Again With the Grue! (talk)

John T. Roberts

Draft of February 19, 2014

As everyone knows, there are two big problems about induction: the old one and the new one. The old one is Hume's problem: What good reason do we have to expect logically and conceptually contingent patterns in what we've observed so far to continue into what we're going to observe in the future? Evidently, we could have no a priori reason to expect this, because otherwise the patterns in question wouldn't be logically and conceptually contingent. And if we have an a posteriori reason to expect it, it seems that this reason would have to be something like, "Because expecting such patterns to continue has worked out well for us in the past." But to treat that as a reason to expect such patterns to continue into the future would be just a special case of expecting a logically and conceptually contingent pattern in what we've observed so far to continue into the future. So it would be at best a viciously circular reason. Hence, we could have neither an a priori reason nor an a posteriori reason to expect patterns in our observations to continue into the future. But those are the only kinds of reason there are. So we have no good reason at all to expect such a thing. Therefore, we have no good reason to believe a huge portion of what we ordinarily take ourselves to know. The old problem of induction is the problem of figuring out how inductive knowledge could be possible in spite of this apparently compelling argument – or, failing that, to learn how to live with it.

The new problem of induction is Goodman's problem. There are always indefinitely many different patterns in what we have observed so far, and it's logically impossible for all of them to continue into the future. So even if we are justified in projecting past patterns into the future, we need some way of picking and choosing among all the patterns we might project. But how can we do this, without just making an arbitrary and therefore unjustified assumption about which sorts of patterns are most likely to persist? There are different things you might call by the label "Goodman's problem." But this is what I'm going to mean by it: It's the problem of finding an epistemically justified way of choosing which patterns in our observations-to-date to expect to continue into our future observations.

Of course, Goodman's problem is often called "the Grue Problem" or "the Grue Paradox" because of a particular example Goodman used to illustrate it. (ANd this is the 1st thing on the handout.)

Suppose you are convinced that ordinary inductive reasoning is perfectly in order, and you understand this to mean that when you have observed a great many Fs, and all of them have been Gs, then it is reasonable for you to conclude (provisionally and cautiously) that all Fs are Gs. And so in particular, if you have observed a great many emeralds, and all of them have been green, then it is reasonable for you to conclude (provisionally and cautiously) that all emeralds are green. (This is an unfortunate example, since in fact, emeralds are green by definition; 'emerald' means the same as 'green beryl.' But for the sake of upholding a venerable philosophical tradition, let's ignore that unpleasant fact.)

The trouble starts when Goodman comes along and defines a new predicate, "grue," as follows: x is grue iff x is observed before time T and is green, or else is not observed before time T and is blue. Time T is some time that is still in the future, but not too far off in the future: We expect that we will still be around once time T arrives, and also that there are very many emeralds that no one will observe until after T. Now, all the emeralds you have observed so far are green, and of course, they have all been observed prior to time T, so they are all grue. So ordinary inductive reasoning, as you understand it, should lead you to conclude (provisionally and cautiously) that all emeralds are grue. This conclusion clashes with the one you drew earlier, because the earlier one predicts that once time T comes and goes, any new emeralds to be observed will be found to be green, whereas this new one predicts that any such emeralds will be found to be grue, and therefore blue. So, ordinary inductive reasoning, as you understand it, doesn't give you any consistent advice about what to predict about the colors of emeralds first observed after time T. There's nothing special here about emeralds, or their colors, or time T, so the point generalizes: Ordinary inductive reasoning doesn't give you any consistent advice about what to predict about anything.

In one sense, it's perfectly clear what has gone wrong here: When you draw the inference "I have observed a great many emeralds, and all have been green, so (probably) all emeralds are green," you do good; induction licenses that inference. But when you draw the inference "I have observed a great many emeralds, and all have been grue, so (probably) all emeralds are grue," you do bad; induction does not license that inference. Your mistake was in thinking that induction always licenses you to infer from "I have observed a great many Fs, and all have been G" to "(Probably) all Fs are G," no matter what F and G are. It doesn't.

But in another sense it is very far from clear what is going on here. Why is it that the Green induction is okay but the Grue induction isn't? What does the first one have going for it that the second one doesn't?

I'm going to try out a new way of solving this problem. If it works out, it will be very different from the most familiar ways of trying to treat that problem. For example, unlike Goodman's own solution, it won't make any use at all of the concept of projectibility; it allows that as far as induction is concerned, all predicates are created equal; what is wrong with the Grue Induction isn't that it uses a bad predicate; it's something else instead. The solution I'm going to propose is importantly different from all the other attempts I'm aware of to solve the problem too, but in the interest of time I won't go into the details of how it differs.

The solution I'm trying to work out has its roots in a strategy for solving Hume's problem of induction that was championed by D. C. Williams in his 1947 book *The Ground of Induction*. It seems to me that William's way of solving Hume's Problem just might work: At any rate, it's very impressive; it looks a bit like magic, in fact. But it has received surprisingly little attention in the intervening decades. A number of other philosophers have taken it up, including David Stove (1986), John Pollock (1990), and Timothy McGrew (2001). But it hasn't exactly taken its place in the pantheon of standard responses to Hume's problem. Why not? I suspect that a big part of the reason is this: Williams's way of solving the Old Riddle was first published in 1947, at right about the same time that the New Riddle became well-known. It is widely held that the main lesson we should take away from the New Riddle is that there cannot be a purely formal or syntactic logic of induction. But on the face of it, Williams's solution appears to be offering exactly that: a purely formal or syntactic account of the logic of induction. Bad luck for Williams, then, to have produced his solution right before the appearance of a proof that no solution of its kind could possibly work! Had Williams written, say, fifty years before Goodman, then his solution might have had a chance to have its day in the sun; as it is, it was relegated to the margins almost as soon as it was born. (This is a piece of speculative history on my part; if anyone knows more than I do about the real history here, I would be very grateful to learn about it.)

Anyhow, I think this was all a big mistake. I think that Williams's approach to solving Hume's problem, dressed up just slightly, yields a solution to Goodman's problem as well.

Williams's solution to Hume's problem aims to justify inductive inferences by appeal to something he called *the statistical syllogism*, and which he thought should be counted as part of logic. The statistical syllogism is a form of inference exemplified by this chain of reasoning: Suppose that we are justified in believing:

> 58% of the residents of Snellville are Snells;
> Doris is a resident of Snellville;
> We know nothing else about Doris that suggests she is more or less likely to be a Snell than any other resident of Snellville

then it seems we are entitled to infer that it would be reasonable for us to be 58% confident that Doris is a Snell.

