The Theory of Statistical Decision

L. J. Savage

Journal of the American Statistical Association, Volume 46, Issue 253 (Mar., 1951), 55-67.

Your use of the JSTOR database indicates your acceptance of JSTOR's Terms and Conditions of Use. A copy of JSTOR's Terms and Conditions of Use is available at http://www.jstor.org/about/terms.html, by contacting JSTOR at jstor-info@umich.edu, or by calling JSTOR at (888)388-3574, (734)998-9101 or (FAX) (734)998-9113. No part of a JSTOR transmission may be copied, downloaded, stored, further transmitted, transferred, distributed, altered, or otherwise used, in any form or by any means, except: (1) one stored electronic and one paper copy of any article solely for your personal, non-commercial use, or (2) with prior written permission of JSTOR and the publisher of the article or other text.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

Journal of the American Statistical Association is published by American Statistical Association. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/asa.html.

Journal of the American Statistical Association
©1951 American Statistical Association

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2000 JSTOR

http://www.jstor.org/
Wed Mar 15 19:07:50 2000

THE THEORY OF STATISTICAL DECISION

L. J. SAVAGE
University of Chicago

ABRAHAM WALD'S recent book, *Statistical Decision Functions* [10], presents a new theory of the foundations of statistics.[1] The vigorous exploration of this theory was begun by Professor Wald five or six years ago and is being continued under his leadership. Since almost all published treatments of this theory known to me, including this one, are mathematically forbidding,[2] and since the theory promises to be of great interest to all statisticians, it seems appropriate to attempt an informal exposition of it. The critical and philosophical remarks in this exposition may not accurately represent the views of Professor Wald, for both in writing and lecturing, he prefers to be rather noncommittal on such points.

Traditionally, the central problem of statistics is to draw statistical inferences, that is, to make reasonably secure statements on the basis of incomplete information. This entails other problems, particularly that of designing experiments which permit the strongest inference for the expenditure involved. The new theory under discussion, however, centers about the problem of statistical action rather than inference, that is, deciding on a reasonable course of action on the basis of incomplete information. There is clearly an abundance of situations calling for statistical action. Industrial quality control is a clear cut and familiar domain of examples. The problem of design evoked by the traditional inference viewpoint is also a problem of statistical action. Much more generally, it can be argued that all problems of statistics, including those of inference, are problems of action, for to utter or publish any statement is, after all, to take a certain action. Since the consequences of statements regarded as actions, especially their influences on those to whom they are directed, are often unusually difficult to analyze and appraise, for example, when the statement refers to academic science, it is well to point out that many statements, especially of applied science, are tantamount to deciding on a concrete action. Thus, medical diagnosis cannot in principle be separated from the choice of treatment. The typical agronomic experiment, though

[1] This paper was undertaken as a review of Wald's book. In requesting the review, the Editor suggested that some exposition of the field to which the book pertains be given prefatory to the review. This exposition so extended itself that the paper is no longer exclusively a review.

This paper is in part based on research supported by the Office of Naval Research.

[2] Some general exposition of the theory and an original example are given elementary treatment in [7].

commonly discussed in the language of inference, is plainly concerned almost exclusively with deciding on a course of action.

As has already been suggested, the term "course of action," or briefly "act," is to be understood in a flexible sense. The following chain of examples illustrates how different may be the scopes of the acts associated with a single situation, depending on context:

(i) Decide whether to accept, i.e., to buy, a certain lot of screws.
(ii) Decide among all possible rules for accepting or rejecting the lot, based on the thousand and one possible outcomes of inspection of a thousand screws chosen from the lot at random and classified as defective or nondefective.
(iii) Decide among all possible experiments, including sequential ones, and all possible associated acceptance rules for accepting or rejecting the lot.

