RRIS workshop on 26-28 July 2013 at Gladstone's Library, Wales:

How Do or Should Religious Beliefs Impact our Scholarship?

Presuppositions, faith and statistics: an ecologist's view

Richard Gunton

Abstract

A common theme in the history of statistics is the desire to know the true value of a quantity that can only be measured with error. Statistical inference therefore concerns the numerical accuracy of beliefs. In many natural and social sciences inferential methods are dominated by the school of frequentism, which arose early in the 20th century. In this paper I argue, following Andrew Hartley, that frequentism embodies a reductionistic motive whereby beliefs are reduced to mathematical inequalities arising from the numerical nature of data. The problems of this reductionistic approach, however, tend to elicit a humanistic reaction in the "indirect frequentist" approach to inference, where expert judgement intervenes in arbitrary ways. A more promising alternative has long been offered by the Bayesian school of statistical inference. Again following Hartley, I will argue that the Bayesian paradigm, in which inferences are explicitly based on prior beliefs, is a logically and experientially coherent approach to inference. Because it is not inherently reductionistic, I will also argue that it does better justice to the multi-faceted nature of reality and is therefore more consistent with a theistic worldview as advocated by the neocalvinist Christian tradition.

Introduction

The science of statistics has diverse origins, but a common theme is the desire to know the true value of a quantity that can only be measured with error, and to use this knowledge for technical purposes involving decision-making. The discipline of statistics therefore concerns both the accuracy of beliefs and the justification of actions, and is now widely used in diverse areas of science, engineering and policy. This paper concerns the first of these two goals: inference. The approach to statistical inference that currently predominates in many natural and social sciences is based on the school of frequentism, which arose early in the 20th century in the context of evolutionary biology. The older tradition associated with Bayesian inference, however, concerns the modification of prior beliefs based on the evaluation of new data. This paper argues that not only does the Bayesian approach reflect everyday common reasoning, but it is also more logically coherent and more consistent with a Christian theistic worldview.

First I will offer a definition of "religious belief", to which I return at the end of this paper in assessing different approaches to statistics with respect to a Christian theistic worldview. Next I give a very brief introduction to the purpose of statistical inference, outlining the Bayesian and frequentist approaches. Then I give two short primers on the approaches each of these schools may take for dealing with a simple inference problem, pointing out some of the key assumptions involved. Finally I mention ethical and practical concerns, and discuss more explicitly the religious nature of these approaches.

Religious beliefs

Defining "religious" is notoriously difficult. A simple but broad definition of "religious belief" would be any basic belief that is foundational to thinking and action. A more technical definition that I wish to consider is Roy Clouser's (Clouser 2005): a belief is religious if it concerns what a person holds (consciously or not) to be ultimately and non-dependently real – which for Clouser defines "the divine" for that person – or how one relates to that divine. We might consider how Christian theology proposes the God revealed in Jesus Christ to be before all things, the source of all creation while being separate from it, whereas faiths classified as "pagan" seem to propose that divine spirits are located within some part of the world. More pertinently for us, Clouser's view entails that philosophical theses about the nature of reality are also religious beliefs. In particular, reductionistic ontological claims are religious, in that they posit one aspect of reality as constitutive of another. I realise that that definition of "reductionistic" may need clarification; I mean by it to describe radical claims that some aspect of experience can be explained in terms of another aspect of experience, when the two aspects have different statuses in everyday thought (e.g. they can appear in separate considerations used in decision-making, like feelings, economic reasons and law) and in discourse (e.g. the vocabulary of physical forces is rarely conflated with the vocabulary of sensations, however much we may be told that sensations are really neural impulses). I realise I have not given an absolute definition, but this may be appropriate, because I also suggest that "religious" may apply to any claim that turns out to contravene a belief genuinely held by *someone* about the ultimate nature of reality.

