



Economics Division
School of Social Sciences
University of Southampton
Southampton SO17 1BJ
UK

Discussion Papers in Economics and Econometrics

INFORMATION AND ECONOMICS IN FISHER'S DESIGN OF EXPERIMENTS

John Aldrich

No. 0502

This paper is available on our website

http://www.socsci.soton.ac.uk/economics/Research/Discussion_Papers

Information and Economics in Fisher's Design of Experiments*

John Aldrich

Economics Division

School of Social Sciences

University of Southampton

Southampton

SO17 1BJ

UK

e-mail: john.aldrich@soton.ac.uk

Abstract: In 1951 R. A. Fisher described what had been achieved in the 20th century so far: “we have learnt (i) To conserve in its statistical reduction the scientific information latent in any body of observations. (ii) To conduct experimental and observational inquiries so as to maximise the information obtained for a given expenditure.” This paper asks what Fisher meant and, in particular, how he saw his work on experimental design as contributing to the objective of maximising information for a given expenditure. The material examined ranges from detailed work on issues like “the information lost in measurement of error” to polemics against decision theory.

* This paper is a revised (March 2005) version of one presented at the Open University Conference *The History of Statistics 1900 ± 30* in December 2004.

Introduction

The themes usually associated with Ronald Fisher's work on experimental design are replication, randomisation and local control, *not* information and economics; see e.g. Atkinson & Bailey (2001, p. 57). Yet consider Fisher's (1951, p. 54) summary of what statistics had achieved in the first half of the twentieth century:

we have learnt (i) To conserve in its statistical reduction the scientific information latent in any body of observations. (ii) To conduct experimental and observational inquiries so as to maximise the information obtained for a given expenditure.

In this formula, spanning what statisticians do with observations and how they choose them, Fisher is clearly referring to his own work—in (i) to the theory of estimation (1922, -25a, -34), and in (ii) to *The Design of Experiments* (1935) and publications going back to 1925.

In another retrospective piece Fisher (1947, p. 435) described the “development of the theory of experimental design” in terms of the relationship between (i) and (ii):

[The finding that the amount of information extracted in the process of estimation could never exceed the quantity supplied by the data], combined with the practical fact that directly available processes of computation would extract almost always a very large fraction of the total available shifted the *moral balance*. ... The weight of his responsibility was thrown back on to the process by which the data had come into existence. So armed with amount of information as a practical tool, statisticians came to study what forms of experiment or what types of observational programs would yield the most information for a given expenditure in time, money and labor.

Was there anything behind this story and the spanning formula, or were they the froth

celebrations call for? The standard developments, expositions, appreciations and critiques of Fisher's design work do *not* treat replication, randomisation and local control as expressions of fundamental concerns with information and economics because, presumably, the authors believed those ideas played no real part in the work or were better dropped; see e.g. Cochran & Cox (1950), Youden (1951), Finney (1955), Yates (1964), Pearce (1979), Preece (1990) or Street (1990). I will argue that Fisher made a genuine effort to found his design work on those ideas but, because he failed, the formula and story sound hollow.

Fisher first published on information in 1922 and, according to Box (1978, p. 156), designed his first experiment in 1924. His early estimation work is described in §1 and his early design work in §§2-5. How he brought information and design together is discussed in §§6-8. §9 picks up the economic theme and compares Fisher's work with Wald's decision theory, which he knew well enough to attack, and with comparison-of-experiments theory, which he did not know at all, and which appears not to have known him. The concluding remarks in §10 include some speculations about why these themes have been so invisible.

1 “The total available relevant information”

The theme of information and its conservation distinguished Fisher's estimation from earlier work in the theory of errors (least squares) and on biometry (the method of moments). Today “Fisher information” is either a heuristic device or a technicality in the asymptotic theory of maximum likelihood, but the intuitions and the asymptotic theory are survivals from a larger information scheme. We need to notice the scheme to make sense of the “conservation” slogan and to indicate the ideas Fisher applied to design. Only the basic ideas are needed for Fisher did not apply the more sophisticated concepts, such as ancillarity and second-order

efficiency, to design. For more on the origins of the theory, see Aldrich (1997) and, for more on the later, more sophisticated theory, see Hinkley (1980 and -80a). Edwards (2000) discusses Fisher’s thoughts on information and entropy in relation to the fundamental theorem of natural selection. For general information on Fisher and further references, see Aldrich (2003).

The theme of “conserving scientific information” and, indeed, the notion of an amount of information first appear in the “Mathematical Foundations of Theoretical Statistics” (1922). However, the technique of information was more thoroughly developed in the “Theory of Statistical Estimation” (1925a). There (p. 709) the ‘intrinsic accuracy’ of the frequency curve f “as a means of estimating θ ” is defined as the mean of

$$-\frac{\partial^2}{\partial\theta^2} \log f(x; \theta),$$

with the gloss that this “may equally be conceived as the amount of information in a single observation belonging to such a distribution”.

The “efficiency” with which information is “isolated” (or “extracted” or “conserved”) was also formalised. The first formulation (1922, p. 316; 1925, p. 13) was restrictive but the second (1925a, p. 714) had “the advantage of applying to finite samples and to other cases where the distribution is not normal”:

The efficiency of a statistic is the ratio of the intrinsic accuracy of its random sampling distribution to the amount of information in the data from which it has been derived.

The illustrations include the efficiency of the median for estimating the location of a Cauchy for sample sizes between 1 and 11. This information theory was not an asymptotic theory.

Fisher would say much more about conserving information but the chief practical lesson

was already there in 1922: don't use the method of moments (exposed as inefficient in §§8-11), use maximum likelihood! Maximum likelihood underwrote the "practical fact that directly available processes of computation would extract almost always a very large fraction of the total available."

The 1922 paper did not only supply tools for the theory of estimation, it cast the problem of Statistics as the problem of estimation. The formulation passed unchanged into the *Statistical Methods for Research Workers* (1925, pp. 1 & 7): Statistics is "the study of methods of the reduction of data." This is *not* today's "descriptive statistics" for the statement, "It is the object of the statistical processes employed in the reduction of data to exclude ... irrelevant information, and to isolate the whole of the relevant information contained in the data", is to be understood in terms of sufficient statistics and parameters of "hypothetical infinite populations".

This definition reappears in mid-century as lesson (i), redrawn to accommodate the design of experiments. The redrawing is less surprising than the original for the goal of isolating the whole of the relevant information contained in the data never came near to covering the full range of Fisher's non-experimental statistics. There is an alarming mismatch between the information objective of the introductory chapter of the *Methods* and the tests of significance (χ^2 , t and z) that fill the rest of the book. In that chapter (pp. 8-14) and in the "Mathematical Foundations" (p. 304) the business of testing is to check the validity of the specification (the "mathematical form of the population"). Of course, Fisher's ideas on testing and on information were broadly compatible in that the test statistics were usually based on quantities with good information credentials but he never found it necessary to make really tight connections. We return to these issues in §§7-8 below.

There was more of a journey for (ii) and, by one test, it never arrived! Fisher (1925,

p. 9) stated that, when the consequences of a hypothesis agree with observations, it is “justified at least until fresh and more stringent observations are available.” The treatment of experiments later in the book made that passivity obsolete on publication. In later editions an advertisement was added, “The author has attempted a fuller examination of the logic of planned experiments in his book, *The Design of Experiments*” (Fisher (1936, p. 9)). Point (ii) materialised in the 1930 edition of the *Methods* (see §6 below) in the last chapter. It never progressed to the first and statistics remained “the study of methods of the reduction of data.”