As the third example suggests, the whole design of a complicated statistical program can be regarded as a single decision to adopt one of an enormous number of possible acts. In a highly idealized sense, a man might be thought of as making only one decision in his whole life, namely the decision to conduct himself according to some set of maxims envisaging all eventualities. This is manifestly far-fetched, but the theory does call attention to the propriety of considering very large decision problems as organic wholes. In practice, it will presumably continue to be necessary to break such big decision problems into pieces, rigorously or approximately independent of one another.

Acts have consequences for the actor, and these consequences depend on facts, not all of which are generally known to him. The unknown facts will often be referred to as states of the world, or simply states. Suppose for example, that on a given occasion you face the decision of whether or not to carry an umbrella, and consider for simplicity only two possible states, future rain and future shine. The consequences of each act in each possible state might be given by the following table:

Act \ State	Rain	Shine
Carry	Inconvenience and wet feet	Inconvenience and slight embarrassment
Don't carry	Miserable drenching	Bliss unalloyed

In some situations the possible consequences of each act in each state are cash incomes, or at any rate expressible in terms of cash incomes. Thus in the rain-shine example above, you might perhaps consider the monetary values expressed in dollars by the following "income matrix" reasonable:

Act \ State	Rain	Shine
Carry	4	5
Don't Carry	−10	10

It will be assumed in this exposition, except in the next paragraph where some other possibilities are mentioned in passing, that in the contexts to be discussed the consequence of any act is to all intents and purposes cash income (possibly zero or negative) which may depend on which of the possible states of the world is in fact the real one, and that a situation with a known expected (or actuarial) cash income is as valuable to the actor as a pure cash income of the same amount. The terms are to be understood in such a sense that the actor knows how his income depends on his action and the unknown state of the world, i.e., he knows the income matrix of the situation.

The assumptions just made are more restrictive than is logically necessary but are adopted here for definiteness. It is enough that to each possible consequence there should correspond a numerical score such that higher scores are preferred to lower scores, and a given expected score is as acceptable as the corresponding sure score. Cash income, which has here been adopted as the epitome of such a score, does in many situations have the postulated properties. Other possible scores are expected number of lives saved, and probability of over-all success in some venture. The von Neumann-Morgenstern theory of utility[3] suggests that very generally there will exist some suitable score.

If in a given situation the actor assigns probabilities to the various unknown states of the world, he can calculate unambiguously the expected cash income associated with any action. By the assumptions already made, he then acts in such a way as to maximize his expected

[3] This theory in its modern form is discussed by its authors in [9]. The more elementary [5] may also be consulted. Paul A. Samuelson has pointed out to me that [5] errs in asserting that three propositions listed on p. 288 are equivalent to the von Neumann-Morgenstern system.

cash income, and the decision problem is relatively trivial and old-fashioned.[4] If the actor does not assign probabilities, this trivial solution does not apply and the problem is newer. In fact, the problem of dealing with uncertainty when probability does not apply to the unknown states of the world has been the main theme of modern, or un-Bayesian, statistical theory. The actor may attach probability to some aspects of the unknown states but not to others. For example, he might have to do something about a coin known to have an unknown probability of heads. This seemingly more general situation can be shown to reduce to that in which no probabilities are assigned. Suppose, for example, the problem is to guess how a certain coin will fall on a single trial, possibly after some experimentation, and that it is known that for this coin heads had probability either one-third or two-thirds. Clearly the problem reduces to that of guessing which of the two possible probabilities applies.

Even if probabilities are not assigned, there is one unquestionably-appropriate criterion for preferring some act to some others: If for every possible state, the expected income of one act is never less and is in some cases greater than the corresponding income of another, then the former act is preferable to the latter. This obvious principle is widely used in everyday life and in statistics, but only occasionally does it lead to a complete solution of a decision problem.