A word must be said about the background to Clouser's definition. The philosophy of the cosmonomic idea, or Reformational philosophy, is a tradition drawing on the work of Herman Dooyeweerd (Dooyeweerd 1953), which recognises that all scientific work is based on abstractions from human experience: the consideration of a certain type of properties that pertain to things or situations encountered in lived experience, to the exclusion of other types of properties (Strauss 2011). So, for example, the science of physics concerns the physical properties of things – and, we might say, the physical aspect of experience – to the exclusion of biotic, sensory, linguistic, social and other properties and aspects – which are each the concern of another kind of science (Clouser 2005). However, no-one can conceive of a physical (or any other) property of a thing in isolation from any other aspect (Dooyeweerd's "transcendental critique", Clouser 2005), and the practice of science invariably involves human functioning in all kinds of aspects simultaneously. For present purposes, we particularly note that most natural sciences depend upon numerical analyses (even though mathematics is the proper science that analyses numerical relations *per se*) and that all scholarship is concerned with the accuracy of opinions and convictions, insofar as these pertain to the goal of knowledge. Clouser's definition of religious beliefs, then, is made in respect of the observations that (i) Christian interpretations of the Bible generally affirm that the triune God is behind all lived experience, yet do not show any aspect of the created order to belong unambiguously to God's true nature; and (ii) the history of Western thought may be characterised by a series of claims about the nature of reality being ultimately located in a particular abstracted aspect (e.g. claims that everything is ultimately numerical / physical / sensory / logical / linguistic / etc) (Dooyeweerd 1979). We might also find that (iii) a claim that fundamental religious convictions lie behind scholarly controversies is compatible with the depth of paradigm conflicts seen in many disciplines, and also theologically compatible with Protestant teaching on the radical nature of sin and the power of idolatry.

Knowledge and statistical inference

How do we use observations to learn about the world? Any scientific or pre-scientific understanding of the world posits something like general laws or kinds, spirits or intentions, natures or essences that cause predictable patterns, yet our flow of consciousness only takes in a single stream of perception, and no matter how many situations and events we observe, we do not observe the causes behind them. Any model of the world is said to be *underdetermined* by any experience, and knowledge results from trying to harmonize ideas with new experience, in a cyclical process of learning. The science of statistics helps with this challenge when quantifiable ideas need to be tested against numerical data that are influenced by extraneous factors that cannot be controlled. So the challenge of statistical inference is to use "noisy" data to develop more accurate knowledge about some quantitative law or pattern. We may wonder how this relates to basic, religious beliefs. I will argue for a cyclic process of refinement of a community's scientific knowledge that is analogous to – indeed a part of – the cyclic process of development of a person's basic, or religious, beliefs (Fig. 1). In modern scientific learning, statistical inference plays an increasingly important role, and I will argue that it depends upon prior scientific beliefs in an essential way that is disputed by some of the paradigms that I will sketch below.