To summarise, by 1922 Fisher had a vision, or at least a slogan, and the basis for believing that “directly available processes of computation would extract almost always a very large fraction of the total available”. How statisticians interested in design came to be “armed with amount of information as a practical tool” is a more involved story.

2 “Decreasing the standard deviation and increasing our certainty”

From his first published work on design Fisher emphasised “the two *desiderata* of the *reduction of error* and of the *valid estimation* of error”, as he put it in 1926 (p. 508). The “valid” estimation of error—and randomisation as a way of guaranteeing it—was a new concern but the reduction of error was not; it was generally seen as the main goal in designing experiments.

From joining Rothamsted in 1919 Fisher was in continuous contact with W. S. Gosset (‘Student’) and, from 1923, the design of experiments was a topic in their correspondence (McMullen (1970)); see Pearson’s ‘*Student*’ (1990) for an account of ‘Student’'s work and his interaction with Fisher. Although Gosset had only one pre-war publication on design, he had been giving advice on the statistical aspects of experiments from 1904 and, insofar as there

was a statistical authority on the design of experiments, he was *it*. The publication was the appendix he wrote for Mercer and Hall's "The Experimental Error of Field Trials" (1911). This paper provided important information on the variability of yields in uniformity trials.

The appendix gave "a method of arranging plots so as to utilize a given area of land to the best advantage in testing two varieties." 'Student' used data from the paper to illustrate his reasoning as others, including Fisher (1925)—see §5—and Engledow (1925), did later. 'Student' imagined a field divided into plots sown "chequer fashion" with the two kinds of seed. The objective is to "decrease the standard deviation and to increase our certainty." On the size of plot and the advantage of comparing adjacent plots, he concludes

Roughly speaking, one-twentieth acre plots of mangolds would require at least twice as much land as one-two-hundredth acre plots in order that we may place as much confidence in the results....

Hence it clearly of advantage to use the smallest practicable size of plot.

Also the advantage of comparing adjacent plots is apparent in these example since with roots less than two-thirds of the land is required to give the same accuracy as random comparison...

"Confidence" and "accuracy" are both determined by the size of the standard deviation of the difference between adjacent plots. In 1923 'Student' entertained the possibility of using calculus to determine the size of plot that will give the "minimum probable error per acre" but considered the data inadequate (p. 279).

Fisher kept the goal of decreasing the standard deviation, re-phrased as increasing the information, but that rephrasing took rather a long time.

3 “Modern methods” and the analysis of variance

Fisher’s ideas on design appear fully-formed in *Statistical Methods for Research Workers* (1925). The “principles underlying modern methods of arranging field experiments” are put in the “further applications” of the analysis of variance for Fisher (p. 224) maintained that, “The statistical procedure of the analysis of variance is essential to the understanding of [those] principles.” Because the analysis of variance provided the framework for all of Fisher’s work on agricultural experiments, it requires examination. “Modern methods” was an outrageous, but characteristic, tag for methods that hitherto had existed only in Fisher’s head!

In 1923 Fisher introduced the analysis of variance in a paper with his assistant Winifred MacKenzie on “the manurial response of different potato varieties” (1923). This was an extraordinary inauguration for, after giving the most basic information on the linear model, Fisher abandoned it for a multiplicative model; see Cochran (1980) for a detailed analysis of the paper. Alas, none of Fisher’s later publications contains a clear statement of the assumptions underlying the analysis of variance. The paper on the z -distribution (1924) does not treat the distribution as it arises *in the experimental situation*. However, Fisher explained the assumptions in an exposition he wrote for Gosset and which appears as a footnote in ‘Student’ (1923).

In the passage in ‘Student’ (1923, p. 283n) Fisher begins by saying:

The yield obtained in any experiment is the sum of three quantities, one depending only on the variety; a second depending only on the ‘trial’; and a third, which may be regarded as the ‘experimental error’ varying independently of variety and trial in a normal distribution with a standard deviation which it is desired to estimate.

The model explaining the yields X_{pq} is thus

$$X_{pq} = A_p + B_q + \varepsilon_{pq}, \quad \varepsilon_{pq} \sim IN(0, \sigma^2), \quad p = 1, \dots, m; \quad q = 1, \dots, n$$

(Apart from the error symbolism—and the term “model”—the notation is Fisher’s.)

Fisher estimates the quantities A_p, B_q by least squares, minimising

$$\sum (X_{pq} - A_p - B_q)^2 \quad (1)$$

“Evidently (1) will be a minimum if

$$A_p + B_q = \bar{X}_p + \bar{X}_q - \bar{\bar{X}}$$

where \bar{X}_p is the mean of the values obtained with variety p , \bar{X}_q the mean of the values obtained with trial q , and $\bar{\bar{X}}$ is the general mean.” He next introduces the analysis of variance table:

The actual evaluation is most conveniently carried out in the following form of the analysis of variance:

| Variance | Degrees of freedom | Sum of squares |
|----------------------|--------------------|--|
| (a) Due to variety | $m - 1$ | $n \sum_1^m (\bar{X}_p - \bar{\bar{X}})^2$ |
| (b) Due to trial | $n - 1$ | $m \sum_1^n (\bar{X}_q - \bar{\bar{X}})^2$ |
| (c) Random variation | $(m - 1)(n - 1)$ | $\sum_1^m \sum_1^n (X_{pq} - \bar{X}_p - \bar{X}_q + \bar{\bar{X}})^2$ |
| (d) Total | $mn - 1$ | $\sum_1^m \sum_1^n (X - \bar{\bar{X}})^2$ |

The sum of squares in line (c) being calculated by subtracting the values of lines (a) and (b) from the total. If either variety or ‘trial’ were without significant effect

on the yield, the corresponding mean square would not differ significantly from that of line (c). To test the significance of such a difference we may use the fact that the estimates of variance in (a), (b) and (c) are all independent, and when m and n are fairly large the natural logarithms of the mean square has standard deviation $\sqrt{2/n_1}$ where n_1 is the number of degrees of freedom.

Here is the familiar table with the variance and degrees of freedom identities and also the idea of testing whether distinct estimates of a variance are significantly different. The normal approximations were soon replaced by the exact z ($= \frac{1}{2} \ln F$) distribution. Fisher introduced z in the “On a distribution yielding the error functions of several well known statistics” (1924) and used it in when he re-examined some of the Fisher & MacKenzie data in Ex. 41 of the *Methods* (1925, p. 203).

When Fisher took this framework from *analysis* and applied it to *design* he introduced randomisation at the same time. The next section takes a quick look at the status of normal theory when there is randomization. This may seem a detour from our quest for information in design, especially when Fisher (1947, p. 435) wrote,

The only modification of the art of experimentation, as developed on statistical principles, which does not flow directly from the requirement to maximise the amount of information, obtainable from limited resources, consists in the introduction into the experimental procedure of what is called *randomization*.

However, the status of normal theory matters for the information project because Fisher knew how to do normal theory information calculations but he never even conceptualised information in nonparametric situations.

4 Randomization and the theory of errors

Fisher went on using the normal theory methods of 1923—with the z -refinement—after he adopted randomization. Numerous authors, including Welch (1937) and Kempthorne (1952), have investigated the relationship between randomization and normal theory and naturally commentators on Fisher have asked what was in *his* mind when, e.g., he wrote

The estimate of error is valid, because, if we imagine a large number of different results obtained by different random arrangements, the ratio of the real to the estimated error, calculated afresh for each of these arrangements, will be actually distributed in the theoretical distribution by which the significance of the result is to be tested. (1926, p. 507)

Kempthorne (1966) and Box (1980) argue that Fisher considered the normal theory analysis an approximation to an exact test based on the randomisation distribution, while Hinkley (1980b, p. 583) appeals to several authorities, including Yates, to support the view, “the purpose of randomization in the design of agricultural field experiments was to help ensure the validity of normal-theory analysis”. I think the latter better represents Fisher’s position—whether, or not, it is a defensible position!

