Statistical Science
1992, Vol. 7, No. 3, 358–368

R. A. Fisher's Fiducial Argument and Bayes' Theorem

Teddy Seidenfeld

1. INTRODUCTION

In celebration of the 100th anniversary of Fisher's birth, I want to raise the subject of fiducial inference for our reflection. Shortly after Fisher's death in 1962, my teacher and friend, Henry Kyburg, addressed a conference on fiducial probability. I find it appropriate to begin with some of Kyburg's (1963) remarks:

> I am a logician and a philosopher; I have not studied statistics for very long, and so I still very quickly get out of my depth in a discussion of the technicalities of statistical inference. But I think it is important, none the less, for people whose interests lie in the area of inference as such to do the best they can in reacting to – and in having an action upon – current work in that particular kind of inference called "statistical." That this interaction is difficult for both parties is the more reason for attempting it. (p. 938)

My purpose in this essay is to try to assist that interaction by focusing on the rather vague inference pattern known as the "fiducial argument." I hope to show that, though it really is untenable as a logical argument, nonetheless, it illuminates several key foundational issues for understanding Fisher's disputes over the status of Bayes' theorem and thereby some of the continuing debates on the differences between so-called orthodox and Bayesian statistics.

Begin with the frank question: What is fiducial probability? The difficulty in answering simply is that there are too many responses to choose from. As is well known, Fisher's style was to offer heuristic examples of fiducial arguments and, too quickly, to propose (different) formal rules to generalize the examples. By contrast, among those who have attempted to reconstruct fiducial inference according to a well-expressed theory, one must single out the following five, in chronological order: Jeffreys (1961), Fraser (1961), Dempster (1963), Kyburg (1963, 1974) and Hacking (1965). So, instead of beginning with a formal definition, let us try to find out what fiducial inference is supposed to accomplish and then see whether, in fact, any answer to our question can be accepted. That is, first we shall determine whether the goals Fisher set for fiducial probability are even mutually consistent.

A convenient starting place for our investigation is the distinction between *direct* and *inverse* inference (or *direct* and *inverse* probability). These terms date back at least as far as Venn and appear often in those passages where Fisher attempts to explicate fiducial probability, for the reason I will give below. As a first approximation, the difference between direct and inverse probability is the difference between conditional probability for a specific (observable) event D given a statistical hypothesis S, $p(D|S)$, and conditional probability of a statistical hypothesis given the evidence of sample data, $p(S|D)$. For example, an instance of a direct probability statement is, "The probability is 0.5 of 'heads' on the next flip of this coin, given that it is a *fair* coin." Here the conditioning proposition may be understood as supplying the statistical hypothesis S that there is a binomial model (a *hypothetical population*) for flips with this coin, $\theta = 0.5$, and the event D is that the next flip lands "heads." An inverse probability statement is, "The probability is 0.4 that the coin is *fair*, given that 4 of 7 flips land 'heads'." Direct and inverse *inference*, then, denote those principles of inductive logic which determine or, at least, constrain the probability values of the direct and inverse probability statements. They would explain the probability values 0.5 and 0.4, assuming those values are inductively valid.

A maxim for direct inference which is commonplace, in so far as any inductive rule can be so described, is what Hacking (1965, p. 165) labels the *Frequency Principle*. It states loosely that, regarding the direct probability $p(D|S)$, provided *all* one knows (of relevance) about the event D is that it is an instance of the statistical "law" S, the direct conditional probability for event D is the value specified in the statistical law S. If all we know about the next flip of this coin is that it is a flip of a *fair* coin, then the direct probability is 0.5 that it lands "heads" on the next flip.

To this extent, direct probability is less problematic than inverse probability. There is no counterpart to the Frequency Principle for inverse probability. At best, $p(S|D)$ may be determined by appeal to Bayes' theorem, $p(S|D) \propto p(D|S)p(S)$: an inverse inference which involves both a direct probability $p(D|S)$ and a *prior* probability $p(S)$. It is here that Fisher sought relief from inverse inference through fiducial probabil-

Teddy Seidenfeld is Professor in the Departments of Philosophy and Statistics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.