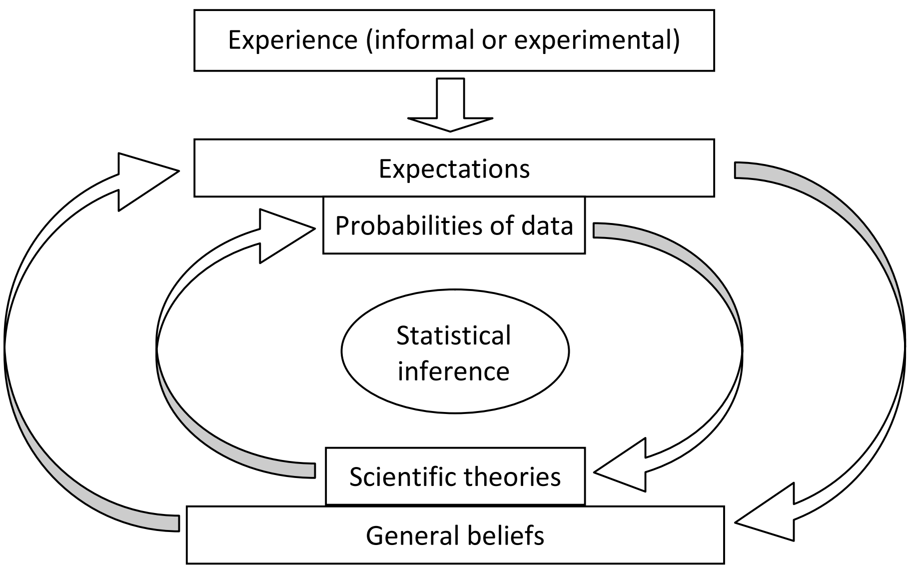

Fig. 1. The cycle of knowledge, showing the proposed analogy between general knowledge (outer arrows) and special scientific knowledge (inner arrows). The latter can be refined with the aid of statistical inference. Experience (top box) provides the contents of knowledge but cannot generate it *de novo*; rather it has a regulating and reinforcing role.

Early work on statistical inference sought to use numerical data to determine the true value of a parameter – say, the density of gold, or the number of tree species in the world. Given that data can reflect a plethora of unknown and varying influences (they are "noisy"), we clearly can't expect a finite set of data to tell us the true value of the number they theoretically tend towards (known as the *expectation*). So we might want to ask what value of that expectation is *most likely* to be true, given the data. But it turns out that the simplest thing is to suppose that the expectation could take any positive value at all and then try and find a value that would make the *observed data most likely*. This is known as the *maximum likelihood* method of estimation, and for many problems, it can be shown that simple statistics, like the mean of a set of repeated measurements, give the maximum likelihood estimate. It's clear that the maximum-likelihood estimate – suppose we commission some estimates based on global tree surveys, and get a mean of 12,054.5 species – is unlikely to be *exactly* right, but in fact the logic of maximum likelihood doesn't say anything at all about probabilities for different values of a parameter; it only gives us a precise number that would make the *data* most likely (Fig. 2). So, rather as a theory is underdetermined by any experience, probabilities of parameter values cannot be given by any data.

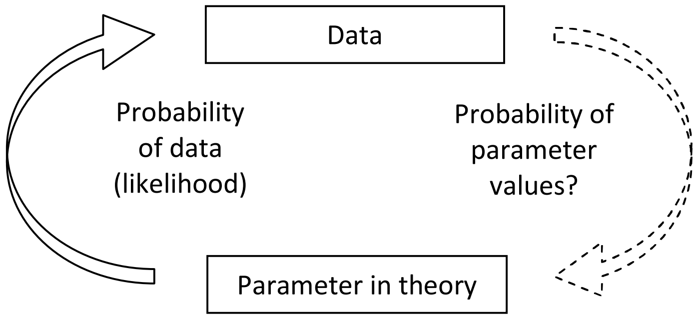

Fig. 2. Given a set of data, inferring the most probable value for the parameter that lies behind them is a much more challenging problem than finding the parameter value that would make the data most likely.

Pierre-Simon Laplace (1749–1827) was a French natural philosopher who worked, among other areas, on celestial mechanics, medical records and jurisprudence. He noted that the probabilities of parameter values cannot be obtained *de novo* from data, but that data can be used to modify *prior probabilities* obtained from another source. Thus Laplace formalised the method of "inverse probability" for inferring what scientists really want to know: the probability of various parameter values, based on data and prior beliefs. He did this by drawing on Bayes' theorem, which concerns the probability of one condition *given* another pertinent condition that is known to hold. Thomas Bayes (1701-1761) was an English mathematician and Presbyterian minister; we'll see later how his theorem is used in Bayesian statistics.