The statement at the beginning of chapter X of the *Experiments* seems clear:

The foregoing chapters, III to IX, have been devoted to cases to which the theory of errors is appropriate ... In this kind of hypothesis all discrepancies classified as error, and not eliminated by equalisation or regression, are due to variation in the material examined, following the normal law of errors with a definite and constant, but unknown variance. (1935, p. 187)

In support of their view Kempthorne and Box quote a well-known passage from earlier in the book, chapter III (p. 51)

the physical act of randomisation, which ... is necessary for the validity of any test of significance, affords the means, in respect of any particular body of data, of examining the wider hypothesis in which no normality of distribution is implied.

The “wider hypothesis” was not something that required investigation because of randomisation. It considered the possibility that the “variation in the material examined”—the ε_{pq} of the last section—may be non-normal. While Fisher liked the ingenuity of his randomisation test, he insisted that non-normality was not a genuine issue. He (1926, p. 505) held that the “assumptions about the nature of field errors” used in the theory of errors have been “extensively verified in the examination of the results of uniformity trials.” There are similar statements in (1929) and in (1935, p. 50) when he presented the randomisation test.

5 “Accuracy” “precision” “information”

The exposition of the “technique of plot experimentation” in the *Methods* (1925, pp. 224-229) demonstrates two devices: randomisation to produce a valid estimate of error and the use of randomised blocks to reduce error. Fisher rejected any “pre-arranged system” of choosing plots, going against what ‘Student’ had recommended in 1911 and -23.

Fisher makes his points, not by a theoretical analysis, but by, what was, in effect, a single replication Monte-Carlo experiment inspired by the real conditions of Mercer and Hall’s (1911) uniformity trials. It is worth setting this out in detail for it was the canonical example, the means by which Fisher expressed his meaning in 1925 and re-expressed it in 1938.

Fisher takes 20 strips of land and assigns 5 different treatments each in quadruplicate to

them. The allocations are “found by shuffling 20 cards thoroughly and setting them out in order.” The letters are just labels attached to plots for there is no difference in treatment. A table gives the weight of mangold roots harvested from each plot and the corresponding one-way analysis of variance table is

| Variance due to | D. F | Sum of Squares | Mean Square | St. Dev. |
|--------------------|------|----------------|-------------|----------|
| Treatment | 4 | 58,725 | 14,681 | 121.1 |
| Experimental error | 15 | 231,041 | 15,403 | 124.1 |
| | 19 | 289,766 | 15,251 | 123.5 |

Fisher comments, “It will be seen that the standard error of a single plot estimated from such an arrangement is 124.1, whereas, in this case, we know its true value to be 123.5.” The point is that the standard error of a single plot is close to the true standard deviation of this finite population of 20 numbers.

He next shows how, by blocking, the “accuracy of our comparisons is much improved; in fact the remaining variance is reduced to almost 55 per cent of its previous value”:

we may divide the 20 strips into 5 blocks, and impose the condition that each treatment shall occur once in each block; we shall then be able to separate the variance into three parts representing (i.) local differences between blocks, (ii.) differences due to treatment, (iii.) experimental errors; and if the five treatments are arranged at random within each block, our estimate of experimental error will be an unbiased estimate of the actual errors in the differences due to treatment. As an example of a random arrangement subject to the above restriction, the following was obtained:

A E C D B | C B E D A | A D E B C | C E B A D.

Analysing out, with the same data as before, the contributions of local differences between blocks, and of treatment, we find

| Variance due to | D. F. | Sum of Squares | Mean Square | St. Dev. |
|--------------------|-------|----------------|-------------|----------|
| Local differences | 3 | 154,483 | 51,494 | .. |
| Treatment | 4 | 40,859 | 10,215 | .. |
| Experimental error | 12 | 94,424 | 7,869 | 88.7 |
| Treatment + error | 16 | 135,283 | 8,455 | 92.0 |

The value of “almost 55 per cent” for the reduction is obtained from the 8,455 of this table and the 15,251 of the previous one. The criterion of “improvement” is the same as Gosset’s (§2 above).

When Fisher gave a more extended account of the principles of design in his 1926 “The Arrangement of Field Experiments” he did not use the analysis of variance as scaffolding. As in the *Methods*, he compared a fully randomised scheme and a randomised block scheme, linking the two with the statement “On most land, however, we shall obtain a smaller standard error, and consequently a more valuable experiment, if we proceed otherwise.” (p. 509)

From Fisher’s words quoted in the Introduction—“armed with amount of information as a practical tool, statisticians came to study ...”—it may be surprising that information made only slow progress into Fisher’s publications and that its installation was completed only in the *Experiments* of 1935. In the meantime, e.g. in Eden and Fisher (1927), “accuracy” and “precision” were used as glosses on the standard deviation. Yet for some years before 1935 information seems to have been in waiting. When Fisher & Eden (1929, p. 205) introduced a “precision index”, they denoted it by I ; it was calculated according to the formula

$$I^{-1} = 100 \left(\frac{\sigma}{m} \right)^2$$

where “ σ is the standard deviation of the mean yield of each treatment, and m is the mean yield of all.” When the index re-appeared in the *Experiments* it was in §60, *Precision regarded as amount of information* (1935, p. 180) and I was identified as the “quantity of information.” As such, it went into the *Methods* (7th edition, 1938).

Fisher (1938, §48, p. 275) applies the 1929/35 information formula in a reconsideration of 1925 example from the viewpoint of the “amounts of information elicited by different methods”,

If we take as having unit value an experiment giving comparable yields subject to a standard error of 10 per cent., the value of such an experiment as this, in quadruplicate, may be found by squaring one-tenth of the mean yield [$3285.75 = m$], multiplying by four (giving 431816) and dividing by the mean square obtained by each method of procedure. For randomisation without blocks we have then

$$\frac{431816}{15250.8} = 28.18$$

units of information. Using randomised blocks we have 51.07 ...

| | Degrees of freedom for error | units of information |
|---------------------------|------------------------------|----------------------|
| Randomisation of 20 plots | 15 | 28.18 |
| Randomisation in 4 blocks | 12 | 51.07 |

In 1938 there was also an adjustment for position within the blocks based on the analysis of covariance but, as it does not involve any new principles, it is omitted from the table. There was also a refinement in evaluating these arrangements, which *did* involve a new principle. The refinement, introduced in the *Experiments* (1935), is considered in §8. But first we return to 1930 when the information objective in design was *first* made patent.

6 Information: “design and precision”

Besides agricultural experiments, based on the theory of errors, the *Experiments* considered two other kinds of experiment. The Lady tasting tea (ch. II) was a nonparametric one-off but the other discussed in ch. XI represented a substantial body of work integrated into the theory of estimation. The examples (1935, pp. 218-44) involve discrete distributions and the design interest ranges from the choice of sample size, when estimating the probability of success in Bernoulli trials, to the choice of schemes for estimating linkage in man. There was no randomization, as though we might reasonably “act in the faith that Nature has done the randomization for us” to use Fisher’s words from 1947 (p. 436). This information-based design theory first appeared in the estimation chapter of the *Methods*. In 1930 there was a new sub-section extending the running example of estimating linkage in plants to the choice of which crosses to make.