Before going further, the idea of a mixed act, or mixed strategy, as it is commonly called, should be introduced. Instead of choosing a particular act outright, the actor might prefer to let chance play such a role in his choice that each act has a probability determined by the actor of being chosen. For example, in a choice between two acts, the actor might flip two coins, agreeing in advance to choose the first act if both coins fall heads, and otherwise to choose the second. The admission of mixed acts obviously does not restrict the actor since he can, if he chooses, attach probability one to any particular act, that is, choose it outright. On the other hand, it is not offhand clear that it is of any advantage to consider mixed acts; indeed, the contrary can be seriously argued. The fact that the process of randomization, so essential to modern statistical methods, is an employment of mixed acts, certainly suggests that they may be advantageous.

To state the general rule by which the theory of statistical decision functions tentatively proposes to solve all decision problems, some

[4] In the rain-shine example, if probability applies to the states easy calculation shows that you should carry, not carry, or be indifferent according as the probability of rain is more than, less than, or equal to 5/19.

mathematical notation is appropriate. Let $I(a, s)$ denote the expected income resulting from the (mixed) act a if the world is in state s, and assume for simplicity that the number of possible states s and of unmixed acts is finite, though possibly enormous. Understand by the loss, $L(a, s)$, associated with the act a and state s, the difference between (i) the most that can be made by any act if the state s obtains and (ii) the amount the act a makes if s obtains; that is, in notation which will explain the words to the mathematical and the notation "max" to the verbal,

$$L(a, s) = \max_{a'} I(a', s) - I(a, s).$$

The loss may be said to measure how inappropriate the action a is in the state s. In the rain-shine example, it is easily verified that the loss for the unmixed acts is:

Act \ State	Rain	Shine
Carry	0	5
Don't carry	14	0

If you give yourself probability P of carrying the umbrella and $1-P$ of not carrying it, the loss is given by:[5]

Act \ State	Rain	Shine
P	$14(1-P)$	$5P$

The following general rule, called the minimax principle, is central to the theory of statistical decision functions, at least today. In fact, it is the only rule of comparable generality proposed since Bayes' was published in 1763 [1]. The minimax principle is: Choose such an a that $\max_s L(a, s)$ shall be as small as possible, that is, minimize the maximum loss. There may be more than one act fulfilling the requirement of the principle, in which case other criteria are generally invoked, particularly the trivial one mentioned a few paragraphs above concern-

[5] This easy calculation is made even easier by the reflection that there is always some unmixed act as appropriate to a given state as any mixed one.

ing two acts one of which has in each state as high an income as the other.

Examples are imperative here. Though small and artificial, they will exhibit many salient features of the general situation. First, consider the rain-shine example. If $P=14/19$, $L=70/19$. For lower P the loss is more than 70/19 in case of rain, and for higher P it is more than 70/19 in case of shine. Carrying the umbrella with probability 14/19 is therefore the minimax act in this utterly academic situation.

Turn now to a more elaborate example. Let it be given that of four numbered coins three are pennies and one is a dime, or else one is a penny and three are dimes. There are thus eight possible states s, because any of the four coins may be the singular one, and in two ways. Let an unmixed action consist in this: first, examine one of the coins chosen by number or refrain from so doing; second, guess the denomination of the singular coin. Let your income be computed on the following basis:

(i) If the singular coin is a penny you pay a tax of \$10, if it is a dime you receive a bonus of the same amount.
(ii) If you choose to examine a coin you must pay \$1 inspection cost.
(iii) If your guess was incorrect, you pay a \$4 penalty.

The first step in applying the minimax rule is to compute the loss function. Since item (i) of income does not depend on your act, it plays no role in the loss. On the other hand, if you knew s, you could act in such a way as to avoid both inspection cost and penalty; thus, insofar as you incur them, they are your loss. If you choose not to examine any coin, you might, for example, guess that the singular coin is a penny if a fair coin of your own falls heads. The expected loss for this act is easily seen to be \$2 no matter what s may be. It is almost obvious that any other act avoiding the inspection cost can for suitable s lose more. For example, both unmixed acts in this category can lose \$4.