Inverse probability calculations, however, quickly become difficult for many realistic problems, and they also have some conceptual challenges, as we shall see later. Since it is much easier to determine likelihoods of data given arbitrary parameter values, it is perhaps unsurprising that a new school of statistics arose in the early 20th century which discarded the problem of subjective probabilities and laid out rules for making decisions and inferences by considering how frequently or infrequently a specified statistic would arise from data generated by precisely-defined conditions. This "frequentist" school was developed by R.A. Fisher, Jerzy Neyman and others, several of whom worked in evolutionary biology. We will explore below how this paradigm contrasts with the inverse-probability paradigm. But we start with a primer for this approach to statistical analysis.

A primer for frequentist inference

Here's an example of how a frequentist statistical analysis typically proceeds. I'll make comments about some of the assumptions as we go along...

1. A classic frequentist hypothesis starts in the form, "There is some causal relationship between two variables" – say, the number of tree species, *S*, growing naturally in an area and the latitude, *L*, of the area. Notice that this research hypothesis is characteristically vague – but it is a realistic starting point in many areas of exploratory science. The scientist now proceeds by seeking statistical evidence for some kind of causal link between these variables. Once he has such evidence, he has a validated discovery to report to the world! (Or perhaps he imagines that failure to find such evidence will enable him to play a part in Karl Popper's project for scientific progress by conjecture and refutation? Well, maybe... Let's see where we get with this.)

2. The research hypothesis was not quantitative. However, it can be expressed quantitatively if the relationship between *S* and *L* can be measured on or converted to (e.g. by counting positive outcomes) a

quantitative scale. So we find an appropriate *statistic*, *r*, which describes the trend in the relationship between my variables and is mathematically defined so as to be independent of the number of data (*degrees of freedom*) used in calculating it. If we say "*r*=0", we have a precise, quantitative "null hypothesis".

3. At some stage, data are obtained. In our case we cannot experimentally manipulate *L* while measuring *S*, so an alternative is to observe values of *S* and *L* together at appropriately-selected natural instances (probably different spatial locations in the world). Our statistic *r* is then calculated, which is a maximum-likelihood estimator of the true value.

4. In general, the estimate of *r* will not be exactly zero. At this point, if not already, we may suspect that the null hypothesis was a wildly improbable dream, and conclude that the research hypothesis has been supported. But the problem is this: the apparent support might still be consistent with an imaginary world where *r* really is zero but measured values of *r* take non-zero values because of the noisiness of the data. So we proceed to ask how *significant* the departure from zero is, making use of the variability in the data.

5. The statistician therefore refers to the expected probability distribution (Fig. 3) of values of *r* that might be obtained from random samples of new data with the same intrusion of noise, if the true value of were zero (as in the null hypothesis). Even extreme values can be expected to occur with low frequencies, so a small error rate α is specified, which is to be an acceptable risk (conventionally 5%) of rejecting the null hypothesis were it actually true. The statistician then identifies the range of values of *r* that would be most likely to occur under the null hypothesis, so as to account for a proportion 1-α (e.g. 95%) of the possible outcomes. The limits to this range are the *critical values* of *r*.

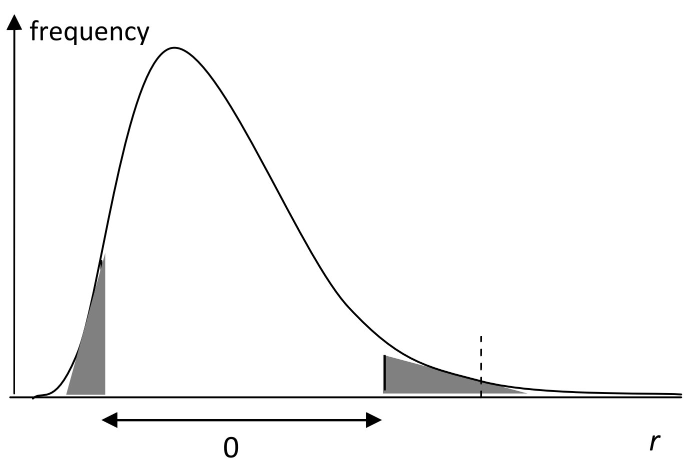

Fig. 3. A frequency graph showing the range of most-likely values that a statistic *r* might take under a null hypothesis that *r* equals 0. The region spanned by the arrows is the 95% confidence interval; values outside this range are expected to occur 5% of the time if the null hypothesis is true. The short dotted vertical line indicates a maximum-likelihood value for *r* estimated from some data, showing that the null hypothesis gives these data a low probability (< 5%). We can go a step further and calculate the so-called *P*-value: the exact probability of obtaining such extreme data under the null hypothesis.