For my purposes in this talk, it will be convenient to use a qualitative version of the statistical syllogism, which goes like this:

(QSS) Most Fs are G;
x is an F;
So, it is more likely that x is G than that it isn't.

The word 'likely' here in the conclusion is just the ordinary English word 'likely'; it doesn't have anything to do with the technical notion of likelihood that plays an important role in Bayesianism. The conclusion here is a judgment about relative likelihood. It is not meant to be a judgment about objective chances or objective probabilities of any kind. It expresses a doxastic state: The person making this judgment has more confidence that x is G than that it isn't, though she might not fully believe that x is G.

I think this is a good form of inference. The sense in which it's good is that if you are justified in believing both of the premises, then you are thereby prima facie justified in making the judgment expressed by the conclusion. This justification is defeasible, though. The most obvious way in which it can get defeated is for us to have additional information about x that justifies us in judging that x is less likely to be G than many other Fs are. So for example, we might be justified in believing both of the premises of this inference:

(2) Petersen is a Swede;
Most Swedes are socialists;
So, Petersen is more likely than not a socialist

and so, they give us a prima facie justification for making the judgment expressed by the conclusion. But if we also happen to know that Petersen is a member of the Stockholm Ayn Rand Society, then this justification is defeated, because we know something about Petersen in the light of which we should judge that he is less likely to be a socialist than most other Swedes. In general, an inference of the form QSS has its justification defeated whenever we (the subject drawing the inference) is justified in judging that x is less likely to be a G than some other Fs are. Let's call this rule *Defeat Rule 1;* there are presumably also other defeat rules, too, that statistical syllogism is subject to.

Some people would say that what I've just said is too lax, and that stricter standards have to be imposed on statistical syllogisms. They would say that in order for this direct inference to be justified, we really need to add an additional premise, namely:

> Petersen was selected for our consideration via some random procedure.

If all we know is that Petersen is a Swede and that most Swedes are socialists, according to this view, we cannot justify forming any opinion whatsoever about whether Petersen is a socialist. Because, for all we know, Petersen might have been selected for our consideration by a method that is most likely to select a non-socialist. In order to justify the judgment expressed by the conclusion here, we have to be justified in believing that this isn't so, and that Petersen was selected by some method whose objective probability of selecting a socialist is proportional to the fraction of Swedes who are socialists.

I think that's wrong. This is an issue that Henry Kyburg and Isaac Levi had a very interesting dispute about in the 70s; I take Kyburg's side in it. According to Kyburg, and me, the sense of randomness that's relevant to statistical syllogism cannot be defined or analyzed in probabilistic or statistical terms: It is an irreducibly epistemic concept. In particular, what it is for Petersen to be a random Swede, with respect to socialism, relative to us, is for us not to know anything that suggests that

Petersen is more or less likely to be a socialist than Swedes in general are. So long as that's the case, on the Kyburgian view, you are free to treat Petersen as a random Swede for purposes of this inference.

There's a lot to be said about this, of course. But since time is limited, let me just give you one persuasive example. I'm walking down the street, and I am approached by a cat. I wonder whether the cat would like me to scratch behind its ears or not. I think I know that not all cats like being scratched behind the ears, but most of them do. So, I infer that, most likely, this cat likes it too, and so I proceed to scratch the cat behind the ears. I don't think I just did anything irrational. But I don't know that this cat was selected and presented for my consideration by a random process. In fact, for all I know, I could be walking through a neighborhood full of people who like to take in deeply disturbed cats that hate being touched. Nevertheless, it seems obvious that it's reasonable for me to make the judgment I did. Of course, if I had extra information about this cat that suggested that it had a better-than-average chance of not liking to have its ears scratched – for example, if I knew that I met this cat during a walk through the valley of the deeply disturbed cats, then that would defeat the justification of my inference. But as it is, I don't have any such extra information, and the mere fact that I don't know that I'm not in the valley of the deeply disturbed cats shouldn't veto my perfectly reasonable inference that this cat probably likes having its ears scratched.

So, following Williams and Kyburg, I'm going to assume that statistical syllogism is a justified form of inference. Now, you might be saying to yourself, "Aha, this is where the rabbit is being snuck into the hat: In order to justify induction, obviously you're going to have to assume something that is at least as dubious as induction itself. This is where the Williams solution makes its dubious assumption." But I think that would be unfair. Notice that although the statistical syllogism is not a deductively valid form of inference, it is possible to know a priori – on the basis of just a little reflection – that the statistical syllogism is necessarily reliable, in a certain sense: In any possible world, the majority of instances of this inference pattern in which the premises are true will be instances in which the proposition that gets judged more likely than not is also true. If what we mean by a reliable pattern of inference is one that usually yields a true conclusion when fed true premises, then there is no doubt that the statistical syllogism is reliable in that sense. So to regard this as a justified pattern of inference is a far cry from, for example, boldly taking it for granted that the course of nature is uniform, or that like causes always lead to like effects, or that the simplest explanation of the facts is most likely to be correct – in making this assumption, Williams does not assume anything synthetic or contingent about the way the world is at all.

Here's how you can use statistical syllogism to justify induction. Suppose we're interested in finding out what proportion of emeralds are green. So we gather together a large sample of emeralds, and find that R% of the emeralds in the sample are green. We wonder whether it would be a good idea to infer that about R% of emeralds in general are green.

At this point, we can get some help from mathematics. A useful theorem of combinatorics assures us that, roughly speaking, most large samples from a finite population are very close to that population in their frequencies of any given characteristic. For example, suppose we have a population with 100 members, 80 of which are green. Then there are very many possible samples of 10 that we could draw from this population – in particular, there are

$$100!/90!10!$$

of them, which is on the order of 10^{13}. 32% of them have 8 members; 27% have 9 members; 21% have 7 members. Hence, most of the possible samples have greenness samples that differ from 80% by at most 10% – in particular, 80% of them do. So almost all of the possible 10-member samples we might draw from this population are quite close to the whole population in terms of their greenness frequency.

More generally: If P is any finite population, however big, and Q is any characteristic that some of the members of P have, then for any standard of closeness you like, and for any standard of "almost all" you like, there is a sample size N such that almost all samples of size N drawn from P will have a Q-frequency that is close to the Q-frequency within P as a whole. In particular, for example, suppose we let our standard of closeness be "within 2%," and we let our standard of "almost all" be "at least 97%"; there is a population size N such that no matter how many emeralds are in the emerald population we are interested in (so long as it is finite), at least 97% of the possible samples of size N drawn from that emerald population will have a greenness frequency that is within 2% of the greenness frequency within the emerald population. To be specific, N is 631[i]: We can be sure that at least 97% of the samples of emeralds of the same size as ours have greenness-frequencies that are within 2% of that of the emerald population, so long as our sample has at least 631 emeralds in it. That would be a lot of emeralds to carry around with you. But it's really not that big a sample, for serious mineralogical enquirers such as ourselves.