Can you do better by examining a coin, despite the inspection cost of \$1? Not if you choose a definite coin, for this does not change the situation with respect to possible penalty, and it adds the inspection cost. If, however, you choose one of the four coins at random, and guess that the singular coin differs from it in denomination, you will be right three times in four and your expected loss will be $\$1+\frac{1}{4}\times\$2=\$1.50$, no matter what s may be. It will, therefore, be better to pay the inspection cost (at least some of the time).

In fact, the act just described is minimax, as may be seen by the following line of argument, which is of very general applicability. If the

loss be averaged with equal weights over each of the eight possible states, it can be seen with patience that no unmixed act can have a smaller average loss than $1.50. If this is true for unmixed acts it follows immediately for mixed ones. Since the average loss of any act is at least $1.50, the maximum loss of any act must be at least that great, which was to be proved. The effective weights will not in general be equal; their equality in this example was contrived for simplicity of exposition.

Little, if any, reason for adopting the minimax principle seems to have been published, though the literature suggests that some individuals find it intuitively acceptable, and demonstrates that almost all ordinary principles of statistics can, with suitable special assumptions about the loss function, be exhibited as special cases of it. This fact must not be taken too seriously as argument for the principle, since by postulating a sufficiently arbitrary loss function literally any rule of action can be seen as a special (or limiting) case of it. The principle that unbiasedness is a desirable characteristic of estimates is an example of a principle which does not follow from the minimax principle without straining the interpretation of terms, but it is noteworthy that the principle of unbiasedness in estimation is currently held in low esteem by the consensus of statisticians.

It seems clear that no categorical justification for the minimax rule can be given. Suppose, for example, that you must choose sides in an even money bet that Switzerland will become a monarchy sometime in 1951. If the concept of probability is not applicable to that event, and it is typical of the sort of event to which modern statisticians generally consider probability inapplicable, the rule says you should bet on monarchy if a fair coin falls heads.

Stimulated by the ideas of Bruno de Finetti, as expressed in [3] and [4], I would tentatively suggest the following as one possible motivation for the minimax rule. I here suppose, contrary to the opinion of most modern statisticians, that an *individual* actor faced with a choice among actions will personally find some more attractive than others. More fully, he will feel each to be equivalent to a certain cash income. Such an individual will simply choose the act which he feels is equivalent to the highest cash income. You, for example, would reject an even money bet that Switzerland will become a monarchy in 1951. (If not, please wire collect.) If, however, the actor is a group of individuals who must act in concert with a view to augmenting the income of the group, the situation is radically different, and the problems of statistics may often, if not always, be considered to be of this sort.

Thus, for example, if an experimenter claims to have evidence that a certain agronomic technique is preferable to others, this evidence should, so far as possible, look convincing to everyone concerned. Again, one might idealize at least part of the activity of academic science in terms of decisions reached jointly by the whole group of people interested and competent in the science in question.

Suppose then, that the individual s in a group evaluates the action a to be equivalent to the income $I(a, s)$. The parallelism between the notation $I(a, s)$ here and in its earlier contexts is more than formal, for the individual s may well be said to believe that the world is in such a state that the expected income of a is $I(a, s)$. If the action a should be chosen, the individual s would consider that the income $L(a, s)$ had been lost. Application of the minimax principle from this point of view means to act so that the greatest violence done to anyone's opinion shall be as small as possible. Such an action may sometimes be justifiably adopted, especially if the greatest violence resulting from it is too small for anyone to wrangle over. Where effective experimentation is a component of some of the possible actions, practical agreement may well be reached in this way. For unless two opinions are originally utterly incompatible, enough relevant evidence will bring them close together, and even if they are utterly incompatible the holder of each will feel that he has little to lose in agreeing to a sufficiently extensive fair trial. For example, you might be rather sure that a particular coin is fair and I just as sure that it will fall heads three times out of four, and we might both agree that these are the only two possibilities worthy of consideration. If the cost of observing a thousand tosses of the coin is very small compared with the stakes on which our joint action depends, we may well be satisfied to abide by an experiment based on the thousand tosses, because for each the loss is little more than the small cost of the experiment. In connection with the interpretation of the minimax principle as a principle for group action, it is relevant to mention the technical fact that in many cases, under the minimax rule the same amount will be given up by each member of the group.