6. By comparing the observed value of *r* with the critical values, it is possible to decide whether the data from which it was derived would be reasonably likely under the null hypothesis. The statistician knows how unlikely it is that values outside the critical range would occur if the null hypothesis were true. So the conclusion of the analysis is a judgment about how unlikely the observed data would be if the null-hypothesis were true – which is expressed as a so-called *P*-value. The point is carefully made to students of statistics that we have a probability of data given a made-up parameter value, not a probability of a parameter value given the data.

7. That kind of conclusion seems logically valid and impressively impartial, but its practical value is not obvious – we want to make *inferences about the parameter* (the actual relationship between latitude and species richness), and here we only have a *conclusion about our data*, based on a pipe-dreamed null hypothesis. However, users of the approach normally make an intuitive leap here: If the data we observed would be very unlikely under a null hypothesis, then surely the null hypothesis is unlikely to be true? Typically the conclusion is subtly re-phrased to say that the data are "inconsistent" with the null hypothesis, as if the conditional probability statement worked both ways.

That primer for frequentism was fairly mechanical. It might even sound as though a computer could be made to do science this way. But there are some problems, which lead Andrew Hartley (Hartley 2007) to distinguish what he calls the "direct frequentist" paradigm from an "indirect frequentist" paradigm for statistical inference. This hinges on the role that is granted for intervention of expert judgement. Consider my analysis of the relationship between latitude and number of tree species. I find that the estimate of r summarising my data is so far from zero that the probability of getting data like these under the null hypothesis is estimated at 0.06. This does not meet the conventional threshold criterion of 0.05 because the 95% of values that would be most likely if $r = 0$ is quite broad. I'm disappointed with my lack of "discovery" – but then I manage to find a large amount of additional data from monotonous forests in the far north, and by increasing my sample size, I arrive at a new P-value of 0.007. The change has occurred not because my new estimate of r is very different, but because the 95% confidence envelope has shrunk. So now I can tell the world that species richness varies with latitude – it decreases as you go further north. Now, for the next stage in my investigation I want to determine whether species richness actually declines with distances both north and south from the Equator, or from some other latitude close by. So I set up a new null hypothesis for L*, the latitude of maximum species richness: L* = 0°. And even with my expanded data set, I cannot quite reject this, because $P = 0.05$. So I will have to use L* = 0 for the calculations in the next stage of my analysis, which will be about how the area of the sampled region affects the number of species. However, I cannot help noticing that my actual maximum-likelihood estimate for L* was 4° S, and the only published paper I can find that estimates the same parameter, although based on all kinds of plants rather than just trees, reports it to be 6° S. Surely I could justify using my exact maximum-likelihood estimate rather than 0, to improve the accuracy of my subsequent analyses? After all, there is rather little land-mass on the Equator itself, and much more just a few degrees south, which may support a greater diversity of habitats and species.

This example shows how expert opinion tends to intervene in a frequentist analysis, in various ways. So one key reason we might not trust a computer to do science under frequentism is that arbitrarily-small effects, so long as they are real, become statistically significant given a large enough sample size. All sorts of quantities in the world may be linked to each other in very weak and indirect ways - like the density of gold and the time of year when it is measured – so one could argue that most null hypotheses are probably false. Paradoxically, therefore, scientific progress can be made by the rejection of null hypotheses which are ultimately trivial. In practice, therefore, the concept of "scientific significance" is important. One reason why the frequentist paradigm works so well may be that there is intuition (tacit knowledge) among practitioners about the minimum size of an effect that would be meaningful in the context of the research hypothesis being considered. This feeds into intuition (or it can be formalised with methods known as *power analysis*) of how to achieve appropriate statistical power (e.g. by controlling sample size) to detect an effect of about this size or larger.