Originally (see §1 above) Fisher used information for comparing the sampling distributions of estimators based on the same data and the idea of 1930 was that it could be used for comparing distributions arising from different experiments, provided the same parameter were involved. However there was already something like a comparison of experiments in the “Mathematical Foundations.” Its Fig. 1 (1922, p. 322) displays a Pearson Type VII curve—the Cauchy distribution—and a normal curve of equal intrinsic accuracy. The point of the pairing is that for any (troublesome) Cauchy there is a (tractable) normal that is as difficult to locate; here intrinsic accuracy is interpreted as a measure of the difficulty a particular distribution poses for estimation.

The new subsection of the *Methods*, §57.2 “The Amount of Information: Design and Precision”, compares the information to be gained from the first generation of plants, represented by (A), and that from breeding the second generation from different kinds of plant, repre-

sented by (B) or (C). The interest is in a recombination fraction p and Fisher is comparing three ways getting information about p . The three experiments amount to sets of Bernoulli trials, where the probability of a success is related to the parameter p in three different ways. As usual, the argument unfolds through a numerical example based on actual data. Having estimated the recombination parameter, Fisher gives values for the additional information per plant under the three schemes: “At 6.345 per cent recombination [the estimated value] the numerical contributions under (A), (B) and (C) are .006051, 31.77 and 35.58. The second year’s classifications thus give nearly 5000 and 6000 times as much information per plant as the first year’s classification.” (1930, p. 267).

The information/resources theme, i.e. point (ii) from mid-century (Introduction), appears when Fisher describes the value of such calculations (1930, p. 268):

The advantage of examining the amount of information gained at each stage of the experiment lies in the fact that the precision attainable in the majority of experiments is limited by the amount of land, labour and supervision available, and much guidance may be gained as to how these resources should best be allocated, by considering the quantity of information to be anticipated. In the experiment in question, for example, it appears that progenies from **OhAr** plants [(C)] are somewhat more profitable than those from **ohAr** plants [(B)].

Fisher’s use of the estimated value in computing the expected information dramatises the comparisons but is only applicable in situations where the implications of the results of a pilot experiment are being considered. The examples in chapter XI of the *Experiments* are all set at the pre-data stage of the investigation. One arrangement is better than another, although by how much will depend upon the (unknown) value of the parameter.

The subsection's sub-title "Design and Precision" indicates a second theme, the separation of points (i) and (ii) in a way that resembles Raiffa and Schlaifer's (1961) posterior and preposterior analysis:

If ... our object is merely to assign a standard error to a particular result, we may estimate the amount of information available directly by differentiating the expression for $\partial L/\partial p$ in the equation for maximum likelihood, using the actual numbers recorded in the classes observed... [This] gives 3725 as the total amount of information upon which our estimate has been based and 1.638 as the standard error of the estimate of the recombination percentage. It should be noted that an estimate obtained thus is in no way inferior to one obtained from the theoretical expectations; only that it gives no guidance as to the improvement of the conduct of the experiment. It might be said that owing to chance, ..., the experiment has given the small amount of 11 units less information than should be expected [3736] from the numbers classified.

In the 4th edition (1932, p. 193) Fisher re-did the calculations and changed his conclusion about the informativeness of this experiment, compared to expectations. However, this did not affect the general point.

In 1930 and -35 Fisher showed by example how to use information in design. He did not provide a general account of where the principle of using expected information in choosing between experiments would lead, or the obstacles it might face, e.g. if there is no experiment that produces more information in all parts of the parameter space. By the standards of later comparison-of-experiments theory—see §9 below—this is very primitive work.

1930 saw the introduction of the theme of information as a guide to the best allocation of resources in experiments. By 1933 (p. 50; cf. 1935, p. 217) the arithmetic on which the

“shifting moral balance” was based had been done:

By applying statistical methods not only to the interpretation but also to the design of experiments it is not uncommon for the value of the experiment to be increased five or tenfold, a result which could not be obtained from improved methods of interpretation only, unless previous methods had been excessively inefficient.

Here the commensurability of the “value” of an experiment and the “efficiency” of a “method of interpretation” is taken for granted.

Perhaps “methods of interpretation” was meant to suggest something more comprehensive than estimation. Fisher’s agricultural experiments usually led to a test but the lesson of 1930 was how better design produces better estimates—on average at least. Fisher’s attempt to use information to show how better design produces better tests is discussed over the next two sections. Here it is much less clear what he was trying to do, or, what he thought he had achieved! By contrast, the link between design and testing in Neyman-Pearson theory through power is clear. Neyman, Iwazskiewicz & Kołodziejczyk (1935) describe the link and, although Fisher opened—very aggressively—the discussion of this paper, he seems to have registered this aspect of the paper only 20 years later. See §8 below.

7 Precision and degrees of freedom

In §5 we saw the arithmetic of the effectiveness of blocking in reducing the standard deviation given in Chapter VIII of the *Methods* (1925). However, there is also a discussion of blocking—and replications—in Chapter V, on the comparison of means, where the design issues are linked more clearly to obtaining a better test and a tighter interval statement.

Examples 17 and 18 (1925, pp. 103-5) present 95% intervals for the difference between men's and women's heights. Fisher comments on the first unpaired comparison,

we have treated the two samples as *independent*, as though they had been given by different authorities; as a matter of fact, in many cases brothers and sisters appeared in the two groups; since brothers and sisters tend to be alike in stature, we have overestimated the probable error of our estimate of the sex difference. Wherever possible, advantage should be taken of such facts in designing experiments. In the common phrase, sisters provide a better "control" for their brothers than do unrelated women. The sex difference could therefore be more accurately estimated from the comparison of each brother with his own sister. (p. 104)

The first interval was from $4\frac{1}{2}$ to 5 inches but with the paired comparison (Example 18), "we may estimate the mean sex difference as $4\frac{3}{4}$ to 5 inches". The calculations are based on large-sample normality and Fisher presented no interval inferences using Student's distribution, which is the main topic for this chapter; for more on interval inference in Fisher's work, see Aldrich (2000).

Having introduced Student's distribution, Fisher (p. 111) discussed the impact of changing the number of replications:

The use of "Student's" distribution enables us to appreciate the value of observing a sufficient number of parallel cases; their value lies not only in the fact that the standard error of a mean decreases inversely as the square root of the number of parallels, but in the fact that the accuracy of our estimate of the standard error increases simultaneously. The need for duplicate experiments is sufficiently widely realised; it is not so widely understood that in some cases, when it is desired to place

a high degree of confidence (say $P = .01$) on the results, triplicate experiments will enable us to detect with confidence differences as small as one-seventh of those which, with a duplicate experiment, would justify the same degree of confidence.

The figure of “one-seventh” reflects the .01 points on t with 1 degree of freedom (63.657) and with 2 (9.925). In Neyman-Pearson terms, Fisher is *not* describing how the power function changes but how the .01 critical region of the t -test changes shrinks as the sample size changes. Fisher emphasised the importance of increased numbers of degrees of freedom elsewhere too, e.g. in Fisher & Wishart (1930, p. 12)

A valid test could be made based on only one degree of freedom for error, but as the table for z shows the values of z exceeded in 5 per cent. or 1 per cent. of trials are very large, and in consequence relatively large real effects must be judged insignificant, or in other words, escape detection.