So, back to induction: Here we are, with our large sample of emeralds. We know from combinatorics that most large samples of emeralds have greenness-frequencies that are very close to that of the emerald population.[ii] For short, let's put that like this: Most large samples of emeralds are *green-representative*. Armed with this mathematical fact, we can construct an interesting statistical syllogism:

> Our sample is a large sample of emeralds;
> Most large samples of emeralds have green-frequencies are green-representative;
> So, most likely, our sample is green-representative.

In other words, it's more likely than not that our sample's green-frequency is very close to that of the whole emerald population. Now, we know from observation that:

Our sample's green-frequency is R%.

So now we can conclude:

So, most likely, the green-frequency among emeralds is very close to R%.

Now, in particular, suppose that *all* of the emeralds in our sample are green: then R = 100, and our conclusion is that it is more likely than not that approximately 100% of emeralds are green. That's not quite the familiar conclusion that all emeralds are green. But it's very close. And it's arguable that if we really have nothing to go on here except enumerative induction from the emeralds we've observed, the most we can really be entitled to say is that probably, *about* all emeralds are green.

So, using just a comparative statistical syllogism and a little math, we have managed to justify a canonical example of an enumerative inductive inference. We did it without assuming that the course of nature is uniform; indeed, we did it without making any synthetic assumptions about the way the world is at all. All we took for granted in addition to deductive logic and arithmetic was that the statistical syllogism is a justified way of forming judgments of relative likelihood – and as we saw before, it's a priori necessary that inferences of that form usually have true conclusions when they have true premises. So, it looks like we've pulled a justification of induction out of thin air. That's pretty nifty. As I mentioned before, it looks kind of like magic.

Let's see what the New Riddle looks like in this context. Once we have the definition of 'grue' before us, we can see that we are in a position to construct the following statistical syllogism:

Our sample is a large sample of emeralds;

Most large samples of emeralds are grue-representative;

That is, most large samples of emeralds have grue-frequencies that are very close to the grue-frequency among emeralds in general;

So, our sample is most likely grue-representative;
Our sample's grue-frequency is 100%;
So, the grue-frequency of the whole emerald population is most likely very close to 100%.

We're in trouble now, because we also think we know that most of the emeralds in the universe are not observed before T, and we know that any emerald that is not observed before T is green only if it isn't grue, and grue only if it isn't green. So, if the whole emerald population's green frequency is very close to 100%, then its grue-frequency must be less than 50% -- because if all emeralds are green, then the only emeralds that are grue are the ones that are observed before T, and *ex hypothesi* fewer than half of the emeralds are going to get observed before T. So both of the population frequencies couldn't be very close to 100%. But we have two apparently good arguments that seem to justify us in thinking that each one of those frequencies is most likely close to 100%. It seems that something must be wrong with one of those arguments. Of course, it's perfectly obvious which one is the bad one. But it's much less obvious exactly what the flaw is.

It does seem clear, though that the error has already been committed by the time we finish the statistical syllogism that kicks off the argument; once you reach the step that says that our sample's grue frequency is probably pretty close to that of the whole population, the argument is home free; the damage must already be done by the time you reach that step. So it seems there must be something wrong with this inference:

Our sample is a large sample of emeralds;

Most large samples of emeralds are grue-representative;
So, our sample is most likely grue-representative.

But why shouldn't we be perfectly justified in drawing this inference? You might say we aren't justified in believing the second premise: You might say, 'What we know about large samples of emeralds is that most of them have green-frequencies close to that of the whole population, not that most of them have grue-frequencies close to that of the whole population.' But that would be a mistake: What assures us that this premise is true is a mathematical theorem; it doesn't care about the difference between green and grue. In fact, we know both that most large samples are green-representative *and* that most large samples are grue-representative. The problem is not with the second premise. And of course, the first premises is obviously true as well.

So we have here two competing statistical syllogisms, both of which have true premises; call these *The Green Syllogism:*

Our sample is a large sample of emeralds;
Most large samples of emeralds have green-frequencies are green-representative;
So, most likely, our sample is green-representative.

and *The Grue Syllogism*:

Our sample is a large sample of emeralds;
Most large samples of emeralds are grue-representative;
So, our sample is most likely grue-representative.

Of course, the statistical syllogism is only a defeasible form of justified inference. So there is no threat of inconsistency here; at least one of these two statistical syllogisms must be defeated. Either there is some independent reason why one of them is defeated, leaving the other to stand, or else the prima facie availability of each one of them defeats the other one.

Of course, we know in our hearts that the Green Syllogism is all right, so there must be some reason why the Grue Syllogism gets defeated. What is that reason? What consideration defeats the Grue Syllogism while leaving the Green Syllogism standing? If we can answer that question, then we can solve the New Riddle.

The interesting thing here is that we have managed to transform the original New Riddle of Induction into a problem that really has nothing to do with induction as such: It's a problem about why a certain statistical syllogism gets defeated. And the very same problem can be illustrated with examples that have nothing to do with induction at all.

The general form of this problem is like this:

We are defeasibly justified in believing the premises of two different statistical syllogisms:

x is an F; most Fs are G; so most likely x is G

x is an F; most Fs are G'; so most likely, x is G'

But we also know that at most one of the conclusions of these two syllogisms is true. In this situation, we cannot be, all things considered, justified in accepting both syllogisms. Whenever you face a situation that takes this form, it is possible that one of these two syllogisms has its justification defeated by some flaw it has that it doesn't share with the other one. In that case, no problem; the flawed one falls and the other one stands. A second possibility is that there is nothing independently wrong with either of the two syllogisms – except that it happens to conflict with this other one that looks equally good. In that case, it seems to me, the thing to say is that the two syllogisms cancel each other out: The otherwise-undefeated prime facie justification of each syllogism defeats the justification of the other one. So there is a general problem here that has nothing in particular to do with induction: It's the problem of saying under what conditions a statistical syllogism gets defeated. If we could solve that general problem, then we could say what it is that's wrong with the Grue Syllogism that makes it worse than the Green Syllogism. And once we learn that, we'll be able to say what's wrong with the Grue Induction that isn't wrong with the Green Induction. So here is a possible path to a solution to Goodman's problem.