Another key justification for expert intervention that appears in the above example is the inferential weight that may be given to prior studies and theoretical expectations. In fact, given the essential roles for expertise throughout the research process, a sceptic might categorise even the most ardent proponents of conventional, direct frequentism as indirect frequentists. Perhaps there is no chance of computers becoming the scientists of tomorrow - but instead we may have unaccountable geniuses following whims of inspiration that lead them to make authoritative pronouncements about “scientific truth” – with impeccable statistical validity!

A more logical problem is the inferential sleight of hand that seems necessary in order to get any useful conclusion from a frequentist analysis. Conspiracy theorists can have a field day. If there were no gremlins playing bowls in my attic, it seems very unlikely that I would hear the regular bumping noises I’ve just noticed! Or: if we were in Wales, it’s unlikely that we would be having two dry days in a row...

Is there any alternative? Well, let’s go back to the good old school of inverse probability...

A primer for Bayesian inference

Here is how a Bayesian analysis typically proceeds. Again, I will point out some assumptions and contrasts with the frequentist approach as we go...

1. We start in a similar way to the frequentist primer: a situation about which the researcher would like to know more can be described by some measurable parameter. For example, species richness and latitude are thought to be related to each other, and the parameter *r* can describe a relationship between these variables *S* and *L*. Note, however, that we do not start with the black-and-white question of whether or not such a relationship exists; our world picture and previous experience lead us to believe that this relationship can be quantified, and our aim is to gain an improved estimate of *r*.

2. Thus we have some prior belief, however vague, about what kind of values the parameter *r* might take. We don’t need a single precise value; we can use a probability distribution (Fig. 4). This is the *prior* distribution of *r*, describing our uncertainty about its real value. In contrast to the frequentist approach, we talk about *subjective probabilities* for different values of *r*. *r* itself has a single, real value of course, but we don’t know it.

3. Data are collected in a Bayesian research programme, much as in a frequentist programme. The more, the better!

4. The data are analysed by combining them with the prior for *r* in such a way as to give a new, *posterior* distribution for *r*. Bayes’ Theorem states that the probability of condition A *given* condition B is the same as the probability of condition B *given* condition A multiplied by the ratio between the prior probabilities of the two conditions. Taking one condition to be a particular value of the parameter *r* (with prior probabilities given by the prior belief) and the other to be the data observed (which may be said to have a probability of one), we can calculate an improved, posterior probability for each parameter value on the basis of the data observed (Fig. 4). (For a given parameter value, the posterior probability is simply the prior probability multiplied by the likelihood of getting our observed data from that value.)

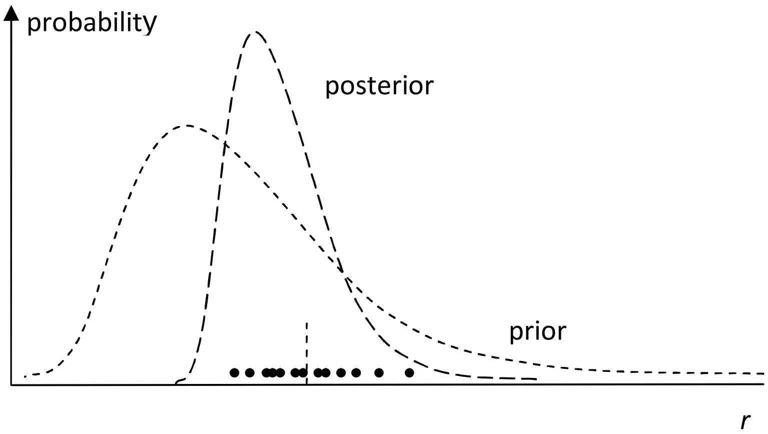

Fig. 2. A probability graph showing the prior and posterior probability distributions for a parameter *r*, with the data points (black dots) that were used to convert the prior to the posterior. The short dotted vertical line indicates the maximum-likelihood parameter value estimated from the data.

