to the same particular argument. And this is what is happening here.

Unfortunately I have no idea what to do about such problems. See Part 11!

Now this problem about the sampling analysis arose and distracted us while I was discussing the effect that background probabilities for A, B, and C had on the validity of the convergence argument.

I concluded that such probabilities for A and B have little effect, but that for C switches us from evaluation by F to evaluation by F (C.)

Now this leads to a strange result. Although F – or some other formula – gives the validity of the convergence argument, it is hard to imagine any instance of the convergence argument being evaluated by F - or whatever the right formula is. Well, let us say the correct formula is F'. The strange point is F' gives the validity of the convergence argument, but no actual convergence argument is evaluated by F'.

For whenever we consider any actual argument the letters A, B, C are replaced by specific class concepts. And looking at these concepts, we immediately estimate the sizes of the classes involved.

---

[4] One might object that in an ordinary sampling argument, the conclusion is that the population is about 70%C, for instance + 3%. However as U and therefore W approach infinity, the expected error approaches + O, or just plain O, and the probability of the sample matching the population approaches 100%.

Even if we have no empirical information, our very conceptual scheme will suggest a priori estimates. According to Plato and Aristotle, we start with being and divide it into living and non-living and then divide living into animal and plant and then divide and divide until we get to adult human being and then we divide two ways: into man and woman. So we think that women are 50% of adult people. If we then divide man into husbands and bachelors, we think husbands are about 25% of adult people. Since there are many occupations, we think professors or plumbers are a very small percentage of adult people, and so forth.

So in effect, we will virtually always use F (C) rather than F' to evaluate a particular convergence argument.

I might note by the way the non-nice part of the validity of the chain argument is unstable in this same way. A background probability for C would change it.

Well, further discussion of this matter is deferred to Part 11.

End of Part 1