Examples 19-20 of the *Methods* continue the paired versus unpaired comparison using the Cushny and Peebles data on the reaction of ten patients to two drugs. This gave Fisher a second example showing the advantage of pairing, an increased P -value. However there was a possible drawback of more elaborate designs: they would consume degrees of freedom so that, while the standard error being estimated was smaller, the estimate of it is less precise.

A more precise comparison is obtainable by [pairing] only if the corresponding values of the two series are positively correlated, and only if they are correlated to a sufficient extent to counterbalance the loss of precision due to basing our estimate of variance upon fewer degrees of freedom.

Example 21 illustrates how pairing may bring a loss of precision. There are 4 paired observations and with 3 degrees of freedom the value of P is between .01 and .02. If the observations

are unpaired, there are 6 degrees of freedom and the value of t is so large that it is beyond the range of the table, showing that P is “extremely small”.

These considerations are not applied to design when that topic is discussed in Chapter VIII of the book but the question is raised in relation to the Latin square in Eden and Fisher (1929, p. 204). There the matter rested until 1935.

8 “The information lost in the estimation of error”

“The information lost in the estimation of error”, the last section of the information chapter (XI) of the *Experiments* (1935, pp. 247-50), addresses the issue of precision and the number of degrees of freedom: “There is ... one further point in connection with experiments involving measurements, to which the theory of errors is applicable, which may be cleared up by the methods of this chapter.”

Fisher (p. 248) explains the issue, using the language of “absolute precision,”

When, as a result of an experiment, a value x has been assigned a sampling variance, s^2 , validly and correctly estimated from n degrees of freedom, the position is not the same as if the variance were known with exactitude. ... We need ... in considering the absolute precision of an experimental result, to take into account, not only the estimate s^2 derived from the data, but also the number of degrees of freedom upon which our estimate, s^2 , was based.

In this formulation, the “amount of information supplied by an observed value, x ” about μ in the case of *known* σ^2 (given a normal density for x) is $\frac{1}{\sigma^2}$.

“The methods of this chapter” (§6 above) are not applied mechanically for now there are two parameters and, besides, there is scope for using the fiducial argument, which Fisher

thought inapplicable in the case of discrete sample spaces. He had first used the fiducial argument in 1930—see Aldrich (2000)—and he described the fiducial distribution for μ , obtained from the t -distribution for $(x - \mu)/s$, twice in 1935—in the *Experiments* (ch. X) and in “The Fiducial Argument in Statistical Inference”. For a brief guide on matters fiducial, see Stone (1983).

Fisher (1935, p. 248) calculates the information given by the experiment about μ using the t -distribution in place of the normal curve, arguing fiducially:

the probability that the quantity t , defined by the relationship

$$x - \mu = st,$$

where x is the observed value and μ the hypothetical value of which it is an estimate, shall lie in any assigned range, dt , is given by the formula

$$df = \frac{\frac{n-1}{2}!}{\frac{n-2}{2}! \sqrt{\pi n}} \cdot \frac{dt}{\left(1 + \frac{t^2}{n}\right)^{\frac{1}{2}(n+1)}}$$

or in terms of x and μ , by

$$df = \frac{\frac{n-1}{2}!}{\frac{n-2}{2}! s \sqrt{\pi n}} \cdot \frac{dx}{\left(1 + \frac{(x-\mu)^2}{ns^2}\right)^{\frac{1}{2}(n+1)}}.$$

This derivation of a “ $df\dots dx$ ” density for x from the “ $df\dots dt$ ” density was distinct from Fisher’s other fiducial efforts, the familiar “ $df\dots d\mu$ ” distribution for μ , also from 1935, and the “fiducial prediction” from a normal sample in *Statistical Methods and Scientific Inference* (1956, ch V, §4). The “ $df\dots dx$ ” argument does not seem to have been discussed in the literature on the fiducial argument.

To finish, Fisher does what he had been doing since the 20s, he uses the density for x to calculate the information relating to μ . This turns out to be

$$\frac{n+1}{(n+3)s^2}.$$

He comments on the result

It appears that the true precision of our estimate is somewhat lower than it would have been, had the variance been known with exactitude to be s^2 . In the extreme case, when $n = 1$, and the estimate is based on only 1 degree of freedom, the precision is halved. And in general, the true precision is less than it might be thought, if the uncertainty of our estimate of the variance were ignored, by the fraction $2/(n+3)$. It may thus be worth while to sacrifice, to some small extent, the aim of diminishing the value of s^2 , if this diminution carries with it any undue reduction in the number of degrees of freedom, available for the estimation of error.

There are two concerns here: how is the precision *achieved* by an experiment is to be measured and how is the choice between experiments to be made. The second is a continuation of the precision and degrees of freedom theme discussed in the last section.

In the seventh edition of the *Methods* Fisher (1938, pp. 274-5) incorporated the correction in his reworking of the canonical analysis of 1925 (see §5 above) to find the “amounts of information elicited by different methods”:

In this case, as in many others, the lower mean square is obtained at the expense of some reduction of the number of degrees of freedom on which the estimate of error is based. This makes the test of significance somewhat less stringent. If n is the number of degrees of freedom for error, the loss of information due to this cause is

found... to be the fraction $2/(n + 3)$, so that taking this factor into consideration we may summarise the results as follows:—

| | d.f. for error | units of information/crude | adjusted |
|----------------------------|----------------|----------------------------|----------|
| Randomisation of 20 plots | 15 | 28.18 | 25.05 |
| Randomisation in 4 blocks | 12 | 51.07 | 44.26 |
| Eliminating order in block | 11 | 80.07 | 68.03 |

Even when allowance is thus made for the degrees of freedom absorbed, it is clear that in this case both the use of blocks, and that of order within the block have been extraordinarily profitable.

I have presented the construction of this table in two stages (the first was described in §5) but Fisher’s readers got both stages together. It is possible that Fisher only went fully informational when he had found a resolution of the degrees of freedom puzzle.

The analysis provoked some irritated private discussion. M. S. Bartlett (1936) was also working on information and he wrote to Fisher (Bennett pp. 46-51) about the analysis in the *Experiments* and, for a few weeks around the New Year of 1936, they tried unsuccessfully to make themselves understood. Fisher returned to the subject in the 50s, or rather it returned to him. In 1955 Finney wrote to him with a “little paradox” that “the average of the precision is greater than if σ^2 were known.” (Bennett (1990) pp. 95-103)). Finney added a P.S. to one of his letters, “I see no reason why you should spend time in putting me right, but I thought the paradox might amuse you.” (p. 101) Fisher did not rise to the challenge, nor did he add anything about what he hoped to achieve with the measure.

As far as Fisher was concerned, his two contributions were the final word, or rather the *only* word. He learnt otherwise when John Nelder reviewed Federer’s textbook, *Experimental*

Design. Nelder (1956, p. 341) criticised the book for not evaluating the work it précised and he instanced its presentation of Fisher's measure for "the relative efficiency of two designs with errors based on different numbers of degrees of freedom". Federer was at fault because the measure depends on the fiducial argument "not everywhere accepted (and of a type which the author nowhere else uses)" and because "other intuitively reasonable measures give rather different results." Fisher felt *he* was being smeared and complained to Nelder. At the same time he wanted to know what those "other" measures were. Nelder and Finney (also consulted) referred him to Cochran & Cox (1950); see Bennett (pp. 257-9, 280-3). Cochran and Cox (p. 18) mention the "complicated tables" in Neyman, Iwazkiewicz & Kołodziejczyk (1935) and then give an argument "which though logically faulty, leads to an approximation that is good enough for most practical purposes." They (p. 28) also report Fisher's argument.