Thus the conclusion from a Bayesian analysis is an updated description of the probabilities of the parameter *r* taking a range of values. We are able to make real inferences about parameter values and hypotheses by using a method of "inverse probability". If some parameter values were deemed impossible by the prior assessment, they will still have zero probability when multiplied by the likelihood of getting the observed data from those values. This removes the gremlins problem. If some parameter values were deemed highly likely by the prior assessment, their new value will essentially be constrained by the likelihood of getting the observed data from those values - which removes the Welsh sunshine problem.

At this point, Hartley distinguishes two forms of Bayesianism. One makes full use of the Bayesian toolkit, and starts by trying to describe the current state of belief in proper, informative priors. To go back to the example of investigating the latitude and which tree species richness peaks, I could take previous work on other kinds of species to help shape my prior. I could even take the amount of land mass at each latitude to determine my prior, if I have good reasons to believe this to be an important determinant. In general, the challenge of forming a prior opens up many options: does the investigator try to describe his current personal beliefs, or the consensus of a community of which he is part, perhaps using some accepted theory? Or does he consider the beliefs of a (perhaps sceptical) audience that he intends to convince? This seems to be the approach taken qualitatively in philosophical discourse, where the proponent of an argument considers a range of possible objections in order to respond to them. We will return to these challenges below. Hartley calls this full Bayesian approach "subjective Bayesianism".

The other option is *objective Bayesianism*, which was the method developed by Laplace with the principle of "insufficient reason". This approach uses trivial, or "uninformative" priors (we might also call them "lazy"): simply identify the full range of possible parameter values, or set far-out limits if it would be infinite, and start by declaring all values in this range to be equally likely. This conveniently avoids the difficult task of deciding on realistic levels of prior subjective belief about the parameter values, which makes it attractive, especially nowadays to researchers who are used to the (even lazier) methods of frequentist inference. But the posterior probabilities it gives may be misleading: if we have a spurious data set, the estimate will be just as spurious, because the prior distribution was fake!

The impact of religious beliefs in statistics

The four paradigms of statistical inference I have outlined (as identified by Hartley) differ radically in many respects: what inputs they require, what role prior beliefs may have, how they consider unknown parameters, what information they yield, how they lead to inference and decision-making, etc. It is clear that there may be deep-seated logical and practical challenges to using each of these paradigms. These considerations may suggest ethical reasons for preferring one or other approach. Indeed, I've shown how the different approaches can give different conclusions in similar situations; specifically, direct frequentism can be shown to give unreasonable conclusions in cases where prior information is clearly important.

Besides all this, however, there are more fundamental questions to ask. The divergence of the perspectives, together with the depth and persistence of controversy among their advocates, may suggest that there is a religious dimension to their differences. This would widely be denied under popular notions of "religion". Let me therefore advance this thesis by laying out some of the fundamental characteristics of the paradigms and noting the sense in which these are "religious" under the definition of "religious beliefs" that I proposed at the start.