There might well be some argument for preferring to minimize the average of $L(a, s)$ over the group, or for adopting some other "democratic" principle. But it should be remembered that in many applications of statistics, it may be impractical to define membership in the group, let alone to poll the members. In such a case it is generally much more feasible for the responsible statistician to treat all "reasonable" opinions by the minimax principle.

The mathematical problems posed by the minimax rule are in principle exactly the same as those posed by the theory of (zero-sum) two person social games formulated by von Neumann.[6] But studies of the minimax rule have a special mathematical tone because there are frequently infinitely many states and unmixed actions in statistical problems, whereas the analogous problem of games in which the players have an infinite number of moves is rather academic. The close parallelism between the theory of games and that of the minimax rule is very stimulating to both theories, but it has tended to direct attention away from the great conceptual differences between them.

There is a general principle for finding minimax actions when the number of states and mixed actions is finite, but it leads in general to very extensive computations and is not applicable at all when either of these numbers is infinite. Devices for solving special classes of minimax problems are therefore much sought after, and even now many of the most commonplace situations of statistics lead to difficult minimax problems.

Few, if any, new minimax solutions of immediate practical importance have yet been found. This is due partly to the difficulty just alluded to and partly to the understandable preoccupation of Wald and others working in the field with fundamental theoretical questions. On the other hand, the theory of statistical decision, and more particularly the minimax rule, has led to some profound theoretical results, especially in sequential analysis, and has greatly unified the whole field of theoretical statistics.

It is often said that the minimax principle is founded on ultra-pessimism, that it demands that the actor assume the world to be in the worst possible state. This point of view comes about because neither Wald nor other writers have clearly introduced the concept of loss as distinguished from negative income. But Wald does frequently say that in most, if not all, applications $I(a, s)$ is never positive and that it vanishes for each s if a is properly chosen, which is the condition that minus $I(a, s) = L(a, s)$. Application of the minimax rule to minus $I(a, s)$ generally, instead of to $L(a, s)$, is indeed ultra-pessimistic; no serious justification for it has ever been suggested, and it can lead to the absurd conclusion in some cases that no amount of relevant experimentation should deter the actor from behaving as though he were in complete ignorance. In the more elaborate example discussed a few

[6] von Neumann opened the subject in [8], which is still in many ways the best introduction to it; but chapters III and IV of [9] are newer and fuller.

paragraphs above, for example, if you applied the minimax principle to negative income your behavior would be not only affected but actually dominated by part (i) of the income, which is irrelevant to the loss, and you would always guess "penny" even if inspection cost were reduced to zero.

It is often pointed out that knowledge of $L(a, s)$ implies economic knowledge not often available to the working statistician. This is true, and often means that the statistician cannot apply the minimax principle or otherwise (no matter what existing theory he may adopt) advise his clients how to act. In such a case, he will generally have to recommend an action for each of several different loss functions, leaving it to his client to determine which loss function applies. Occasionally, however, it happens that some one action fits, or nearly fits, all "reasonable" loss functions. For example, in estimating the mean of a normal distribution from N observations, if the loss function depends only on the absolute difference between the true mean and the estimated mean and is bigger for bigger absolute differences, then choice of the sample mean is a minimax act.

If a new possible act is adjoined to those originally available, a new minimax act will generally arise, and it may happen that this new act, though different from the old, assigns no probability to the newly available act, i.e., does not represent the new act at all in the mixture of unmixed acts which make it up. This has been regarded as paradoxical, for why should an individual shift his preference from one act to another just because he can now act in a way which was not available before, but which in fact he does not choose at all [2]? From the point of view about the motivation of the minimax rule sketched above, there is no paradox, for when the new act is admitted the group may well change its choice to arrive at a compromise with some members who prefer the new possibility, without actually adopting the new possibility itself.