Nelder's reply included a general objection to the fiducial argument. Fisher did not expect to be answered back and he dismissed Nelder like a schoolboy, who had not paid enough attention to his teachers, or who had been taught badly. (Bennett p. 282; see also Senn (2003)). All this activity had a precipitate in the next editions of the *Methods* and the *Experiments*. In the *Methods* Fisher (1958, p. 267) made a last addition to the analysis of 1925 (§5 above):

It should be stated that the adjustment ... is correct and exact, since more complicated and inexact adjustments have been proposed by later writers who evidently have not understood the problem.

Fisher wrote at greater length in the *Experiments* (1960, p. 244). He made a new point, that the omnibus character of the adjustment/allowance made it an advance on what Gosset could have done in 1908 and what he (Fisher) had done in 1925 (§7 above):

Writers averse to the use of fiducial probability (Neyman, Walsh) have attempted solution of this problem by other paths. The results are immensely complex, and these writers seem to overlook the fact that for *given* levels of significance, *Student* has already solved the problem in 1908. What is needed in experimental design is knowledge of the allowance to be made for the limitation of the number of degrees of freedom, before the level or levels of interest are known. The allowance found in this Section have been in use now for 25 years without any serious alternative. It is not an approximation, or a provisional artifice, and has been unintelligible only to those who over a long period resisted the cogency of the fiducial argument.

This section has followed Fisher's attempt to ground in an information measure the relation between degrees of freedom, precision and the sensitivity of tests, which had first exercised him in 1925. The point was "cleared up" to his own, if not to anybody else's, satisfaction.

9 Decisions: economics and polemics

Finally we turn to the economic element in Fisher's teaching, how he had taught us to conduct experimental and observational inquiries so as "to maximise the information obtained for a given expenditure". We also examine his criticisms of Wald's decision theory and note some similarities between his work and that of others on comparing experiments.

In 1926 John Russell, the Rothamsted director, complained that many experiments "do not give as much information as might have been hoped, considering the cost"; he then showed "how experiments can be arranged to give the maximum return for the work and money spent on them." (p. 989) Whether Fisher took this line from Russell or vice versa is unclear for the methods Russell was describing were inspired by Fisher. On that day, however, "information"

appears to have been a technically innocent term.

In the *Experiments* Fisher (1935, p. 189) wrote earnestly—and more concretely than in 1930 (§6 above)—about the need for “information costings”:

it is important that the items of labour and skilled supervision chargeable to a particular method of experimentation, shall be fairly and carefully recorded and calculated. For any time and labour devoted to experimental work must be regarded as having been diverted from other work of scientific value, to which they might otherwise have been given. Even rough costings of this kind will usually show that the efficiency with which limited resources can be applied, is capable of relatively enormous increases by careful planning of the experimental programme, and there is nothing in the nature of scientific work which requires that the allocation of the resources to the ends aimed at should be in any degree rougher, or less scrupulous, than in the case of a commercial business.

Fisher probably thought that a detailed cost analysis, as a piece of accounting, would have been out of place in the *Experiments*. The closest he comes is when he evaluates a year of uniformity trials prior to the experiment proper being done:

It may be noted ... that with annual agricultural crops, knowledge of the yields of the experimental area in a previous year under uniform treatment has not been found sufficiently to increase the precision to warrant the adoption of such uniformity trials as a preliminary to projected experiments. Such a procedure necessarily doubles the experimental labour, and as it is not found to double the amount of information but to increase it, perhaps by 50 per cent., it is clearly unprofitable. For, by the application of twice the expenditure in time and money in the experimental

year, the amount of information recovered may with confidence be expected to be approximately doubled. (1935, p. 166)

The parallel discussion in the *Methods* appears in the section on the analysis of covariance, the technique by which the previous year's experience would be used. There he (1934, p. 271) remarks, "The chief advantage of the analysis of covariance lies ... in the guidance it is capable of giving in the design of an observational programme, and in the choice of which of many concomitant observations shall in fact be recorded."

The economic theme, unlike the information theme, was uncontroversial for it was regarded as a truism that design mattered because resources were limited—it was the premiss for Gosset's work, described in §3 above. Fisher did not develop a rich *economics* of experimentation because he believed that, for given expenditure, there was a best way of increasing information; all that was needed was a calculation giving the information yield of each design. There would have been more to say if he had asked how much the information is *worth* but, as we now see, he held such an evaluation impossible in the case of scientific research.

Fisher reacted strongly to the statistical decision theory of Wald. He attacked *Statistical Decision Functions* (1950) in an article, "Statistical Methods and Scientific Induction" (1955), in a book *Statistical Methods and Scientific Inference* (1956) and in passages he added to the *Methods* (1958) and the *Experiments* (1960). Only in the article (1955, p. 70) did he mention Wald's version of the spanning formula, "the treatment of the design of experimentation as a part of the general decision problem", but he was too outraged, at the idea of decision theory and at Wald's ignorance in writing about design without mentioning replication, control or randomization (the three "elements") to discuss Waldian spanning.

Fisher's criticisms sometimes have a surreal quality, as when he (1956, p. 100) conjures a dystopia where "a giant computer programmed with Decision Functions" replaces "responsible

and independent” scientific thinkers and the Cold War is not far away when he writes about the totalitarian potential of decision theory; see Marks (2003). Equally remarkable is the standpoint of the purest of pure scientist. Although in 1956 Fisher’s stint at Rothamsted was more than 20 years in the past, he continued to follow the work there and to identify with its objectives. The point of testing varieties and fertilisers was to find better yielding varieties and more effective fertilisers. “The object ... is to find out which will pay the farmer best” declared ‘Student’ in his “On Testing Varieties of Cereals” (1923, p. 271). Treating expenditure on experiments as an investment, Gosset estimated that new varieties of barley in Ireland had produced a gain of not less £250,000 per year and “as the cost of experiments from the commencement to the present time cannot have reached £40,000 the money has been well spent.” Gosset was in business but Fisher (1926, p. 503) was not above considering the confidence of the “purchasing public” when establishing the effectiveness of a manure.

Statistical Methods and Scientific Inference has the familiar phrases about experiments, information and resource cost (pp. 5 & 139) but design is not treated in detail. Decision theory is treated in the chapter, “Some Misapprehensions about Tests of Significance.” Fisher (pp. 75-78) begins by arguing that the Neyman-Wald account of testing is based on a false analogy between “tests of significance” and “acceptance decisions”. Tests of significance were originally means by which the “the research worker gains a better understanding of his experimental material.” However the tests have been reinterpreted “as means to making decisions in an acceptance procedure.” Fisher holds that the differences between the two situations are “many and wide”, One of the “deepest” lies in the population: for the quality controller, the population has an unequivocal existence as a particular batch of goods, to be accepted or rejected as substandard, while in scientific inference the population is the “product of the statistician’s imagination” (p. 77). Most of the chapter—and all of the analysis—is given to

discussing what that population should be in the case of regression, the 2×2 table and the Behrens-Fisher problem.

The other “dissimilarity” is in the use of a loss function for making the wrong decision. “It is important that the scientific worker introduces no cost functions for faulty decisions, as it is reasonable and often necessary to do with an Acceptance Procedure. To do so would imply that the purposes to which new knowledge was to be put were known and capable of evaluation.” (p. 102, cf. (1955, p. 77)). Gauss, who Fisher greatly admired and whose goals were as scientific, used a loss function in his second theory of estimation. In Fisher’s decision theory there was an incommensurability between costs measured in money and value measured in units of information.