To make it plausible that this general problem is an important problem in its own right, and that it is completely independent of worries about induction as such, let's look at another example that has the same basic form.

Suppose we have a big urn full of balls, and we know that a large majority of the balls in the urn are pink, though not all of them are. This is our plan: We are going to give the vat a good hard, long shake, and then one of us is going to reach in without looking and pull out a ball.

It seems clear that before we carry out our plan, we are justified in forming the judgment that it's most likely that the ball we draw is green. This can be supported by a statistical syllogism, that's on your handout, labeled *the Pink Statistical Syllogism*:

> Our ball is from this urn;
> Most balls from this urn are pink;
> So, it is more likely than not that our ball is pink.

That seems clear enough.

Now, suppose that our plan has a bit more to it than I told you about at first. After the ball is pulled out, we're going to put it in a special basket we have with us. We're going to leave it there for a little while, and then put it back in the vat. No other ball is ever going to go into this basket.

Now we can define a new predicate:

> x is *pinknobasket* =def x is pink iff x never goes in our basket

or equivalently:

> x is pinknobasket =def (x is pink and never goes in our basket) or (x is not pink and at some time it does go in our basket)

It follows from this definition, and from what we know about our own plan, that the ball we draw will be pink iff it is not pinknobasket.

We know that a large majority of the balls are pink; what do we know about how many of the balls are pinknobasket?

Well, a large majority of them are pink, and at most one of the pink balls isn't pinknobasket: Every pink ball that isn't the one we're going to draw is pinknobasket. So, we know that at least a majority of the balls in the urn are pinknobasket.

With this information in hand, we can draw another direct inference:

> Our ball is form this urn;
> Most balls in the urn are pinknobasket;
> So, it is more likely than not that the ball we draw is pinknobasket.

Let's call this inference *the Pinknobasket Syllogism*.

Now we have something of a paradox. We know all the premises of both of these syllogisms are true. One tells us to judge it more likely than not that the ball we draw will be pink, and the other tells us to judge it more likely than not that the ball we draw will be pinknobasket. But since we know the ball we draw is going in the basket, we also know that our ball is pink iff it is not pinknobasket. So, we can't consistently make both of the judgments these two statistical syllogisms appear to justify.

What's going on here is that we have two statistical syllogisms that are clashing with one another; given our background knowledge, we know that their conclusions can't both be true, though we also know that all their premises are true, and that they have the form of good statistical syllogisms. This is exactly what was going on in the case of the emeralds; the two cases have exactly the same structure. The difference is that here, we are looking at individual balls rather than large samples of emeralds, and this example has nothing to do with induction. But it seems very clear that we have here just two different instances of the same general problem. There is every reason to hope that if we find the solution to one of these problems, it will give us the solution to the other one.

So, how should we go about trying to solve the Pinknobasket Problem? Let's start by stepping back and surveying the logical space of possible solutions.

1. Predicate elitism. You might think that the problem with the Pinknobasket Syllogism is that it uses the predicate Pinknobasket, which is not the kind of predicate that you want to allow into your family. The Pink Syllogism, on the other hand, uses the predicate Pink, which is a perfectly respectable predicate. Statistical syllogisms with nice predicates in them are justified, but statistical syllogisms that use sketchy, shady predicates like Pinknobasket are not. The sketchiness of their predicates serves to defeat their prima facie justification – or perhaps, precludes

them from even having any prima facie justification in the first place.

If you go for a predicate-elitist solution to the Pinknobasket Problem, it's going to lead naturally to a predicate-elitist solution to the Grue Problem. And of course, most of the popular attempts to solve the Grue problem in the literature – starting with Goodman himself – are predicate elitist solutions. Having decided to be a predicate elitist, you still have some choices to make: Because there are a lot of different ideas you might have about what makes some predicates the elite ones. Goodman famously made eliteness of predicates depend on the past history of human science: The elite predicates are the ones that have been successfully projected many times in the past. Quine, David Lewis, and others have thought instead that what makes the elite predicates elite is something objective: They pick out natural kinds, or carve nature at the joints, or something like that.

Another choice you still have to make is whether to be what I'll call a *simple elitist* or a *confrontational elitist*. A simple elitist says that there's a certain threshold of goodness that the predicates in a statistical syllogism have to meet in order for that syllogism to be any good at all. A confrontational elitist says that in principle, any predicate might appear in a perfectly good statistical syllogism. But whenever we are in the situation where two different statistical syllogisms are in competition with each other relative to our background knowledge, the syllogism with better predicates wins: It gets to be justified, and not the other one. If the two syllogisms feature predicates of equal quality, on the other hand, then they cancel each other out. Incidentally, though I think this point is rarely noticed in the literature, I think Goodman himself is a confrontational elitist.

Unfortunately, I don't think any of these versions of predicate elitism is going to help with the Pinknobasket Problem. This isn't because I don't believe that all predicates are created equal. For all I know, nature might really have joints, and some predicates might really carve nature along them; but even if it's true, this won't help solve the Pinknobasket Problem.

To see why, first of all consider an extension of the Pinknobasket Case (This is the "Extended Pinknobasket Case 1" on your handout): After we draw our ball, and put it in the basket for a while, and then take it out and put it back in the urn, we go away, and then later in the day Adam Cureton comes along, and he draws a ball randomly from the vat. Now suppose you are interested in the question of whether the ball James draws will be Pinknobasket or not – maybe someone has offered you a bet at even odds on the question, and you're trying to decide whether to take the bet. You are tempted to reason as follows:

> The ball Adam draws is a ball from this urn;
> Most of the balls from this urn are pinknobasket;
> So, most likely, the ball Adam draws is pinknobasket.

That's plainly good reasoning. You should use it, and endorse its conclusion. If you hesitated on the grounds that pinknobasket is a weird predicate, you would lose your chance at a great bet. And that would just be plain silly. In fact, it seems obvious that the reasoning in this syllogism is exactly as good as the reasoning in the Pink Statistical Syllogism.

What this shows is that simple predicate elitism cannot be true. But it doesn't show that confrontational predicate elitism can't be true: In the extended Pinknobasket case, the syllogism about Adam's ball doesn't get into a confrontation with another apparently justified statistical syllogism. So if you are a confronational predicate elitist, then you can consistently say that in the original case, the Pinknobasket Syllogism is defeated, since it comes into conflict with the Pink Syllogism and the Pink Syllogism uses nicer predicates, but in the extended case, the Pinknobasket Syllogism doesn't get defeated. Because it doesn't get into any fights with other syllogisms.