- Direct frequentism seeks to induce beliefs directly from mathematical quantities. Inference from data to parameters, under this paradigm, is only achieved by making an intuitive but non-logical re-interpretation of the formal conclusion of an analysis. Thus it appears to be reductionistic insofar as it seeks to reduce beliefs to mathematical quantities (assuming that at least some people experience their beliefs as not inherently mathematical in nature – which I do, for one).
- Indirect frequentism assumes an essential role for expert intervention in processes of scientific induction without any need to formalise or constrain this intervention. The mathematical side of what starts out as a numerical analysis is ultimately subsumed under the judgement of the expert. Thus indirect frequentism (and arguably also direct frequentism at the point of inference from data to parameters) attributes a determinative role for special human judgement, without formalising this. It may be seen as a humanistic construal of reality in defiance of experience.
- Subjective Bayesianism acknowledges radical differences in people's subjective beliefs, which are an indispensable grounding point. Analysis proceeds according to clear mathematical principles that appear to provide a legitimate way to combine numerical and fiducial aspects of human experience.
- Objective Bayesianism conflates mathematical and (fiduciary) properties at the start of an analysis, which may be reductionistic as in the case of direct frequentism. Thus objective Bayesianism substitutes mathematical equality for prior belief, arguably making a category error at the start of an analysis.

It may clarify some of these points if we return to the analogy with the development of general knowledge outlined earlier on (Fig. 1). If someone I do not know tells me something that surprises me, I consider it less likely to be true than if someone I trust tells me something that broadly fits my prior beliefs. Thus, arguably, our everyday assessment of information is informally "Bayesian": the nature of our conclusions from new information are heavily dependent on prior beliefs and the perceived reliability of the source of the new information. Ironically, indeed, research in experimental psychology (a discipline which has conventionally relied on direct frequentist inference from data) repeatedly argues for the importance of prior assumptions in human perception [ref needed]. Even within the sciences that champion frequentist approaches, then, we might argue that all inference (statistical and informal) is not just "indirectly frequentist" but in fact "subjectively Bayesian", though it be performed under a smokescreen of something supposed to be more objective.

I finish with a personal conclusion. Together, the above considerations mean that I take very seriously the status of statistical science as a tool for refining and changing beliefs – both my own and other people's. Since I am not committed to the autonomy of mathematical reasoning for the generation of true knowledge, I do not wish to use frequentist methods as a way of denying a role for genuine prior belief, and I am concerned about the need for all stages of the inferential process to be properly acknowledged and described – as well as the risk of drawing inaccurate conclusions by means of illogical inferential protocols or inaccurate priors. However, there are practical challenges to adopting subjective Bayesian methods. In presenting conclusions from a scientific study – such as my work on the latitudinal gradient in tree species diversity – I seek to provide an analysis of the data that can be interpreted by a wide range of audiences, whose prior beliefs I do not accurately know. The safest policy would be to present no inferences other than my own, reducing the report's conclusion to a personal testimony of modifications to my own beliefs on the matter. But readers expect to be offered a stronger conclusion, one which challenges them to make a certain inference too. So, for the time being, I am likely to follow conventional analytical protocols, knowing that readers will ultimately draw their own inferences in any case.

This conventional "objectivist" approach perhaps runs into problems most seriously when a wider audience is engaged. For example, scientific studies that are picked up by media channels often involve their authors in anxiety about the interpretation of their conclusions by non-specialist audiences. I am also concerned about teaching: that statistical inference is often taught in ways that encourage reductionistic beliefs and the idolatries that ensue from them. This was part of my burden in revising an advanced statistics course that I was appointed to teach to postgraduates at the University of Leeds recently. Thus I can say that the development of my own Christian beliefs is starting to influence the way I approach the teaching and application of statistical methods in the university and beyond. I hope other scientists too will increasingly consider the presuppositions and faith commitments that shape their inferential practices, and be prepared to examine them critically and openly, whether or not they would agree with the analysis given here and by Andrew Hartley.

References

Clouser, R. (2005). The Myth of Religious Neutrality.

Dooyeweerd, H. (1953). A New Critique of Theoretical Thought: The necessary presuppositions of philosophy, H.J. Paris.

Dooyeweerd, H. (1979). Roots of Western Culture: Pagan, secular and Christian options. Toronto, Wedge.

Hartley, A. M. (2007). Christian and Humanist Foundations for Statistical Inference: Religious Control of Statistical Paradigms, Resource Publications.

Strauss, D. F. M. (2011). Philosophy: Discipline of the Disciplines.