There is a natural problem of estimating the parameter of a binomial distribution for a "large" sample in which the loss as a function of P, which is the unknown state in this case, for the minimax estimate differs strikingly from that of the usual estimate, namely the success ratio. In this case, the loss is much lower for the usual estimate, except in the immediate neighborhood of $P=\frac{1}{2}$, where it is a trifle higher. Many say, and I with them, that the minimax estimate is in common sense not as good here as the usual estimate.[7] One way to analyze this dissatisfaction is to admit that in the minimax problem too many opinions have been admitted as "reasonable." In particular, an almost

[7] See especially page 190 of [6].

unshakable prejudice that the true value of P is almost precisely one-half, does not seem reasonable in the usual situations giving rise to this problem. On the other hand, a clear-cut criterion of the reasonable opinions seems quite impractical even in this simple problem, so this example seems to me to raise a real difficulty for the theory.

Wald's book presents almost all the technical information heretofore known on the theory of statistical decision, together with some new material. The treatment and many of the topics are quite advanced mathematically. The mathematical background expected of the reader is described in the Preface as "some knowledge of probability, including probability distributions in the infinite dimensional space . . ." and "a knowledge of calculus and some familiarity with the elements of set, measure, and integration theories . . .". This knowledge will not be possessed, or even be easily attainable, by the great majority of competent practicing and creative statisticians today; and, as in so many other prefaces, these nominal requirements are gross understatements, except for that purely nominal reader of unlimited intelligence and perseverance. The book is, then, directed toward a small audience today; as a progress report on a field rapidly progressing under Wald's own leadership, it will be obsolete tomorrow.

It seems to me unfortunate that Wald has avoided such extra-mathematical questions as possible motivations for adopting the minimax rule and the distinction between negative income and loss. He does mention that mathematical study of the minimax rule leads, through a technical artifice, to a certain theorem of interest in itself, irrespective of interest in the minimax principle as a rule for action; but this is not, and is not presented as, a reason for adopting it. Reading between the lines, it appears that Wald means to work with loss and not negative income. For example, on p. 124 he says that if a certain experiment is to be done, and the only thing to decide is what to do after seeing its outcome, then the cost of the experiment (which may well depend on the outcome) is irrelevant to the decision; this statement is right for loss but wrong for negative income.

I find the book difficult and irritating to read, which is especially disappointing to one who has had the pleasure of hearing Wald lecture. The technical nomenclature and notations seem unnecessarily complicated and unsuggestive. Such judgments are at best subjective, but a few examples may convey my meaning: (i) On p. 4, allusion is made to a "finite set of integers $i_1, \cdots, i_k$ which are pairwise different." In no context could "pairwise" modify the meaning of "different," and in this context "different" does not modify the meaning of "integers," so the technical-looking phrase beginning with "which" is doubly redundant. On the other hand, it is not mentioned that the set in question

is supposed to be non-vacuous. (ii) The technical term "minimal complete class" is introduced on p. 15, where a fact is immediately demonstrated which renders this term superfluous. (iii) In effect, an action a is called "uniformly" better than a', if for every s, $L(a, s) \leqq L(a', s)$, with inequality holding for at least one s. This is not at all consistent with the established mathematical meaning of "uniformly." The word "distinctly," for example, would be better. (iv) The subject matter of Chapter 2, zero-sum two person games, is such that every theorem which can be stated about it has, so to speak, a mirror image. The mirror image will generally differ from the theorem itself, but in so uniform a way that it is superfluous and therefore detrimental to state it separately, as Wald seems in several instances to agree, but of four theorems he states the mirror images in their full distracting glare of italics and special symbols. (v) Chapter 3 is seriously cluttered up with parallel treatments of discrete and absolutely continuous cases. This old source of inelegant bifurcation in statistics can always, in my experience, be unified by assuming simply that all probability measures in any given problem are absolutely continuous with respect to some one fixed measure of a certain type.[8] If that device does not work here, it would have been well worth reporting why; if it does work but was refrained from because it is relatively advanced, that was bad judgment, for any reader who had come thus far could take it in his stride. I hasten to say that I found no serious mistakes, and that blemishes such as these four do not detract a jot from Wald's achievement. Many will think it pedantic to call attention to them at all; but statisticians are human, and they understand better if allowance is made for their little human weaknesses, such as finiteness of memory, and a tendency to attach old meanings to old terms.