So we need a different case to show what's wrong with confrontational predicate elitism. It's called "The Strange Cave Case."

Suppose that while exploring a canyon in a remote desert, we discover a large cave, filled with artifacts. We send our assistants in with flashlights to make a thorough inventory of the artifacts in the cave. They come out and give us a summary of what they found: There are all kinds of different artifacts in the case: Cars, books, clothing, fountain pens, costume jewelry, coffee pots, musical instruments, and so on and so forth. They seem to have little to do with each other. However, two things stand out: First, approximately 90% of the artifacts in the cave are green; second, approximately 90% of the artifacts in the cave were formerly owned by one Henrietta Jackson of White Plains, NY. Intrigued, we do a little research on Ms. Jackson, and we learn many interesting facts about her, one of which is that though she has owned many green objects, and she has owned many motorcycles, she has never owned a green motorcycle.

Later on, one of our assistants brings an object out of the cave, covered by a sheet so that we can't see it. She tells us that the object is a motorcycle. A compulsive gambler among our party then offers to make some bets with us: He gives us a chance to bet on whether the motorcycle under the sheet is green, and he gives us a chance to bet on whether the motorcycle was formerly owned by Henrietta Jones.

Reflecting on what we now know, we see that we know all the premises of the following two statistical syllogisms:

Most of the objects found in the cave are green
The motorcycle is an object found in the cave
So, most likely, this motorcycle is green

Most of the objects found in the cave were formerly owned by Henrietta Jones
The motorcycle is an object found in the cave
So, most likely, this motorcycle was formerly owned by Henrietta Jones.

But we know that we cannot believe the conclusions of both syllogisms, because we know that Ms. Jones has never owned a green motorcycle. (By the way, this story is set in a possible world in which it is impossible to paint a motorcycle a new color.) So, we have these two statistical syllogisms, each of which looks pretty good, and each of which looks quite relevant to the question of whether we should take the bets. However, we cannot trust both of these syllogisms; given what else we know, we're going to have to reject one of them. Which should we reject?

Unfortunately I haven't done any experiments (yet!) to find out what the folk think. Let me just report my own considered judgment about this case, and urge you to be sensible and agree with it. We should reject them both, and we should be no more than 50% confident either that the motorcycle is green or that it was formerly owned by Ms. Jones. We've just got nothing to go on here. It's probably one or the other, but it can't be both, and we've got no more reason to think it's one than we do to think it's the other. But by anybody's standards, green is a much more natural predicate than "formerly owned by Henrietta Jones of White Plains, NY." So according to confrontational predicate elitism, the green-motorcycle syllogism should win this confrontation. But it doesn't; it's a draw. So, confrontational predicate elitism is wrong. It cannot be right that which of two conflicting statistical syllogisms prevails is a function of which one uses more natural predicates.

If you're not going to be a predicate elitist, then how are you going to solve the problems before us? The most obvious way to go is to be what I'll call a *relationist*. A relationist says that some statistical syllogisms get defeated because they use a predicate G that is somehow inappropriately related to the object x that the syllogism is about. It's not that there's anything wrong with G as such; it's just that in the case at hand, it bears some special relation to x that makes it inappropriate to use in a statistical syllogism about x. It might be perfectly fine to use it in other statistical syllogisms, though.

Relationism seems like a promising way to go to solve the Pinknobasket problem. For the predicate "Pinknobasket" obviously is related in a pretty specific way to the particular ball we're going to draw out of the urn. It's defined in terms of going in the basket, and the particular ball we're going to draw is distinguished from all the others by the fact that it is going to go in the basket. That's a suspicious connection. In the extended pinknobasket case, when we consider the syllogism whose conclusion is that the ball Adam draws will most likely be pinknobasket, things are different: Pinknobasket is still a weird predicate, but it is no longer connected in any special way with the ball in question. For we have no particular reason to think that Adam's ball ever goes in the basket.

Relationism also seems like a promising strategy for dealing with the Grue problem. In the Grue Syllogism, the predicate G is "grue-typical," and the item that plays the role of x is our sample of emeralds. One thing that's special about our sample of emeralds, which isn't true of most samples of emeralds in the world, is that it consists entirely of emeralds that were first observed before T. And grue-typicality is a property that obviously has something important to do with being observed before T. So there, too, its plausible that's what's going wrong is that we're drawing a statistical syllogism of the form "x is an F; most Fs are G; so, most likely, x is G," where the particular x we're talking about is distinguished from most of the other Fs by being related in a conspicuous way to the predicate G.

If we go for relationism, though, we still have another question to answer: What is the particular relation between x and G that undermines the justification of a statistical syllogism? Maybe the most obvious answer is given by what I'll call *counterfactual relationism:*

The statistical syllogism … is unjustified whenever we know that x has a special feature H that distinguishes it from all or most of the other Fs, such that whether or not x is G depends counterfactually on whether it is H.

That seems to get the Pinknobasket Cases right: We know that the ball we will draw will go in the basket, which none of the other balls in the urn will; we also know that whether our ball is pinknobasket depends counterfactually on whether it is going in the basket. Later on, when Adam draws his ball out, things are different: We don't think that whether or not Adam's ball is pinknobasket depends on whether or not Adam draws it out of the urn or not.

Unfortunately, this thought runs into trouble. Let's extend the story a bit: You are the one to draw the ball from the vat; you stick your hand in, grab one of the balls at random, and start to pull it out. Before you have caught a glimpse of the ball you are drawing out, a trustworthy friend stops you. "Wait," she says, "I have new information to give you. If you had chosen any other ball than the one you did just choose, then all the balls in the vat would have had their colors changed just a second ago – the pink ones would have become non-pink and the non-pink ones would have become pink. As it is, though, none of the balls underwent any color changes." Why should that strange counterfactual be true? Well, here is one way the story might go: Standing nearby is an eccentric wizard. The wizard has a favorite ball in the urn – let's call it ball z. Your informant has no information about what color ball z is. The wizard has been standing by watching you, planning to play a trick on you: She was going to cast a spell that switched the colors of all the balls in the urn just as you were drawing your ball out. However, she decided, if by some chance you happened to choose ball z, her favorite ball, then she would not do anything – she would leave all the balls alone. As it happened, you chose ball z. So the wizard did nothing; the colors of the balls have not changed. But if you had chosen any other ball instead, then all the balls would have changed their colors. Again, your informant has no information about what color ball z is.