An interesting side-light on Fisher’s decision theory is provided by Lindley’s “On a Measure of Information Provided by an Experiment” (1956). This paper belongs to the comparison of experiments literature which Blackwell (1951) and others developed from Wald’s theory; for a recent review see Goel & Ginebra (2003). Lindley integrated the Shannon (1948) notion of information and Bayesian decision theory by putting information into the utility function. Lindley (p. 999) refers to Fisher, the critic of decision theory, not the theorist of information and experiments. With a nod to Fisher (1955), Lindley suggests the object in experimentation is “not to reach decisions but rather to gain knowledge about the world.”. Replacing the decision objective by an inference objective, Lindley measures success by expected information gain. His (p. 987) maxim “perform that experiment for which the expected gain in information is the greatest and continue experimentation until a preassigned amount of information has been attained” sounds like a close relative of Fisher’s “to conduct experimental and observational inquiries so as to maximise the information obtained for a given expenditure.” While Lindley mentions Fisher’s objections to the view that “modern statistics *is* decision theory”

and to the interpretation of significance tests as Wald decision problems, he does not mention that Fisher had views on the “measure of information provided by an experiment.” In the later literature “Fisher information” is used as a technical tool, e.g. in Kiefer (1959), but with no awareness that Fisher was first to use this tool to compare experiments. (§6 above)

10 Summary and conclusions

This investigation started from a description of where statistics stood in the mid-20th century and a story of how it had got there. I have argued that description and story reflect an information turn in Fisher’s thought on experiments in the course of the 1930s. I now summarise the argument and speculate on why the development left so little mark on the literature. The only paper I know linking Fisher’s ideas on design and his ideas on estimation, Seidenfeld (1992), does not mention that Fisher made his own links.

One reason why Fisher’s information turn had so little effect was that statisticians were turning in other directions. They were following Neyman (1935, p. 75) in his judgement that the “frequency of errors in judgement” provides a “sufficiently simple and unquestionable principle in statistical work” while the “amount of information is too complicated and remote to serve as a principle”. Fisher’s information addition to Gosset’s (and his own) standard deviation was the proverbial fifth wheel. In this connection it is worth noting that when Finney (1955, p. 16) treated Fisher’s situations of a “more general character” (those described in §6 above) he worked entirely in terms of standard errors.

Statisticians were not interested in looking and was there much to see? There was some stamping of INFORMATION on old work and a flow of aphorisms but only two new technical contributions described in §§6 and 8. These were noticed: one had a life in the specialised world

of genetics—see Mather (1938)—and the other went into textbooks on design. The *Methods* was an old book and the contributions were additions to existing analyses but the *Experiments* was new and could have been built around information. However there were other things the author wanted to emphasise. The information principle never found expression in a structure comparable to Wald’s *Statistical Decision Functions* or Raiffa and Schlaifer’s *Applied Statistical Decision Theory*. Fisher’s last book, *Statistical Methods and Scientific Inference* (1956), emphasised the *variety* of inferences and information is confined to the chapter on the principles of estimation

To say that Fisher faced an unreceptive audience, or that he did not write a book on information, is not to say that the statements, or the work behind them, had no significance for him. In the '20s he had been hugely impressed by what he had discovered about information; his “reduction of data” definition of statistics (§1) applied his deepest insight in the broadest possible way. The spanning formula of 1951 was a natural development of this definition to accommodate the design of experiments. The two technical contributions were a design addition to his canonical statement on estimation and an information addition to his canonical statement on the advantage of design. Fisher’s practice was to use examples to make general points and these were important additions to major examples. The work described in §8 attempting to link information and testing had significance for Fisher, even if his readers were unpersuaded by it.

The story and the spanning formula date from 1947/51 but by 1935 they would have seemed to Fisher natural ways of expressing his achievement. His account of how “armed with amount of information as a practical tool, statisticians came to study what forms of experiment ... would yield the most information” is literally false, given that so much of his work on experiments was done without apparent benefit of this “practical tool”. Yet, when

he started to bring information into experiments it was with the understanding that two of the “three elements” of replication, local control and randomization were manifestations of the information principle. He had been using the tool without giving it its proper name.

References

Note: all of Fisher’s papers referred to appear in J. H. Bennett (ed) *Collected Papers of R. A. Fisher*, 5 volumes, Adelaide: Adelaide University Press (1971-4) and many, with Bennett (1990), are available from the University of Adelaide R. A. Fisher Digital Archive

<http://www.library.adelaide.edu.au/ual/special/fisher.html>

Aldrich, J. (1997) R. A. Fisher and the Making of Maximum Likelihood 1912-22, *Statistical Science*, **12**, 162-176.

_____ (2000) Fisher’s ‘Inverse Probability’ of 1930, *International Statistical Review*, **68**, 155-172.

_____ (2003) *A Guide to R. A. Fisher*, website

<http://www.economics.soton.ac.uk/staff/aldrich/fisherguide/rafreader.htm>

_____ (2004) E. S. Pearson’s reviews of R. A. Fisher’s *Statistical Methods for Research Workers*, website

<http://www.economics.soton.ac.uk/staff/aldrich/fisherguide/esp.htm>

Atkinson, A. C. & R. A. Bailey (2001) One Hundred Years of the Design of Experiments on and off the Pages of “Biometrika”, *Biometrika*, **88**, 53–97.

Bartlett, M. S. (1936). Statistical Information and Properties of Sufficiency, *Proceedings of the Royal Society, A*, **154** 124-137.

Bennett, J. H. (1990) (ed) *Statistical Inference and Analysis: Selected Correspondence of R. A. Fisher*, Oxford, University Press.

Blackwell, D. (1951) Comparison of Experiments, *Proceedings of the 2nd Berkeley Symposium*

on *Mathematical Statistics and Probability*, pp. 93-102. Berkeley: University of California Press.

Box, J. F. (1978) *R. A. Fisher: The Life of a Scientist*, New York, Wiley.

_____ (1980) R. A. Fisher and the Design of Experiments, 1922-26, *American Statistician*, **34**, 1-7.

Cochran, W. G. (1980) Fisher and the Analysis of Variance, pp. 17-34 in Fienberg & Hinkley (1980).

Cochran, W. G. & G. M. Cox (1950) *Experimental Designs*, 2nd. edition 1960, New York: Wiley.

Eden, T. & R. A. Fisher (1927) Studies in Crop Variation. IV. The Experimental Determination of the Value of Top Dressings with Cereals, *Journal of Agricultural Science*, **17**, 548-562.

_____ (1929) Studies in Crop Variation. VI. Experiments on the Response of the Potato to Potash and Nitrogen, *Journal of Agricultural Science*, **19**, 201-213.

Edwards, A. W. F. (2000) Fisher Information and the Fundamental Theorem of Natural Selection. *Rendiconti (B) Istituto Lombardo di Scienze e Lettere*, Milano.

Engledow, F. L. (1925) Theory of Experimental Error, *Journal of the Ministry of Agriculture*, **3?**, 326-333.

Federer, W. T. (1955) *Experimental Design: Theory and Application*, New York: Macmillan.

Fienberg, S. E. & D. V. Hinkley (1980) (eds) *R. A. Fisher: An Appreciation*, New York: Springer.

Fisher, R. A. (1922) On the Mathematical Foundations of Theoretical Statistics *Philosophical Transactions of the Royal Society, A*, **222**, 309-368.