Wald's report on the current state of the theory of statistical decision is of great scholarly value, and its possible influence for the good on statistics, through the enthusiastic few who are able to study it, is inestimable.

NOTE ADDED IN PROOF:

Abraham Wald and his wife died in an airplane accident in India on December 13, 1950. The preceding article had gone to press shortly before. On hearing of Wald's death, I thought at first that the article ought to be revised in proof or withdrawn altogether, because it would be inappropriate in several respects to write the same thing now. But

[8] Namely, a sum of a denumerable number of disjoint finite measures.

on further thought, it seems right to let stand verbatim what was written when I believed Wald's great work on statistical decision to be but a fraction of what he was about to achieve, adding these few words isolated from the paper itself.

Wald's death gives *Statistical Decision Functions* an altogether new significance. When it appeared, it was threatened with rapid obsolescence by the activity of Wald himself. Now progress will be immeasurably slower. For a long time to come, scholars in the field will turn to this book for what is stored up in it of the guidance and stimulation which used to flow so copiously from its author.

A short obituary of Wald by Harold Hotelling, who introduced Wald to modern statistical theory, suggested much of his early research, and was responsible for his appointment at Columbia, will appear in *The American Statistician* early in 1951. It is planned that the totality of his work, which bears on many important topics in statistics in addition to these mentioned in my paper, will be reviewed in some detail by Jacob Wolfowitz, Wald's most distinguished student and later his colleague, in the *Annals of Mathematical Statistics*, probably in the first issue of 1952.

REFERENCES

[1] Bayes, Thomas, *Facsimilies of Two Papers by Bayes.* Prepared under the direction of W. Edwards Deming, Washington, The Graduate School, the Department of Agriculture. Pp. xvi, 52 numbered in facsimile.

[2] Chernoff, Hermann, *Remarks on a Rational Selection of a Decision Function,* Cowles Commission Discussion Papers, Statistics, No. 326 (Jan. 10, 1949). Unpublished.

[3] De Finetti, Bruno, "La Prévision: See Lois Logique, Ses Sources Subjectives," *Annals de l'Institute Henri Poincairé,* 7 (1947), 1–68.

[4] De Finetti, Bruno, "Recent Suggestions for the Reconciliations of Theories of Probability," *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability, 1950.* (In Press).

[5] Friedman, Milton and Savage, L. J., "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy,* 56 (1948), 279–304.

[6] Hodges, J. L., Jr. and Lehmann, E. L., "Some Problems in Minimax Point Estimation," *Annals of Mathematical Statistics,* 21 (1950), 182–97.

[7] Sprowls, R. Clay, "Statistical Decision by the Method of Minimum Risk: An Application," *Journal of the American Statistical Association,* 45 (1950), 238–48.

[8] von Neumann, J., "Zur Theorie der Gesellschaftsspiele," *Annalen Mathematische,* 100 (1928), 295–320.

[9] von Neumann, John and Morgenstern, Oskar, *Theory of Games and Economic Behavior. Second Edition.* Princeton (Princeton University Press) 1947. Pp. xviii, 641.

[10] Wald, Abraham, *Statistical Decision Functions.* New York (John Wiley and Sons, Inc.) 1950. Pp. ix, 179.