So now, you have this additional information, and you have not yet looked at the ball you have chosen, though you can easily do so. Does this additional information you have make any difference to what you ought to expect the color of the ball to be when you look at it? It seems clear that the answer is No. Beforehand, you knew that you were randomly drawing a ball from an urn containing a large majority of pink balls. Your new information doesn't tell you anything new that bears on the question of what color the drawn ball is. You know now that the ball you have drawn is ball x, the wizard's favorite ball, but that isn't relevant to its color. You know if things had gone differently, then the balls would have had their colors changed, but as it is, the colors of the balls didn't change. So it seems that this new information does nothing to undermine your original inference; before you had this new information about the wizard, it was reasonable for you to think that the ball is most likely pink, and now that you have this new information about the wizard, it is still just as reasonable for you to think that the ball is most likely pink. Your new information does not defeat the Pink Syllogism. So, this is a counterexample to counterfactual relationism.

There are a lot of other solutions you might try, but since we're running out of time, let me skip to the one I like the best. I'm going to call it the *epistemic relationist* solution. The reason why the predicate "pinknobasket" is inappropriately related to the ball we draw out is that we know something special about our ball, which we have to find out in order to find out whether our ball is pinknobasket. We know that our ball is going in the basket at some point. And there is no method available to us for finding out whether something is pinknobasket that doesn't involve finding out whether it ever goes in the basket, and then using this information in a calculation. The only way we have of finding out whether something is pinknobasket is by finding out whether it is pink, and finding out whether it ever goes in the basket, and then doing the appropriate truth-table. There just is no other way in which pinknobasketness can manifest itself to us.

More formally:

… is defeated if we know that x is H, that all or most of the other Fs are not H, and whether an F is G or not *depends epistemically, relative to us*, on whether that F is H – in the sense that there is no method in principle available to our epistemic community for finding out whether that F is G that does not involve finding out whether that F is H and inputting that information into a calculation whose output depends on whether that F is H.

Why should we think that this view is right? Well, when this condition is met, then from our epistemic point of view, x is not just like any old F, insofar as the question of whether it is G is concerned. Rather, it has a feature that is special among Fs with respect to whether it is G. So, when we're inquiring about whether x is G, it isn't epistemically kosher for us to treat it as if it were just some generic F. But that is exactly what you do when you run the statistical syllogism: You treat x as if it were just like any other F. Since most of them are G, most likely x is too. So, since, when this condition is met, it isn't okay to treat x as if it were just like any other F for purposes of inquiring about whether it is G, the statistical syllogism is not justified.

That explains why the Pinknobasket Syllogism gets defeated, but the Pink Syllogism is allowed to stand: We can find out whether a ball is pink without finding out whether it ever goes in the basket, but we cannot find out whether a ball is pinknobasket without finding out whether it ever goes in the basket. So the Pinknobasket Syllogism gets its prima facie justification defeated, but the Pink Syllogism doesn't. This also explains why the syllogism … that we draw later on does not get defeated: We can find out whether a ball is pinknobasket without finding out anything about that distinguishes the ball Adam draws out from the rest of the balls in the urn. This also explains why the scenario with the wizard doesn't change anything: All the wizard changes is what the balls would have been like in some non-actual situation; it makes no difference to the fact that we cannot find out whether a ball is pinknobasket independently of finding out whether it ever goes in the basket.

So epistemic-dependence relationism seems to get all these cases right. It also gets the Grue Case right. In the Grue Syllogism, the role of x is played by our sample of emeralds, and the role of G is played by the predicate 'is grue-typical.' Now, there is no way for us to find out whether a sample is grue-typical, even in principle, that doesn't involve finding out whether each of the emeralds in it is grue or not. But we cannot find out whether any emerald is grue without finding out whether it was observed before time T or not. But if we find out whether each emerald in a sample was first observed before time T, then we thereby find out whether the whole sample has the special property of consisting entirely of emeralds that were first observed before T. And of course, we know that's a special property that our sample has. Thus, the Grue Statistical Syllogism is defeated.

But the Green Syllogism doesn't get defeated in this way: To find out whether a sample of emeralds is green-typical, we first need to find out what fraction of the emeralds in the sample are green, and then we need to find out what fraction of emeralds in general are green, and then we need to compare the two fractions. How could we find out what fraction of emeralds in general are green? Perhaps we could do it inductively. It might be thought that it would be inappropriate here for me to assume that we are able to discover this inductively, since the justification of induction is one of the things supposedly under investigation. (I doubt that this really would be inappropriate, but let's not get hung up on this question here, since we don't need to.) There is another way in which it could be done, in principle if not in practice, and that is by exhaustively examining all the emeralds. Now, whichever of these methods we used, it's clear that all we would really have to find out about each individual emerald that we observe is whether it is green or not. It makes no difference what time any of those emeralds were first observed at. (Well, all right, they all need to have been observed by the time we carry out this investigation, but the information about when they were observed plays no role in the determination of whether the emeralds are green, the way it does in the determination of whether they are grue.) So it seems that we can (in principle, anyway) find out whether a given sample of emeralds is greenness-typical without finding out whether it consists entirely of emeralds first observed prior to T.

Thus, according to epistemic-dependence relationism, the Grue Syllogism is not justified whereas the Green Syllogism is. Therefore, we can use the Green Syllogism to underwrite the Green Induction, borrowing Williams's strategy. But since the Grue Syllogism is not justified, we cannot use it together with Williams's strategy to underwrite the Green Induction. Thus, we appear to have solved the new problem of induction along with the old one.

Notice a few unusual things about this solution. First of all, it makes no appeal whatsoever to a distinction between projectible and unprojectible predicates, or between natural and unnatural properties. Second, unlike the solution proposed by Frank Jackson and Robert Pargetter, it makes no appeal to our presumed antecedent knowledge of counterfactuals. Furthermore, unlike Bayesian solutions to the problems, it makes no appeal to our prior probabilities.

It does make an appeal to a feature of ourselves, though, for it refers to the set of methods for finding things out that are in principle available to our epistemic community. I think this relativization to our epistemic community is perfectly in order: What we need to explain, in order to solve Goodman's problem, is why *we* are justified in drawing the Green Induction but not the Grue Induction. If we were put together differently, so that we could tell just by looking at an emerald not only whether what color it was, but also whether it was grue or not, without finding out anything about when it was first observed, then the view I'm proposing implies that the Grue Induction would be just as justified as the Green Induction is. So, they would cancel each other out, and we wouldn't be justified in making any prediction about the colors of future emeralds. But I think that's all right. If we had eyes that could detect grueness directly, the world and its laws and its causal relations would have to be very different from the way we think they are: Whether something is grue or not would have to be able to play a role in the laws and the causes that we don't think it actually plays. Under those conditions, who knows? Maybe 'grue' wouldn't *be* such a disreputable predicate. Thanks.