_____ (1924) On a Distribution Yielding the Error Functions of Several Well Known

- Statistics, *Proceedings of the International Congress of Mathematics, Toronto*, **2**, 805-813.
- _____ (1925) *Statistical Methods for Research Workers*, Edinburgh: Oliver & Boyd.
- _____ (1925a) Theory of Statistical Estimation, *Proceedings of the Cambridge Philosophical Society*, **22**, 700-725.
- _____ (1926) The Arrangement of Field Experiments, *Journal of the Ministry of Agriculture*, **33**, 503-513.
- _____ (1929) Statistics in Biological Research, Letter in *Nature*, August 17th, 266-7.
Reproduced with notes in Aldrich (2004).
- _____ (1930) *Statistical Methods for Research Workers*, third edition, Edinburgh: Oliver & Boyd.
- _____ (1933) The Contributions of Rothamsted to the Development of the Science of Statistics, *Annual Report of the Rothamsted Experimental Station*, 43-50.
- _____ (1934) Two New Properties of Mathematical Likelihood. *Proceedings of the Royal Society, A*, **144** 285-307.
- _____ (1935) *The Design of Experiments*, Edinburgh: Oliver & Boyd.
- _____ (1935a) The Fiducial Argument in Statistical Inference, *Annals of Eugenics*, **6**, 391-398.
- _____ (1938) *Statistical Methods for Research Workers*, seventh edition, Edinburgh: Oliver & Boyd.
- _____ (1938a) Presidential Address First Indian Statistical Conference 1938, *Sankhya*, **4**, 14-17.
- _____ (1947) Development of the Theory of Experimental Design, *Proceedings of the International Statistical Conferences*, (Washington), **3**, 434-439.
- _____ (1951) Statistics, in A. E. Heath (ed) *Scientific Thought in the Twentieth Century*,

London: Watts.

_____ (1955) Statistical Methods and Scientific Induction, *Journal of the Royal Statistical Society, B*, **17**, 69-78.

_____ (1956) *Statistical Methods and Scientific Inference*, Edinburgh: Oliver & Boyd.

Fisher, R. A. & T. Eden (1929) Studies in Crop Variation. VI. Experiments on the Response of the Potato to Potash and Nitrogen, *Journal of Agricultural Science*, **19**: 201-213.

Fisher, R. A. & W. A. Mackenzie (1923) Studies in Crop Variation. II. The Manurial Response of Different Potato Varieties, *Journal of Agricultural Science*, **13**, 311-320.

Fisher, R. A. & J. Wishart (1930) The Arrangement of Field Experiments and the Statistical Reduction of the Results, *Imperial Bureau of Soil Science, Technical Communication*, **10**.

Goel, P. K. & J. Ginebra (2003) When is One Experiment 'always better than' Another, *Statistician*, **52**, 515-537.

Hinkley, D. V. (1980) Theory of Statistical Estimation: the 1925 Paper, pp.685-94 in Fienberg & Hinkley (1980).

_____ (1980a) Fisher's Development of Conditional Inference, pp.6101-108 in Fienberg & Hinkley (1980).

_____ (1980b) Contribution to the discussion of Basu's "Randomization Analysis of Experimental Data: The Fisher Randomization Test", *Journal of the American Statistical Association*, **75**, 582-584.

Kiefer, J. (1959) Optimum Experimental Designs (with discussion), *Journal of the Royal Statistical Society, B*, **21**, 272-319.

Kempthorne, O. (1952) *The Design and Analysis of Experiments*, New York: Wiley.

_____ (1966) Some Aspects of Experimental Inference, *Journal of the American Statistical Association*, **61**, 11-34.

- Lindley, D. V. (1956) On a Measure of Information Provided by an Experiment, *Annals of Mathematical Statistics*, **27**, 986-1005.
- McMullen, L. (ed.) (1970) *Letters from W. S. Gosset to R. A. Fisher 1915-1936* with summaries by R. A. Fisher with a Foreword by L. McMullen, printed by Arthur Guinness for private circulation.
- Marks, Harry M (2003) Rigorous Uncertainty: Why R. A. Fisher is Important, *International Journal of Epidemiology*, **32**, 932-937.
- Mather, K (1938) *The Measurement of Linkage in Heredity*, London: Methuen.
- Mercer, W. B. & A. D. Hall (1911) The Experimental Error of Field Trials, *Journal of Agricultural Science*, **4**, 107-132.
- Nelder, J. A. (1956) Review of *Experimental Design, Theory and Application* by Walter T. Federer, *Journal of the Royal Statistical Society, A*, **119**, 340-341.
- Neyman, J. (1935) Discussion of Fisher (1935a), *Journal of the Royal Statistical Society*, **98**, 73-76.
- Neyman, J., K. Iwazskiewicz and St. Kołodziejczyk (1935) Statistical Problems in Agricultural Experimentation, *Journal of the Royal Statistical Society Supplement*, **2**, 107-154.
- Pearce, S. C. (1979) Experimental Design: R. A. Fisher and some Modern Rivals, *Statistician*, **28**, 153-161.
- Pearson, E. S. (1990) 'Student', *A Statistical Biography of William Sealy Gosset*, Edited and Augmented by R. L. Plackett with the Assistance of G. A. Barnard, Oxford: University Press.
- Preece, D.A. (1990) R. A. Fisher and Experimental Design: A Review. *Biometrics*, **46**, 925-935.
- Raiffa, H. A. and R. Schlaifer (1961) *Applied Statistical Decision Theory*, Boston: Graduate School of Business Administration, Harvard University.

- Russell, J. (1926) Field Experiments: How they are made and What they are for, *Journal of the Ministry of Agriculture*, **32**, 989-1001.
- Seidenfeld, T.(1992) R. A. Fisher on the Design of Experiments and Statistical Estimation, in S. Sarkar (ed) *The Founders of Evolutionary Genetics*, Dordrecht: Kluwer.
- Senn, S. (2003) A Conversation with John Nelder, *Statistical Science*, **18**, 118-131.
- Shannon, C. E. (1948) A Mathematical Theory of Communication, *Bell System Technical Journal*, **27**, 379-423 & 623-656.
- Stone, M. (1983) Fiducial Probability, pp. 81-85 of S. Kotz & N. L. Johnson (eds) *Encyclopedia of Statistical Science volume 3*, New York: Wiley.
- Street, D. J. (1990) Fisher's Contributions to Agricultural Statistics. *Biometrics*, **46**, 937-945.
- 'Student' (1908) The Probable Error of a Mean, *Biometrika*, **6**, 1-25.
- _____ (1911) Note on A Method of Arranging Plots so as to Utilize a given Area of Land to the Best Advantage in Testing Two Varieties, Appendix to Mercer and Hall's "The Experimental Error of Field Trials", *Journal of Agricultural Science*, **4**, 128-.
- _____ (1923) On Testing Varieties of Cereals, *Biometrika*, **15**, 271-293. Amendment and correction **16**, 1924, p. 411.
- Wald, A. (1950) *Statistical Decision Functions*, New York: Wiley.
- Walsh, J. E. (1949) On the "Information" Lost by using a *t*-test when the Population Variance is Known, *Journal of the American Statistical Association*, **44**, 122-125.
- Welch, B. L. (1937) On the *z*-test in Randomized Blocks and Latin Squares, *Biometrika*, **29**, 21-52.
- Yates, F. (1964) Sir Ronald Fisher and the Design of Experiments, *Biometrics*, **20**, 307-321.
- Youden, W. J. (1951) The Fisherian Revolution in Methods of Experimentation, *Journal of the American Statistical Association*, **46**, 47-50.