Appendix: Excised material about projectibilism and why it should have seemed wrong to begin with

The most familiar ways of trying to come to terms with the problem involve the assumption that there is a distinction between projectible predicates and non-projectible predicates – or else a hierarchy of degrees of projectibility – and a formally nice inductive inference is justified only insofar as the predicates involved in it are projectible. Some versions of this approach – like Goodman's own solution – make projectibility depend on us in some important way; others – like David Lewis and David Armstrong – make it an objective matter. But the solution I'm trying to work out assigns no role to projectibility at all. It says that what's wrong with the infamous bad inference from the grueness of all observed emeralds to the grueness of all emeralds is not that there is anything wrong with grueness – instead, it says that what's wrong with grueness is something about the relation between grueness and the particular context in which the bad inference is made. Roughly speaking, what's wrong with the inference is that grueness is connected in an inappropriate way to a special feature of our sample of emeralds, namely the feature that all of the emeralds in it were first observed prior to T. So it's not anything about grueness per se that makes the inference bad, and it isn't even anything about the way grueness is related to us; instead it's something about the way that grueness is related to the particular context in which the offending inference is drawn.

It seems to me that it should have been plausible all along that something like this is the case, because it seems clear that there can be perfectly fine inductive inferences in which the predicate 'grue' plays a crucial role. For example, imagine a time in the far future, the 25^{th} century say, long after the crucial time T that figures in the definition of 'grue' has come and gone. Suppose that in fact, all emeralds are green – so, many of the ones that have been observed by the 25^{th} century are grue, and many others are not grue. Now imagine that a group of future mineralogists have a very large vat full of emeralds, and they want to know what fraction of the emeralds in the vat are grue. (Suppose that they have some way of telling, for any given emerald, whether it was first observed before time T or not; maybe by this time in this distant future all emeralds have had serial numbers engraved on them, and from an emerald's serial number you can tell whether it was first observed before T.) It seems that a perfectly reasonable way for them proceed would be to take a big-enough random sample of emeralds from the vat, see what fraction of the emeralds in the sample are grue, and infer that about the same fraction of emeralds in the whole vat are grue. That would be a perfectly standard statistical sampling procedure, and in this case, it would clearly be perfectly in order, even though 'grue' is a funny sort of predicate. But they would be, in effect, drawing an inductive inference from the grueness frequency within a sample to the grueness frequency of the population from which the sample was drawn – which is precisely what the infamous Grue Induction does. So it can't that there is anything wrong with "projecting grueness" from the emeralds in our sample to the emeralds outside of our sample. Rather, the mistake in the Grue Induction is something like, projecting a predicate that is inappropriately related to the sample you are projecting from.

Appendix: Objection to the Williams Strategy and a Reply

You might object: But even if statistical syllogism is necessarily reliable in the sense that most sets of true premises fed into it would lead to a true conclusion, that's not the sense of reliability that's epistemologically interesting; when we want to know whether we ought to trust a certain pattern of reasoning, what we really want to know is whether it usually leads to a true conclusion *on those actual occasions when we use it correctly*. Now even if "most instances" of this inference pattern with true premises lead to a true conclusion, the "instances" being talked about here are abstract possibilities of using the inference pattern – they aren't actual concrete cases in which one of us actually uses it. It could well be the case that on most future occasions when one of us correctly relies on this inference pattern, it will lead us to a false conclusion. When Williams assumes that it is safe to count on statistical syllogism, he is in effect assuming that the world isn't like that, and that is a synthetic assumption about the way the world is that he has no obvious title to.

This must be conceded. But it isn't true that the sense of reliability that is interesting for epistemological purposes is the one the objection makes use of. When do we actually, epistemically correctly make use of a given pattern of inference? Not when we notice that its premises are true; rather, when we are justified in believing that its premises are true. And no matter what pattern of inference you pick, it is possible that in the future, on most occasions when we are justified in believing that the premises of an instance of it are true, the conclusion will in fact be false (because one of the premises is in fact false; our evidence is misleading). This is true even of the most basic deductive patterns of inference, like modus ponens or disjunction introduction. So the standard that the objection holds induction to cannot even be met by basic patterns of deductive inference. That standard is too high: The kind of reliability that is relevant to justification is not reliability over all concrete instances in which we actually justifiably use an inference pattern – for not even deductive logic is reliable in that sense – but instead reliability over all possible instances in which its premises are true – and in that sense, comparative statistical syllogism is obviously, a priori, necessarily reliable. So no questions are begged if we take for granted the justification of statistical syllogism for the purposes of solving Hume's problem.

Appendix: Some Objections and Replies

But hang on a second: Someone might want to ask here just what it means to "find out" whether a ball is pink, or whether it is pinknobasket. Surely "find out" doesn't mean, come to know with absolute certainty – in that case, we would hardly ever be able to find out anything in the relevant sense, and so the rule I'm proposing would be trivial, since it would count every predicate as inappropriately connected with every other. On the other hand, suppose that to "find out" that a ball is pinknobasket just means to find evidence that makes it quite probable that it is pinknobasket. In that case, since almost every pink ball is also pinknobasket, and almost every non-pink ball is also non-pinknobasket, it seems that we can "find out" whether a ball is pinknobasket just by looking at it in a good light and seeing whether it is pink. But that would mean that we can, ever after all, find out whether a ball is pinknobasket without also finding out whether it ever goes in the basket. This threatens to wreck things for my proposal. (Thanks to Dan Parker for raising this objection.)

I think the thing to say here is that what it takes to "find out," in the sense that's relevant here, that a ball from the urn is G is to find evidence that justifies us in being more confident either that the ball is G or that it isn't G than we already were that this ball is not the one we will draw on this occasion. If there are 100 balls in the urn, then for any given ball, we can justly start out being confident to degree 0.99 that this ball is not the one we will draw out; so in order to "find out" that a ball is pink, or pinknobasket, we've got to find evidence that makes us more than 99% confident that it is. If we look at a ball and see it to be pink, then I assume this makes us considerably more than 99% confident that it is pink; but this same evidence assures us only that the ball is pinknobasket iff it is not the one we draw on this occasion. So our degree of belief that the ball is pinknobasket should be equal to our new degree of belief that it is pink multiplied by our confidence that it is not the ball we choose – and that product cannot be greater than our confidence that it is not the ball we choose. Thus, in the relevant sense, we can indeed "find out" that a ball is pink just by looking at it, but we cannot "find out" that it is pinknobasket in the same way.

Here is another important worry: Presumably, it is impossible to find out that a given ball is pink without, in effect, finding out that that ball is thought about by someone at some time -- I mean that it is thought about on its own, that it is the object of a singular thought. So, the property of being pink is inappropriately connected with the property of being thought about by someone at some time. Now, suppose that I have good reason to believe that most of the balls in the vat will never be the objects of singular thoughts: For most of those balls, there will never be an occasion for somebody to single out that ball in particular and think about *it*. So the predicate of being thought about by someone at some time can play the role of H: We know that the ball we choose will be thought about at some time, most of the balls in the urn won't, and it is impossible for us to find out that a ball is pink without finding out that it gets thought about by someone at some time. This seems to show that the defeat rule I have proposed shows that the Pink Syllogism gets defeated as well as the Pinknobasket Syllogism; this is no good. (Thanks to Ram Neta for raising this objection.)

To respond to this objection it's necessary to refine the proposal a bit. The thought was that G is inappropriately connected to a special feature of the ball we draw just in case in order to find out whether the ball is G, you have to find out whether it has that feature, since the information about whether it does or not bears on the question of whether it is G. So, for example, we can't find out whether a ball is pinknobasket without finding out whether it ever goes in the basket, since the only way for us to find out whether it is pinknobasket involves drawing an inference, one premise of which says whether it is ever in the basket. In the example we just considered, the feature of being thought about isn't like that: We don't have to find out whether a ball ever gets thought about because we need this information in order to determine whether the ball is pink; instead, it's just a sort of inevitable side effect of finding out that a ball is pink that you end up finding out that it was thought about. So to fix the problem, what we need to do is specify that the defeat condition applies only when the fact about whether the ball has the feature H is a crucial premise in an inference, such that the way we find out whether the ball is G depends for its justification on the goodness of that inference:

> Whether x is G is inappropriately connected with whether x is H just in case we cannot find out whether x is G without using a method the reliability of

> which depends on the reliability of an inference in which the fact of whether x is H is a crucial premise.

While it is true that you cannot find out whether a ball is pink without thereby finding out whether it ever gets thought about, you can find out whether a ball is pink without relying on any method that depends on an inference that has “the ball is thought about at some point” as a premise. So this revision to the proposal gets around the worry.

The same revision pre-emptively gets us around another worry that I haven’t raised yet: Surely, you might think, it is *possible* for one of us to find out whether a ball is pinknobasket without finding out whether it ever goes in the basket? For example, we might design a robot with a color-detecting sensor to go through the balls in the urn to detect their colors, examine all the available evidence concerning which of the balls ever go in the basket, and them stamp a big black “P” on all and only the balls that are pinknobasket; then we could find out whether a given ball was pinknobasket simply by checking to see whether it has a “P” on it. If we did this, then we would not be drawing any inference using a premise to the effect that the does or does not go in the basket at some point. And – if we wore special lenses that made it impossible for us to tell whether the ball was pink – then we could do it without even finding out whether the ball ever goes in the basket or not. So doesn’t this show that whether a ball is pinknobasket is not, after all, inappropriately connected with whether it ever goes in the basket? No, it doesn’t – not if we adopt the refinement I just proposed: By relying on the robot, we are relying on a complex method for telling whether a ball is pinknobasket, part of which is the design and building of the robot. The reliability of that method depends on the fact that the inference from “x is pink and never goes in the basket” to “x is pinknobasket” is itself a reliable one. It is important that the refined proposal does not specify the condition that we cannot find out whether x is G without drawing an inference from whether it is H; instead, it says that we cannot find out whether x is G without relying on a method the justification of which depends on the reliability of some inference from whether it is H. In the robot case, we don't draw an inference from whether the ball is ever in the basket, but we do rely on a method whose justification depends on the reliability of such an inference. So the refinement already proposed takes care of this worry.

Bibliography

Goodman, Nelson 1955: *Fact, Fiction and Forecast*. Cambridge, MA: Harvard University Press.

McGrew, Timothy 2001: "Direct Inference and the Problem of Induction," *Monist* 84(2): 153-178.

Stove, D.C., 1986: *The Rationality of Induction*. Oxford: Oxford University Press.

Williams, Donald 1947: *The Ground of Induction*, Cambridge, MA: Harvard University Press.

---

[i] For example: Suppose that the population size N is large enough that the normal distribution with mean nr and variance $(nr(1-r))^{1/2}$ is an acceptably close approximation to the binomial distribution for N, n and r. (For N at least as great as 100, the approximation here is very good; see e.g. Feller 1968, p. 188.) The population of ravens is surely greater than 100, so in the ravens example, this is a reasonable thing to suppose.) Then so long as $n \geq D^2/4e^2$, where D is the smallest positive number such that $\mathbf{N}(-D) \leq \frac{1}{2}(1-r)$ (**N** is the standard normal distribution), the fraction of the n-sized samples from P whose A-frequencies are within e of P's A-frequency is greater than r. (See for example Johnson and Bhattacharyya 2001, pp. 328-329.)

To illustrate, suppose that P is the population of emeralds, and A is the characteristic of being green. Now suppose that r = 0.97, and that e = 0.02. Then by consulting a table of the standard normal distribution we find that D = 1.004. Plugging these values into the above inequality, we get:

$$n \geq (1.004)^2/(4 \times 0.0004) = 630.01$$

This justifies the claim about to be made in the text.

[ii] What if the emerald population is infinite? The combinatoric facts I was just mentioning apply only to finite populations; in the case of an infinite population, "almost all" won't in general be well-defined. Isn't this a problem? -- Maybe, but it can be easily circumvented. Suppose that there are infinitely many emeralds, but only finitely many of them within any finite spatial region. Then the reasoning about to be presented in the text will justify an inductive inference to the conclusion that about R% of emeralds within S are green, where S is any finite spatiotemporal region that includes our sample. Since any emerald in the universe is within some such region, every emerald is within a finite population over which we have a justified inductive inference. So in this case, the inductive inferences that are justified don't include one to the conclusion that about all of the infinitely many emeralds are green, but they do include inferences to conclusions that bear on any particular emerald you like. We will be inductively justified in believing, of each emerald, that it is green, even if not in believing the universal proposition that all of them are. I confess to finding this a little disappointing, but only a little.